BEIJING – Tech Day
Monday, April 08, 2013 – 11:00 to 17:00
ICANN – Beijing, People's Republic of China

SPEAKER: For the transcript this would be Tech Day Session starting 11:00 am going until 5:00 pm.

[background chatter and greetings]

EBERHARD LISSE: So, good morning everybody, it's 11:00, we can start. Please keep your seats, or take your seats. My name is Eberhard Lisse, I'm the Chair of the ccNSO Technical Working Group and this session is the Tech Day, as we call it. I have some housekeeping announcements before we start. For all participants who last until lunch there will be tickets handed out at the gates and then Coremail, our last speaker, will take us to a restaurant five minutes on foot from here and we'll get authentic Chinese, as far as I understand it.

Today we have got a more security-focused Agenda. After my usual boring opening remarks, Ed Lewis will tell us a little bit about threats to security, but more from a conceptual basis, not real attacks, but what is there that needs to be considered. Then Microsoft will introduce… Speak about their Microsoft Registry Scanning Service – and you're more than welcome to give them a bit of a hard time, just for the fun of it.

[laughter] Then Patrick Fältström should be here to give us a briefing on the SSAC.

Then we have Coremail who will talk about their IDN email project and then as they are the lunch sponsors we usually put them before lunch so that it is sort of [flowing singles? 00:03:28] also fair to them. In the afternoon we'll have Tom [Schepatzi? 00:03:38] talk about Escrow. We had Iron Mountain here already so this is the other big player on that market, so it's only fair to give both of them a choice. But I have asked him to do it a little bit more technical so we can figure out how they actually do this. Maybe we'll get some ideas from that.

Then Diego Espinoza is going to speak twice. He's giving a long presentation, fifteen minutes each about his two topics. Then we have Jiankang who will give the host presentation as usual he will also speak about IDN email because that's their big project here, obviously. Frederico Neves will then tell us how they dealt with threats that they had to their Registry. Jaap Akkerhuis will present some research down by some students that looked at rate limiting and [LLAP? 00:04:39].

And then we're going to have a round table, from MarkMonitor, but one of the participants couldn't arrive so McAlister is going to give us a presentation about what MarkMonitor has noticed. And then Jay Daley will be… If he's around he will give his usual closing remarks. Yeah? Then as usual we move to the DNSSEC for everybody. The meeting will be in another room. Whoever wants to go there should be aware that it is there and that we can go and move the venue.

We usually do this that we give a presentation and then we have a little bit of time. I would like it if the speakers could adhere to the tables. We can finish a little bit early before lunch but we can't really finish much later because we really don't have much time for lunch so try to keep within your time limits. There will be a little bit of time after each presentation for questions; we've got two microphones and they're really small, and then we will hover around and take your questions. Ed Lewis.

EDWARD LEWIS: Good morning. We're looking for the slides. I think there are other presentations set up. So I'll start on my talk without the slides, I'll see what I remember from my memory. [clears throat] So about a year ago I went to ccNSO Tech Day and I saw a lot of presentations that covered security topics in a lot of detail, and I thought it would be a good idea to have a presentation which said 'let's look at security as a whole' and try to give you a place where you can take the further presentations you'll see later today and put them into 'this belongs here', 'this belongs there', 'this belongs there'. And it lets you also think about areas that you might not have considered trying to secure for your Registry.

So that's the background of this slide. I'm trying to just give you a framework for how to think about security and then go in later on with more details from… When you see the next presentation come along. [background chatter] [laughs] A few people have already given me comments on my slides already. Any questions so far? [background chatter] The only copy I brought with me is on paper. I didn't bring a

computer with me. There we are. Okay. Okay. So the goal here is to give an overview of security – next slide I guess.

So what is security? When I start working on slides I try to think about what do I mean by security? Security is part of availability, reliability, which are very important to any operation of anything in the world, it doesn't matter what your topic is. Availability means being up. Can you be available? Can people get to you at all times? Security means not being taken down, usually for malicious activity or malicious-like activity. So they work hand-in-hand, but availability is more of a positive way of putting more things are out there. Security is making sure things don't go down on you.

When you think about security it starts blending in with other issues out there. And it's hard sometimes to draw a line between is this a security issue or is it just a capacity issue? So I'm going to try to stick two things up that look like malicious activity. Next slide. What does it do? It tries to limit damage to malicious-like activity. And I say 'malicious-like' because sometimes mistakes look like they are threats. If you mistype a password, that can look like a security threat, so it's very fine line. It doesn't matter if it's malicious or not; if it looks malicious we try to defend against that.

We can never, ever, try to prevent an attack. I hear people talking about trying to prevent attacks. You don't control what other people do. So you're going to get attacked or mistakes sent to you no matter what. It's also… Security is not absolute. There's no level of security you need to be at, it's a matter of how much you want to risk, how much can you afford to risk. Next slide. Okay, we're going to start. When you start

with security the first thing you should think about is what do you want to protect – what's the most important thing you have and what sort of things are important, what things aren't so important.

But keep in mind what it is that's the main purpose of your business, the main purpose of your organization. And also how much can you afford to defend it? There's always a limit to how much you can afford. The first thing to do it analyze the operation and take a look at what your business does and start to understand how you break it down. Define what's normal because 'normal' is something which is very much underlooked in almost everyone's operation. We don't really know what we normally have running.

You determine what activity you'd consider to be a risk, what do you think is a threat to your operation and you have to at least monitor it for that, if not make sure it doesn't get attacked that way. And finally you want to automate your response. You want to be able to prepare for these to come in and be able to clean up whatever mess might happen. Where to stop? It's good to know where to stop in everything you do. Security is something that you really need to have; it's not something that you want to have. So you want to refrain from trying to be very [phased? 00:12:28] of security. You should do just enough to keep yourself safe.

You have to prevent security from stopping your network from being used correctly or your operation to go correctly, and you don't want it to become a burden – and that's actually the hardest part. If it's a burden for people who are legitimate users it falls apart quickly because they're going to circumvent it and everyone circumvents it and you're back to

nothing. You really want to have a goal ahead of time for whatever you do, to say 'here's what I want to get up to' and so you know when you get there you can stop worrying about that part and probably work on something else that you have in front of you.

So now, Domain Name Registries is what I've targeting this talk towards, given some of those observations about security. A Registry's role is to associate an object with an entity. It doesn't matter what kind of a Registry it is out there but you have objects; things that you manage. In our case, Domain Names. And they all belong to somebody. They're registered to somebody, the responsible person. That's the important mapping that we have. We want to be able to perform that mapping in a very quick timeframe – I want to be able to tell you instantly. In fact I call them DNS time – I want to be able to answer the DNS protocol information about that Domain Name; where it is and who has it and such.

What also is important is enforcing some policies I have, and the policies change Registry to Registry; there's no commonality there. Some people said RIRs and Domain Name Registries are very different. In the sense of what I'm breaking it down to here, they're really similar. [inaudible 00:14:05] do have some other factors but for the most part, and seeing as we're talking about them anyway, they pretty much fall into the same categories we have here. Next.

So, Domain Name Registry. The most important thing that most of us have to think about here is the database. It's the number one thing out there. If you were to go bad, if you had trouble, someone wanted to take away your operations; they'd eat the database. The database has

the mapping of who has what. So you want to protect that most of all. That's got everything in it that you need to have. It's got to be available for whatever emergency might come up. Around the database you have a bunch of services that you need to keep in mind.

There's what I call provisioning side, which is the input to the database, where you have Registrars and you've got Registrants coming saying 'I want to put this information into the Registry – I want you to map things in the Registry for me.' I have regulars coming to me saying 'I want to do certain things in a certain way in there'. But I also have the export side of the Registry, which is how do I report these things? And for example DNS is going to be the common theme here, WHOIS is another one out there and there's a few other ways to report what's in the Registry.

You have to understand as an operator how all those things behave in normal situations. How does a database keep track of things? How does it get information in? How does it pass it back out? And until you understand that you can't even begin to try to secure it. Now, I'm going to talk through the rest of the slides. Specifics change from Registry to Registry. I'm trying to say this very high-level so that it applies to just about everything out there, but of course if I say something it may differ from your situation to somebody else's. Next.

So this is a cartoon showing the high-level view of a Registry. You've got a database on top of customers, which are the people who are your Registrants; they're the ones that are paying you to take care of data for them. They may be paying you or maybe it's free, I don't know. A regulator is on one side saying 'this is how I want the operations to

occur' and then you have the Internet on the other side, which is the general populous out there that wants to make use of the information.

Provisioning services, you break it down a bit. What I have to show here is a registration interface, which is how do you actually talk to your customers, how to get information, how do they register their names, is it EPP or such? There's also a website you'll probably have for your business, which is probably just a run of the mill type of website portal and then billing is another interface you have to customers, which is usually overlooked. The database takes a [fridge? 00:16:38] from all of that and puts that all together.

And for Reporting Services the database spits that information to the WHOIS servers and the DNS out there with SSAC. And there are other things that come out of here too. I was thinking of just these right now, and that goes out to the Internet users. So as far as basic security, I'm not going to go into this level of detail, but in order to run anything you have to have some basic ideas of what you want to have. You want to have locked doors, locked windows. You want to have physical access to your securities, monitored and controlled. You want to also have a financial security and make sure that your operation is going to be there, available to answer questions in the future.

Personnel security is important to make sure that the people you hire are going to be acting in a responsible way, that you have ways to make sure they don't run amuck and do things bad to you. Information technology security out there, host security. What do I mean by host security? You have to have routers, firewalls, [formally? 00:17:37] operating system and so on. That's just basic stuff that applies to just

about every operation out there. And also to make sure that it works, security audits. Make sure you go through, make sure things are actually running securely and also testing your security to see if it actually is holding up.

Now, when all else fails there's Escrow. And I added the slide after seeing the Agenda today. Escrow is the last safety belt for an organization. Basically it's a way to keep track of your database somewhere else so that if you have a complete catastrophe you can still get someone… Someone can reconstitute the Registry somewhere else. You have to make sure that you provide information; you want to make sure that it can be tested, but you hope to never use it. Next.

Provisioning services, to give you an idea of what I'm talking about here, EPP for example is a Registration interface that's used by many people. It's not universal; it's just an example. I think most people have heard of EPP. Actually, by a show of hands, who here has heard of EPP? Right. Who has never heard of EPP? Okay. It's a particular protocol for requesting yourself a Registry. There's a general information website that may give people information about your Registry. You might have a portal there; you may just have advertisement type information there about how to register for a Domain Name.

And then billing, which is something that we don't consider much in engineering. In fact, when I went through and did some background work on making these slides I asked people I work with about the billing system and they looked at each other like 'we don't do billing. We're engineers. That's another part of the company'. So a lot of times we don't even think about that. Vulnerabilities that we faced on our service

in trying to register Names is important. You have to watch out for. Or hogging, as I call it here. You want to make sure that one customer can't prevent other customers from getting to stuff. You have to register Domains legitimately or to get new Names.

Many of us aware that when Names become unregistered there's a rush to come back and get those Names. You have some competition between some customers to get there. You want to be fair about that. Poorly formatted data is another big concern. Things called SQL injection attacks have happened where when I ask for information from you off the network, if you give me something that's coded to make my systems fail, it's bad for me. So I have to make sure that any input is checked and make sure that I don't execute what someone tells me to execute at the wrong layer.

And then corrupt data is a whole other category, where you may have stolen credentials; people steal money, people steal ids, steal passwords and then take legitimate Domains and change the name service, for example, to go somewhere else; fraudulent registrations. So techniques that are involved with the front-end. Web security – I didn't mention that in the first slide, I probably could have, but the website has to be made secure through the way you would make any kind of website secure.

For the registering protocol we use traffic shaping, which means that when you… If you have a limited number of customers, if you only have prearranged relationships with your customers you reserve so many resources to handle each one. And you can only have so much. We also restrict the addresses they come from, because at least that help pin

down where they're coming from, and anything else that can be used to really prevent either just them shoving data into you they don't deserve to shove to you, or make sure you accept it only from the right people.

Poorly formatted data is usually defended by better software, basically. If you write software correctly you shouldn't be suffering an SQL injection attack. You should know not to execute what's inside someone's string. And for corrupt data, transaction security is needed and it's also good to have a way to take down Domain Names that are malicious. In the first case you may or may not have a long supply change, you may not talk directly to a Registrar; they may go through Registrars, host sellers and so on, or you may go directly through a Registrar.

So the level at which you can check the Registrant's accuracy changes from place to place and you need to do as much as you can in that area. But also, once you've accepted registrations, you should be able to take reports from people saying this name out here is being abusive. I want you to take it out of a Registry. And different Registries have different obstacles to accomplishing this. Some are easier because they have a closer tie with the law enforcement agencies; they have a clearer chain of custody, of evidence, whatever to take something down. Some places don't have that comfort.

Now, for billing, if you take credit cards there's a whole lot about credit card security out there. Someone may attack you just to get the credit card information that's happened around – maybe not in a Registry as much but it has happened in Congress. And also, any kind of account information you have has got to be protected because if I can steal an

account from somebody else, I can start making registrations at will in the Registry and using other Names and so on and not having it traced back to me, trying to do some bad stuff.

For the internal systems the database, which has everything you have, the database can be one thing or it can be many databases. Some places will throw everything into one huge database, some will use a few different databases. It might have contact information, it might have credentials, passwords that might be in there. And then also there's a rules enforcement engine that may or may not be a database or may just be… It could be whatever it has to be implemented in your Registry.

The threats here beyond fraudulent data in a database, you want to make sure that the database itself is structured appropriately, you want to make sure that information is given out only when it needs to get out. And I'm not talking about to the DNS or to the whomever; I'm talking about inside people. If you have operators who have access to the database, make sure they can only see parts of the database they need to get access to. They can only make authorized changes or changes are logged so you know who made a change and when they made a change.

You want to limit access to the database, especially the credential areas, especially the billing information, make sure that's tightly held so no one can walk around with that information. You certainly don't want to have it sitting on a laptop that you carry to another building or country. You want to limit damage from social engineering too. That means making sure that people who call you on the phone can't just get favors done by people working inside, for one reason or another.

The rules. There's not really a whole lot for me to talk about for rules, as far as I can see. You want to make sure the rules are being properly followed; you want to make sure the rule engines are available and functioning and you're working. And make sure you work with the regulators and the rules are sensible and they have the desired effect on the system. Next. Now, for reporting service, this is generally the area where we start thinking about security, it's the most visible part of a database.

Customers generally don't do a lot of abuse to the Registries; it's the people out there who are relying on information being held in the database. WHOIS is an example of one service in there. The DNS and the DNSSEC is something I'm going to talk a little bit about too. Next. So the WHOIS threats. When I asked around about WHOIS threats, most people said they don't even concern WHOIS as something that gets attacked. But that comes from places where there's no restriction on access to the Domain Name Registry information. You can give out bulk access to the Registry.

Some Registries do not want to let everyone know all the Registrant information, some are not concerned about that; it depends on their contractual obligations. WHOIS comes down to basically a TCP service and TCP is a protocol that's been attacked and defended for quite a while now. We're getting pretty good at being able to throw off TCP based attacks. Data mining is the other concern; if you're worried about people seeing data they shouldn't see.

But in general it's not a target. In some cases, even if you allow bulk access, you may also want to make sure you do not get annoyed by one

source. I know that even in some places where the information is totally free to be given out, if one IP address keeps hitting them all the time they will squelch that IP address or stop its… Target that IP address. The defenses for WHOIS; make sure TCP works well. It's usually an OS level defense. You want to have general availability goals, which gets into capacity planning more than security. Make sure you have enough WHOIS servers out there.

When it comes down to it though, I think for anyone who really needs WHOIS to be up and running, they're going to contact you offline anyway. WHOIS is kind of a convenience, but when it really matters, in situations where they need to have a legal statement about who owns what, it's going to come through other channels anyway. For data mining you want to have a bulk access agreement saying what you use the information for, so that they can't just grab it and can't spamming people because they got the address for free.

Capture on the UI to slow down people from asking over and over again and monitoring and filing the requests that come in. DNS threats are and always have been popular targets. It's changing all the time. DNS is becoming even more of a target these days, actually in the last year and a half or so. There's a [nine? 00:27:45] service, which means taking down servers as a [kill-packet? 00:27:45]. If I send certain kind of queries to some Name servers, they could make the process running just die. That's less likely today than it was ten, twenty years ago. Software has gotten better, but you still have to be aware that you want to make sure that your servers are able to stay up and running.

Packet flood attacks, nowadays, DDoS gets a lot of attention. It's become a buzzword about what DDoS is. There are two ways of looking at DDoS for Registry. One is, Registry could be the victim of a DDoS attack; meaning I'm going to flood all these packets at the Registry itself, at the DNS servers, or it could be that the Registry is an unwitting accomplice in a DDoS, in that an attacker will send packets to the Registry – not meant to knock over the Registry but to have the Registry respond somewhere else.

Cash poisoning has been a threat in DNS for a long time. DNSSEC came along to try and stop that. And data disclosure is being able to go through all the information in the DNS zone as a result of DNSSEC. Go to the next slide. Now, the one thing I wanted to spend a little more time on is the idea of reflection-application attacks. Now, these are attacks that make the Registry no longer just the victim but also an unwitting accomplice in an attack. What these attacks refer to is an attacker, that I have there on the left-side of the screen, sends queries into the DNS server, but the return address is not the attackers address, it's the address of the victim.

Now, when the Registry sees these packets come in, it doesn't know it's an attacker, it thinks it really is the victim asking for some information, so it will respond back to the victim over to the right-side of the screen. Not just the reflection of packets from A to B, there's also an application size. The query in DNS is very small, the response is very big, so the victim gets a lot of traffic from the Registries. So the Registries in this case are actually helping to pump up this attack. In fact, the strategy of trying to have more capacity to survive a DDoS backfires here because

the more capacity that you have, the more attack traffic you can send to the real victim out there.

So you've gone from being the victim to an unwitting accomplice, which is a very significant change in your position in the attack tree. So, it's no longer the question 'can my systems withstand the attack? Do I need more capacity for this?' It's 'what am I doing to others around the world? I have to make sure that I'm not sending out bad traffic everywhere'. What can be done right now is Response Rate Limiting, which will come up later today too, which is a fairly new implementation of throttling requests, throttling responses based on what looks like malicious activity out there.

And Response Rate Limiting itself, if somebody RRL… One thing I should be clear of, it's not Resource Rate Limiting, it's not Request Rate Limiting, it's Response Rate Limiting – it's part of the technique and Jaap will talk about that later in detail. This particular solution is already implemented in at least three sets of software out there that's distributable and others too, but BIND has it, NSD has it and Knot has it out there; I believe those are the right terms. I'm [inaudible 00:31:15] Unbound has it or not.

Next. DS defenses, host security, make sure it's up and running. Make sure the DNS itself limits its share of fate across the board. You want to have as many independent sources so it's hard to take one down and see them all come down. Any [cas? 00:31:37] technique to isolate attacks so it helps break up your network into smaller chunks and then rate limiting response is a good thing to do and then DNSSEC is the last topic I'll go into. DNSSEC considerations. I don't have threats here; I'm

changing the tone of this a little bit. DNSSEC considerations come up now.

The first thing that comes with DNSSEC is key management. That's a whole new topic for many Registry operators out there, having to handle cryptography. Cryptography is not really well understood around the world and I've spent 15, 20 years trying to get this working with DNSSEC and it's still not clear to me, and there are no real hard-sell facts about cryptography as far as I'm concerned, out there. So it's something that takes some time to come up to speed on. The DNSSEC private-key material, without going into what all those terms mean, has got to be kept a secret. That's what DNSSEC depends on; the secret key is a secret.

There are different ways it could no longer become a secret, and you have to watch out for that; it's a poorly made key, somebody takes the key outside the company with them or someone else outside just starts guessing what the key could be. And that's stuff that you can't control. But you have to look to realize 'what am I going to do about the chance that someone will discover my secret key and it'll all falls apart?' Signature generation process is also a concern because even if I'm signing with my right key and it's really a secret out there, I have to make sure the data I get to sign is legitimate.

DNSSEC doesn't know if that's the right answer or not, it's just what it's told. And that goes back to the database. So if the database is giving bad information, DNSSEC will sign bad information. Some techniques out there, there's use of [ANSAC III or ENSEC? 00:33:23]. This refers to the situation where if you're a Registry and you don't want to let people

know who's registered in bulk, [ANSAC III?] is going to help give you that ability. You cannot go through the zone with ANSAC III to see what's in it.

[ENSEC?] on the other hand, allows people to potentially go through the zone and ask what's each and every name inside the zone. I will go through why there's a debate between the two of them, but that's the difference between ANSAC III and ENSEC. Choosing parameters well is going to be very important. I've studied parameters choices by operators for quite a while now and what I find to this day is many operators out there, the TLDs, tend to follow a herd. We pick popular numbers. I don't think there's a lot known about these numbers because I think cryptography still isn't really well known how strong you need to be. There's still some room for experimentation in there.

But the thing to keep in mind is that if you try to take on too much security, it's a burden. You do too much, it's too hard to do. If you do too little security it's either forgotten or it's going to be broken. Right now, today, I don't know of any attack anywhere that's trying to break DNSSEC and succeeded. Now, it doesn't mean DNSSEC is the greatest thing out there, it means we don't know when DNSSEC would break. Now, I'd like to know if it broke somewhere, I have to be above that. We don't even have that information now. I don't think anyone has that information.

The other question that keeps coming up is whether or not to use an HSM and I'm not going to give you a yes or no answer but the first thing I observe about HSMs, which are Hardware Security Modules, which help protect the private-key really well, is what's really more important that

your private-key though is our database. So, how protected is your database? How well protected should you have your key? HSMs will help you protect it from one level more, but you have to decide if it's necessary.

The downside of HSMs is they complicate high availability. Now, when you start designing a high availability system with multiple replications, whatever, HSMs have to be architected correctly or else they can actually hinder your fallback plans, and that's actually happened in the past. DNSSEC and amplification. DNSSEC is not alone in making larger packets, larger responses. IPv6 is out there and a few other applications out there that want to have larger responses.

I have larger [ENEX? 00:35:56] Domain responses. That actually is the fault of DNSSEC. DNSSEC will make things much bigger on the way out and that where the Response Rate Limiting comes into effect. [ENEX?] Domain responses. What can be done? You could ignore DNSSEC and just be insecure, which is probably not where you'd be encouraged to go, put it that way. There are other ways in DNSSEC that have tried to limit the size of the responses. You might choose different key sizes that make it more efficient in terms of the bits per message out there.

You might be more efficient on how many records you put out there. Fewer records sometimes are better than more records, makes more responses. And the Response Rate Limiting here is really, probably an important factor at this point. So in conclusion, there are a lot of attack services in a Registry, there are a lot of techniques out there to solve this. You're going to hear talks that talk about solutions to a lot of different particular problems.

It's helpful to remember where that fits in the architecture of your system back at home. When you hear these presentations and they're new to you and you're learning this information, think about what it is you're running back at home, what do you think is normal there. And if you don't know that, that's the first place to start. What's my normal? And then see how these different presentations are going to help build up some part of your defenses out there, or other parts of your operations out there.

You don't want to have too little security, because if you have too little, you'll end up panicking when something happens. And something in one of my talks I've prepared later in the week, if you're not prepared ahead of time, all you can do is panic. You have to prepare ahead of time for this stuff. If you have too much security you can also stop the usefulness of your Registry. You don't want to make it so hard that no one can use it; you want to make it available for people to get information out of it. With that I'll open up for questions.

EBERHARD LISSE: Any questions? Come on. [Kristina/Christina? 00:37:59], anything from the [demote? 00:37:58] side? There's one in the back. And everybody speaking should identify themselves for the remote participants.

JASON POLIS: Hi, I'm Jason Polis from Super.Name. A question with regards to the amplification attacks. Is that only over UDP or does it also happen over TCP and STCP?

EDWARD LEWIS: We've only ever seen it over UDP and I think it's… It only really will work over UDP because it requires the sender to be able to send me a packet that I will then send somewhere else. In TCP, if you try to contact me you and I have to exchange some packets before I send stuff back to you. So that means that in the TCP world, if you sent something to me, before I send anything back to you of any size, you've got to tell me some information which tells me I'm talking to the right person. And it's not at a security level; this is actually just the functioning of the stream protocol. UDP doesn't have that back and forth negotiation so it's a place where you can do this stuff at will.

And just a follow-up too; one of the other factors in the amplifications attacks is something called BCP 38, which is something that Registries can do nothing about except ask ISPs to implement, because that says 'don't send me packets with false addresses'. And there's not a lot we can do with that so I didn't mention it in my slides.

EBERHARD LISSE: Yes?

JAY DALEY: Hi, I'm Jay Daley from .nz. Thanks Ed, that was very useful. Can I ask how this is documented and enforced within your own organization?

EDWARD LEWIS: I don't know. I don't know that there is… And 'I don't know' meaning I don't know, not like I don't think there is. We have a set of documents about security and I think when we were applying to some of the new applications they've been answering questions as ICANN ask questions. This is probably as close as we have. I don't think that we have a comprehensive security document out there. I don't think anyone has taken the time to sit down and say 'let's put together a White Paper on how we secure everything', in the sense that that would be a lot of extra work for what's happening practically. That sounds like a non-weasel answer . [laughs]

EBERHARD LISSE: Sure.

CHRISTIAN HESSELMAN: Hi, I'm Christian, Christian Hesselman. I'm with .nl. My question is you mentioned that the BCP 38 can be used to basically fight off these amplification attacks. What are your thoughts on persuading ISPs to enable or to turn on BCP 38?

EDWARD LEWIS: Okay. I saw a reaction from someone else in the room when you were saying that. I don't think BCP 38 alone is the answer. It's a first step. That prevents the ability to do the reflection part, saying you're somewhere else. Amplification attacks wouldn't be stopped… Sorry, because an amplification attack is always going to be that it's going to

give out a bigger response than I'm going to have. In terms of doing work with ISPs, I would have to say that BCP 38 has been someone else's battle cry for at least ten years and it hasn't been terribly successfully.

My opinion is that I'm a Registry person and I have certain things, certain arsenal or defenses. Asking someone or somebody else and they don't do it, I can't worry about that. So like I turn around and say 'what can I do to survive in the face of people not doing what they need to do for defenses?' So I don't… Unless the Registry has some leverage over ISPs, which may be in some jurisdictions, I would say encouraging BCP 38 has got to be a good thing, but if you don't have any leverage you have to learn to live with the fact that that's not being done.

And the fact is that even if I had a ccTLD in one country, where I could say 'all you ISPs do what I say'. There are ISPs who are not in your country who will not do BCP anyway, so it's not really going to get you anywhere. That's the harsh reality of that. One thing I'll throw out there too is in looking at a lot of the major attacks one there, they tend to be international, for legal reasons. And I'm not going to say more than that because I'm not qualified to speak that deeply on it, but I've been given public information that generally if there's an activity that's malicious and more country's resources are used – different ISPs, servers – the harder it is for anyone to track that person down.

So most likely when you're getting an attack, you're probably going to have someone from outside your jurisdiction attacking also. So that's why I'm saying…

EBERHARD LISSE: All right. Thank you very much. That was a very interesting presentation. I think what's even better is that we have got some detailed slides, so we can refer to later because there was some issues that I found interesting but I can't place them anymore, so the PDF will be very helpful. Next will be Ms. Kern, Cynthia Kern if I'm not mistaken, from Microsoft. She will speak about the Registry Scanning Service, which as generated a little bit of attention recently and correspondence between the ccNSO and the ICANN Chair.

CYNTHIA KERN: So thank you. Actually, I'm joined by my colleague Nick Whitworth, who is also in the same organization. We are an online services division in Microsoft. I'm a Program Manager and he is our Domains Business Manager. He manages our full corporate Domain Name portfolio. And he's going to talk a little bit about the process for getting set up on the Scanning Service in just a bit. And I guess I will lead with we are not security experts so we're doing our best effort today to represent this service offering that Microsoft has provided recently and in early March we announced the offering. But there may be some questions that we can't feel today and we will do our best to do so and take back what we can't answer.

Really quick, in the room, can I just see a show of hands for anyone that's a Registry operator? Okay, good. So, you can move forward. So to be clear, our target audience is definitely Registry operators, ccTLD operators. Microsoft announced early in March that we are offering a Security Assessment offering and so we'll talk a little bit about that, but the basis for doing so is really that for several years now we've seen a

number of Registry hacks, ccTLD Registry hacks. And Edward did a nice job earlier describing a lot of the vulnerabilities that are out there, and I think that's probably all familiar to the room.

But in general we're talking about this service in a way so we could potentially offer it to help mitigate and early detection on some of that. So, increasing problem, right? Since November there have been, I think, approximately 12 or 13 hacks and they just are continuing to happen and increase in their frequency. And so a call to action on our part of what we could do to help you understand what the findings are and consult with you on how to mitigate or mediate those risks.

We are not taking responsibility for actually doing that for you; it's a consultation at that back-end. But it is a free service and we are excited to offer it to the ccTLD Registries. So Nick, if you want to talk a little bit about how folks can get enroll, and the scanning?

NICK WHITWORTH: So to actually get enrolled it's pretty easy. Initially just send an email to cctldregsec@microsoft.com. We just ask for very basic information up front; basically we just request your name and email contact and where you're coming from, as well as an approver name and their email and contact information that can authorize us scanning their system or your system. Also provide us with the IP and the URL for the public-facing portals that your Registers use.

Once we have that information initially from you we will reply to you, giving you just a little bit more information and also providing a Terms of Use document for you just to agree to. Once we get all that back, within

seven days after that we'll perform the scan. The initial security scans will include a host and a Web application scan. And after that we'll do monthly host scans and then cordially web application security scans, basically until you tell us to stop. You can go to the next slide.

This is just an example of the report you'll get back. It's fairly basis but can have a lot of information if there are issues that we find. We go on a five to one scale – five being high and a serious issue and one being low. The report will also just tell you all the threats we identified, what the impact could be and also give you ways to mitigate that issue, basically a solution for you. As Cynthia said, we won't actually provide that solution but we will tell you how you can go about fixing those issues. And that will also just come in a PDF format that you can get quick and easy. And that's basically it. It's pretty simple.

CYNTHIA KERN: Awesome. Yeah, so just in closing, greater good really, it's about benefiting everyone and what we can do, especially in the international, smaller markets, where we can simply provide this out as a quick scan; it can obviously improve the online experience for all and I think in general, hopefully the room shares the sentiment that it's a shared responsibility, right? Security online. And so we are hoping that the Registry operator here and those who are not able to attend will leverage this as an option as they look at their system, their network, and try to improve that security and stability.

NICK WHITWORTH: Just one other thing, really quick, is that if you want to email cctldregsec@microsoft.com, just for your information feel free to do that as well and the Security Team can probably answer any questions you might have.

CYNTHIA KERN: But based on what you said there could be some questions in the room. Yeah?

EBERHARD LISSE: Boom.

CYNTHIA KERN: [laughs] Boom. We'll do our best.

EBERHARD LISSE: This is after [Alan? 00:52:04] Technical Working Group. Let's start at the front. Let's start with [Jack / Jacque? 00:52:07], he hasn't said anything yet.

JEAN-JACQUE SUBRANET[?]: [extreme audio interference 00:52:10 until 00:53:11].

UM: …We sat through Ed Lewis' presentation earlier and saw the breadth of security issues that a Registry looks at. Would you think that describing your service as a Registry Security Assessment Service is perhaps

stretching a little bit, when the elements that you are looking at protecting are maybe 2% of the overall set of things that a Registry needs to look at from a security posture?

CYNTHIA KERN: Possibly. I mean they scan for a lot but yeah, I mean, we had to name it something. I mean it's… Possibly.

UM: Because that's my concern, that one of the things that we need to do in our community is get the best practice out to as much as possible, and my concern would be that some Registries would look at your service and it's Microsoft after all – huge credibility there – and be able to turn around to their government or other people and say we've had our Registry scanned by the Microsoft Registry Security Service, where's the problem? And the problem is the other 98% of things that you don't look at, that aren't in any way covered in that, of course?

CYNTHIA KERN: Right. I think that's actually a really good point you're bringing up and I think I might take it back to the team who's doing the scanning on our end, of it's probably a good idea for us to include a disclaimer at some level, or here's in each scan we do around here is what's in scope. Yeah, I think that's probably good feedback. Thank you.

ROBERT MARTIN: Robert Martin from Packet Clearing House. The part of the report that I could see that you've generated looked a little bit familiar. Is it based on Nessus or something?

CYNTHIA KERN: Based on…? I'm sorry?

ROBERT MARTIN: Based on Nessus?

CYNTHIA KERN: I don't think so but…

ROBERT MARTIN: Is it something you've developed yourself?

CYNTHIA KERN: No, sorry I think it's a combined internal and also a third party that we work with. It's not Nessus though.

EBERHARD LISSE: Christina, any of the [inaudible 00:55:14]? Okay. Then I think we'll go to the next presentation.

CYNTHIA KERN: Thank you. Thank you for the opportunity. [applause] [background chatter]

PATRICK FÄLTSTRÖM: Fun with static electricity, right? [beeping] No, it's me touching the microphone. So my name is Patrick Fältström and I'm Chair of SSAC. I also have with me Jim Galvin to my left, Vice Chair of SSAC, and then Julie Hedlund, ICANN Staff, that is one of the Support Staff we have. We want to give you an update on our activities. So we'll give you an update on our activities. We'll go through basically three different things. We'll briefly explain a little bit what SSAC is and what we have done and then we'll go through… We were thinking of two different reports. We'll see whether we have time to go through both, it's the third time we've tried to do this here at ICANN since yesterday morning. At the two other meetings, that were also half an hour, we completely failed going through both reports. They also ended up being one, so we'll see how it goes.

Next slide please. Next slide. So, SSAC, the Security and Stability Advisory Committee here at ICANN was initiated in 2001, began operation in 2002 and we provide guidance to the ICANN community and the Board on any kind of matters relating to security and integrity of the Internet's naming and address allocation system. We are currently 38 Members and several of the Members are here and several of the Members are also Members of the ccNSO and some of them came into the room just for this and some… [laughs] Some people came into the room just for this but some other people I know will be here during the whole Tech Day.

Can the people that are SSAC Members raise their hands in the audience, just for people to see? So there you are. Look around, see…

Yeah. Next please. We have a couple of Committees and Working Groups, first of all, that deal with more administrative stuff, a Membership Committee, a special group that work on the DNSSEC Workshop Program that we have Wednesday morning for those of you that want to talk more about DNSSEC-related issues. We'll look at… Actually quite interesting this year, just like last year, I always say that, but actually it looks interesting every year.

We also participate following the DSSA Working Group here in ICANN and then we have SSAC Members that also participate in other kinds of Working Groups here inside ICANN. Next slide please. The more interesting thing may be for the substantial outcome of our work is our Work Parties. The way we work in SSAC is that we identify and issue either by finding the issue ourselves or we get a question from any other body inside ICANN or outside ICANN, 'can SSAC?', 'What do SSAC think about this?' and the question can come from anyone.

The Advisory Committee, we [earn the right? 00:59:35] to ICANN's Board but we have got questions from the Government Advisory Committee and some other bodies and also from ICANN Staff earlier. What we do when we work with those questions, those issues that we find is we create Work Parties. And the Work Parties that currently are active, are more or less active, are the ones that see on the screen. Some of them have not really started yet but it's on its way to start.

That we have Work Parties working on things, doesn't mean that they will actually release a report, because it might be the case that some of those Work Parties are just responses to letters, some of them actually do substantial work but they come to the conclusion that no, this was

actually just a rumor; there was nothing to see here. Move on. So that you see things on the list here doesn't mean that you will see a report. Next please.

When the Work Parties do come to a conclusion that we need to write a report, a report is created. We are creating between four and six, or nowadays it might be six reports a year, and you can see here on the list a number of different reports in the category of social security and abuse, Internationalized Domain Names – next slide please – and WHOIS. So these are sort of the issues that we have been working with lately.

To reference the previous presentations; we do have three different reports that are a little bit older from a couple of years ago, that has to do with the security recommendations, the security and stability recommendations for Registries, Registrars and Registrants. Three different reports. And that might be something that could be interesting for you all to have a look at if you're interested in those kinds of issues. Next slide please.

So let me pause there and see whether there is any direct questions on SSAC operation, because otherwise we are diving into the first report immediately. Going, going, gone. Good. So, SSAC 57. So what happened was that we had this meeting with SSAC Members and we were sitting in this room and we came to the end of the day, we only had dinner and went to the bar, or to bed depending on what jet lag you had, and one person said 'can I get some time during any other business please?' Sure.

The person had a presentation and he could literally hear all of SSAC just saying 'this was pretty bad'. So he immediately decided to start the Work Party to work on this issue, and let me try to explain what this is about. So this is a self-initiated work. Next slide please. Internal Name Certificates is something that we call certificates that are given out by CAs that have an identifying string in them, where we normally have a Domain Name, and the string of course have a syntax of a Domain Name but the trailing token is sort of a TLD that is not allocated. Okay? And what the advisor is talking about is that historically it has been not allocated and now it's not yet allocated.

And the question is, what happens if it is the case that those strings, that have been used for Internal Name Certificates, actually, suddenly ends up being allocated as TLDs? And the finding is basically that those certificates, which are internal certificates, suddenly ends up being certificates that look like exactly the ones that we're using on the public Internet. There are no other bits of identification certificates. So if it is the case that a CA have a practice that they give out Internal Name Certificates today, for TLDs that for example are part of the New gTLD process, someone can – and we did some tests; and yes, absolutely, there's not problem whatsoever to buy a certificate today, for some of the TLDs that are applied for.

There's of course no Domain Name validation possible to do those, instead you just get the certificate if you pay. Now, you just sit there with a certificate, you wait until the TLD is actually live and then you pop up your own server somewhere, where you can do a very beautiful man-in-the-middle attack, and the people using your services will absolutely

believe that they're using the correct services from the Domain Name Holder in that Top-Level Domain. So, this of course could impact the New gTLD Program and the SSAC…

What SSAC did was we wrote a report that we handed over to ICANN around the end of 2012, beginning of 2013, where we advised ICANN that ICANN should take immediate steps to mitigate the risks. Normally our reports go to the Board and our recommendations to anyone, but to the Board. This time, just because we decided to use a different method, we decided to give this report directly to Geoff, as the Head of the ICANN Security Team – and Geoff is also here; can you raise your hand? There's Geoff – so we decided to directly communicate with Geoff and have a discussion on what could be done about this, because it's pretty serious. Next slide please.

So if we go into more detail; we looked at just one of those organizations that are scanning the net for what certificates are out there in the wild and they found that at least – and this is really the low watermark – at least 157 CAs in the world are giving out certificates for TLDs that don't exist, which are those Internal Name Certificates. But the number of certificates that are out there in the wild are of course impossible to know because the CAs don't give out number or information about what certificates they have issued. So we didn't even think… This was so bad so we didn't even feel we had to go into disclosure policy for the CAs to be able to get this information because we found that… We see enough of them out there so we found that we need to issue this recommendation anyway.

We also, by contacting… If new enterprises are looking at recommendations and standard configuration in various software, that is very popular. We also found that enterprises very often use Internal Name Certificates for a variety of reasons, and this is something that is very, very widely used. Next slide please. So as I said, the ability to get a man-in-the-middle attack by getting a certificate before the TLD exist and using it after the TLD exists, that is the actual problem. We also contacted the CA Browser Forum, which is the organization with the CAs and browsers, where they come up with best current practices for their operation – and we found that they were aware of this issue and they already had agreed to stop issuing Internal Name Certificates in October 2016.

And our conclusion from SSAC, which led to our recommendation that we handed over to Geoff and the Security Team, we came to the conclusion that a three-year window is far too long. I'm sorry. Next please. We didn't think that three years for a man-in-the-middle attacks would be a good thing. So far I haven't heard anyone that disagrees with us. So what we did was that we asked the ICANN Security Team to immediately develop and execute a risk mitigation plan. So that's where SSAC job stopped. And just to make this session flow a little bit better, what I now will tell you is what is also in our document; in Appendix A. But it's really what the ICANN Security Team did.

In cooperation with us and SSAC, but now instead of having Geoff running up here I'm just going to go through this. So what ICANN did was to alert the CA Browser Forum Chairperson, on January 23rd. The result of that was that ICANN was invited to the Annual Meeting of the

CA Browser Forum and there was a presentation down on February 5th. So to some degree it was pure luck that their Annual Meeting happened to be at this time of year. That was actually pretty damn good. Sometimes you need to be a little bit lucky. So all the Members of the CA Browser Forum they absolutely understood this.

So what they did, which I think is really impressive, is that they whipped a ballot together, ballot number 96, which is how they count their ballots, and they brought it forward and it was passed already on February 20th, okay? This is pretty quick. When they agreed to change their practice completely and this ballot says that CAs will stop issuing certificates and in applied for gTLD strings within 30 days of ICANN signing contracts with the Registry operator and that CAs will invoke any existing certificates within 120 days of ICANN signing the contract with the Registry operator.

So they immediately changed from waiting until 2016 to more or less do these sorts of things immediately. So that is what the… One more thing I should mention is that after all of this happened, we in SSAC felt that now we can release this report with this additional description, so beginning in March the report itself, with this Appendix that told about what happened between when the report was ready and when it was released, in Appendix A, as I just said.

Another thing that had to happen, that was part of our recommendation, is that ICANN also developed a disclosure policy that then was followed – so a disclosure policy was implemented, some of you might have seen it when it was announced, and that disclosure policy was used – so that's a good thing, we can reuse that disclosure

policy. Hopefully, of course you don't want to see this kind of incident again but of course you must be prepared for it to happen, so that's another thing that is a result of this.

So I think that was the last slide for this. So are there any questions on this please? Yes?

JEREMY ROWLEY: Hi, this is Jeremy Rowley, I'm part of the CAB Forum and actually the author of that ballot 96, but I just wanted to point out that although the CAB Forum has passed ballot 96, that doesn't mitigate the entire problem, simply because it's not effective on all CAs until adopted by the browsers. So although the CAB Forum has down what we can as CAs to mitigate that problem, until it's universally adopted by Opera, Mozilla, Microsoft and Google, you might still see some of these certificates in the wild. So you can't expect that 120 days after you've signed the agreement that every CA's going to stop issuing, especially if they're not Members of the CAB Forum.

PATRICK FÄLTSTRÖM: Thank you.

WARREN KUMARI: Warren Kumari. I also want to mention that revocation is not completely effective, if you can block access to the CRL or the OCSP server, even thought the certificate's been revoked, a lot of things will still accept it.

UM: With this latest change from the CAB Forum, have you done a residual risk assessment to see what now is the position?

PATRICK FÄLTSTRÖM: No, we have not done anything more after this, no.

EBERHARD LISSE: Anybody else?

PATRICK FÄLTSTRÖM: Let me expand a little bit on that because that was maybe a little bit too short an answer. No, we have not done any risk assessment and although we in SSAC have our own meetings when we are discussing these name spaces used, where we will discuss various name spaces used tomorrow, we have so far not been contacted and asked by anyone to do a follow-up on this, and we have so far no consensus internally too, because as I said we took up this issue on our own and so we have the ability to of course continue if [you want to? 01:13:18] our own, but so far we don't have any… We don't see any real reason to do that yet. It doesn't stop us from when we do that but this is really where we are at the moment.

And some of you to continue this name space issue, some of you might also, for example, have seen the letter from PayPal to ICANN about some of the main space [collision / coalition? 01:13:44] regarding some of the applied for TLD strings. Okay, in that case, good. In that case

we're moving over to the next report, SSAC 58, and I'm happy to hand the microphone over to Jim Galvin.

JIM GALVIN: Thank you. Next slide please. So we know that validation is of course a significant topic; both with the Registry Agreements and the Registry / Registrar Agreements and also in the relationship with law enforcement, it's been a big part of the discussion related to accuracy of the data and part of the issues that are being examined by the Expert Working Group that's been created to look at the next generation of directory services. So in the same way that SSAC took on responsibility to look at the taxonomy of WHOIS in general and try to shape people away from using just the term WHOIS, we created a taxonomy for WHOIS, identifying three separate and distinct topics.

We've also looked at validation in an attempt to separate this topic into multiple discussions to provide a framework for hopefully facilitating closure on this topic of registration data validation. So next slide please. One of the things that we discovered and wanted to call out explicitly, of course, is that data quality is relative to the Registrants and their purposes. So obviously different people have different reasons for wanting access to the data, and in fact they then had different reasons for wanting access to different elements because there are different things they want to do with it. I think it's important just to identify this and call it out explicitly.

It affects what people think about directory services in general. There tends to be this desire to simply collect everything and then make

everything available. And it's important to point out that one could make distinctions there in that data instead. And so that's where this finding comes from. The second thing that we looked at was in considering the issue of validation, there are different ways in which validation can be accommodated.

And we wanted to explicitly separate out items or mechanisms that can be automated and mechanisms that cannot be; because oftentimes in these discussions we find people talking past each other. One group will be talking about 'well, this is easy, you can just automate it, let's move on', and yet the group they're thinking with has a different purpose in mind and they have different elements in mind and from their point of view you cannot automate the protection and it really is a manual process.

And so we wanted to explicitly take a look at some of the mechanisms that could be used for verification purposes and what it would take to provide some validation in those categories. And then the last finding here is that different contact data elements have different validation core structures. And this would follow naturally from the idea that you could automate some of the validation and some of it needs to be manual – and I think it's pretty straightforward to state that anything that's done manually is going to cost you a good deal more than things you can do in an automated way.

There'll obviously be an investment in created the automation, but long-term the cost will drive down much closer towards zero and become much more cost effective. So we wanted to expand on that and make that statement and clarify those issues. Next. So the first

recommendation is in particular to identify a taxonomy for validation. And we created three particular categories of validation. I want to emphasis here that these are equivalent categories of validation. They are not intended to be ordered in any way. They are all peers to each other.

So it is not our intent, it is not SSAC's intent to create any one of these terminologies – any one of these terms – and suggest that one is more important than another. Again, given the finding that the purpose of the data and the validation for is dependent on the purpose and the Registrant's need, it would make sense that these things should all be passed to each other. The first is syntactic validation and the particular phrase here to describe it is the assessment of the data, okay, satisfies it's particular constraints.

Now, to a large extent one gets syntactic validation, you get most of it just from the use of EPP, because XML of course has certain definitions that have to be met in terms of the data that you put in the particular elements. So a good portion of that is accomplished in that way. But one of the things that we want to get at here is another level deeper in looking at not just a telephone number, for example, is comprised of Arabic numerals, but is it in fact a valid telephone number from a syntactic point of view? Did I get enough digits? Do the digits make sense within the region for which it has expressed that it might be used?

That kind of thing. So syntactic validation is again just asserting that the data that you've been given has the potential to be useful in the context in which it's being offered to you. Does the postal address really look like a postal address? Does the telephone number really look like a

telephone number? Does an email address really look like an email address?

Operational validation is then the assessment that the data could actually be used, for whatever routine purpose it would be expected to be used for. So if I got a telephone number, is it actually a valid telephone number? If I got an email address, is it actually a valid email address? Could I send an email message to it? In the case of a phone number, could I actually call that number and see that it is in fact in service? So does it function like a telephone number. In the case of a postal address, could I actually send a postal message to it and would it actually be delivered? So is the address actually a reasonable address in the context in which it was taken. So that's operational validation.

And finally in the last case we have identity validation. And this is the next step in operational validation and it's confirming not just that the email address is deliverable, but did it actually get to the real world identity that it's intended to represent. If I send a postal letter to a postal address, does it actually get to the real world identity that's behind that postal address? And similarly with a phone number, does it really belong to that real world identity?

Next slide. So the second place that we went to was to put together a set of questions that we think are important when you're examining validation issues. So as you continue to look at the contact information, in particular, that is collected for registration data, one has to examine the costs and benefits of the validation of that data. And not all validation will be appropriate in all contexts. So these are the contexts that one needs to examine as you think about which categories of

validation you want to apply to the data. So you really need to examine why you're collecting the data and is it really meeting the needs of those who are going to use the data?

So you're going to have a set of Stakeholders who have certain expectations for the data you're collecting. And so are you collecting the appropriate data that these people need to have? That's one question. Is the additional registration processing overhead, okay, is that acceptable for improving the accuracy and quality of the registration data? This gets directly to the question of when you do the validation and when it's supposed to be performed. To a large extent our expectation is – and this is stated in the document – that syntactic validation is something that could be automated and could be done at the moment of registration; it could be done inline. That's the expectation in creating that category of validation.

Operational validation is somewhat mixed. There are some things that you might be able to do inline with the registration and automate and there are other things that you couldn't. So for example to some extent you could check the email address – and there are reasons why that may or may not work and those are detailed in the report –, on the other hand you really couldn't check a postal address, the operational validity of a postal address. You can roughly speaking check some of the syntax – and this is detailed in the report – but you really couldn't hold the registration back whilst you send a postal letter to that address to see if it got delivered. Okay? So that would be inappropriate in that context.

The third question here is for those things that have a higher cost associated with them, is that appropriate? Because you really do need

higher accuracy and greater quality, and that's an appropriate question to ask. And not all validation is appropriate in all contexts, and this is the question that we want to put in front of people to ask when you're considering whether you're going to validate the data or not. And then finally of course; would accuracy actually improve if you had natural persons having the privacy protection, which it is often asserted people want and need and that's asserted as a basis for inaccurate registration data.

It's a commonly held belief that people lie on their registration data, give misleading information, because they're looking for privacy. And so one needs to step back and ask the question as to whether or not this would create an improvement or not. Are we solving the right problem by validating the data? And I believe that's it, isn't it? Oh, okay, right. And so the last one is just in general the point about identifying the validation techniques that can actually automate it, and then to create policies that incent the development of those things.

There is actually, we believe, a fair amount of validation that can be automated. And it would make sense in general for those automation techniques to be applied – to be developed and deployed. It's fairly straightforward to suggest that there's a very high cost/benefit ratio to doing that. If I make a small investment up-front I can improve the quality of the data and in general that should be a good thing and we should do that, and ICANN should seek to incent those things to come into existence and create an environment in which it'll actually be deployed by everyone, and in the right places. Thank you. Any questions?

EBERHARD LISSE: Yes, as the prerogative of the Chair, I have a question, but I don't really understand what good this will all do as long as an anonymous registration or proxy or whatever, WHOIS data services… I have recently been slandered by somebody, I cannot get the Go Daddy or Domain for that proxy to give me the registration data of them. What use is all this work?

JIM GALVIN: It's correct to say that this work does not do any policy development. So the question of whether or not anonymous registrations should be valid or not is not spoken to in this document. This is about validating the data.

PATRICK FÄLTSTRÖM: Let me expand on that a little bit. One of the things that we talked about in SSAC is when looking at, for example, proxy registrations, and this is when Jim talked a little bit about regarding anonymous registrations. Obviously there is an interest of being able to do anonymous registrations. Maybe it is the case that the data that is still collected with the registration, there are reasons why people want to do anonymous registrations. And maybe it is the case that by applying some specific privacy or safeguards around that data, that would increase the interest from people trying to actually register with their correct data, as long as the data is not publicly available.

That's one of the things. And this is not only people who want to commit crimes but also it's common, as you know in you room probably better than me, for example when you have an upcoming trademarked registration that a company might not want to disclose that they are looking for a specific Domain Name. So for example, they register a Domain Name but they don't want to do it under their own brand name, unless they have registered that trademark in all different kinds of systems where Domain Name Systems might be one.

There are also other reasons why people would like to register a proxy; it might for example that they really want to do business in, for example, just because we're in the ccTLD Tech Day, it might be the case that they would like to register their Domain Name in the Domain and they cannot because they don't have an operation within that country, so they try to circumvent, also, registration rules. That's another reason. So what SSAC is saying is that without having a proper taxonomy and talking about these different validation mechanisms, we can probably not move forward in the discussions that you just talked about.

EBERHARD LISSE: Okay, you obviously not [divide forum? 01:27:28] to address this, I just thought that I'd bring it up because I found it pertinent. There is a question over there?

CATH GOULDING: Hi, thanks. Cath Goulding from .uk. I'm just interested if you're looking in just Registries or if you're going to look at how Registrars take the

data? So for example when they pay using credit card information, it's obviously in the Registrar's interest to get payment.

JIM GALVIN: We make no statement as to where validation is implemented. So we actually don't make a distinction between Registries and Registrars. This is about applying validation and the appropriate way to characterize that.

PATRICK FÄLTSTRÖM: And let me just add that another continuation of this, which is something that we and probably you also here in discussion with, for example, law enforcement and others, that for legitimate reasons would like to access the data. Sometimes it's pretty important just to know how valid the data is because you need to know what kind of validation was done. So another thing might be for example, okay, is there something that is needed in the EPP protocol? How do you communicate these kinds of things between the Registry and Registrar? Is it a case that the value of the quality of the validation degrades over time?, etc., etc.

So there are many, many, many large portions of follow-up work that could be done, but we have not done so. We stopped here, at exactly where Jim said.

JASON POLIS: Jason Polis from Super.Name. With regards to the identity validation. In some jurisdictions, for financial services it's required that you know your

customers, that you go ahead and do identity validation. Are there any jurisdictions at the moment where 'know your customer' procedures are required for Domain Names?

JIM GALVIN: Yeah, the question… Let me turn the question into the room. It's not really our place to know the answer to that. I mean, I personally don't, and SSAC certainly hasn't talked about it. Would anyone like to speak to that question?

EBERHARD LISSE: I also wanted to ask, is anybody in the room, where in whose country it is a requirement that you need to know your client? Yes, for financial services it is maybe too, but Domain registrations do not fall under that aspect of the [AC? 01:20:01]. We checked that. Our lawyers for example need to know exactly who we are, our bankers do. But we as a Domain Registry, we don't. And I think the financial intelligence [ACs? 01:30:14] of most countries specify this but is anybody aware of a country where the Domain registrations are required?

SULAIMAN AL RAWAHI: Hello, my name is [Sulaiman Al Rawahi? 01:30:27], I am from Oman, from the .om Registry. We do require that actually, because for .om and the .oman Arabic id and for Oman, it's only allowed for companies and institutions that are registered in Oman, or even international companies that have representatives in Oman; legal representatives.

They don't have to have an office in Oman. So that's why we have to make sure who the client is.

EBERHARD LISSE: Is this policy or is this law?

SULAIMAN AL RAWAHI: It's part of our Domain Name's policy.

EBERHARD LISSE: Now that's not the question. The question is whether there is a country where there is a law that requires this. Our policy is so that we also have to have an existing company with an existing individual, that they cannot [authenticate? 01:31:16] it, and that's not the law. The Financial Intelligence Act does not require us, for example, to establish that the client actually exists if we do it. And I think that was the question.

SULAIMAN AL RAWAHI: Okay, thank you.

EBERHARD LISSE: Okay, this is the last question then.

UM: In .us we have to know who they are, but the difference between law and policy is a fine line because our regulator is the US Government, so it's kind of a law that we can't have anonymous registrations and they

have to be… Within the US to even be there. So I think the US might fall into having that law.

EBERHARD LISSE: All right. Okay. Thank you very much. [applause] So the next would be Ms. Xing, Melody Xing. I hope I say this correctly because my Chinese, as little as it is, it terribly broken. Ms. Xing is from Coremail, our sponsor. She will talk about id an email, which I think in China is, as we all know – have a seat – a well known issue. Let me just get the system connected. Let me get the signal working. Where's your presentation? [background chatter] There you go.

MELODY XING: Thank you. This is [inaudible 01:33:25] right? I'm the last speaker this morning and I will speak up and soon we can go for lunch. Before we start I wonder how many of you are non-native English speakers? One, yeah. I see many faces. It's well known that the Asian languages are quite different from a native language. Some said many years ago that if there was only one language that can survive it must be English, because English is still the official language of computers and programming.

But things are quite different nowadays with people who can transform yellow and white into many languages. The UTF-8, the encoding, known as unique code, can teach the computer to know more languages. And better still, we now have International Domain Names and we have emails to support the International Domain Names. So today I want to share with you about our experience, our interpretation of IDN email. My speech will be divided into three parts.

The first and second parts are the introduction of Coremail and IDN email. And the last and most important part is the problems that we've conquered. And so [inaudible 01:34:54] during the implementation of IDN email. I will discuss three problems. The first – what the system has done to support the standard for the IDN email, and the second – how to deal with IDN email and non-IDN email, and the third – how to know who supports IDN email.

Before we start, please allow me to introduce Coremail. We are a professional email system software supplier in China. We provide overall email system solution, to governments, institutions and enterprises. We now have over 600 million users of email in China. Our major customers include Chinese Academy of Science, [State Consul? 01:35:47] Information Office and the [People's? 01:35:49] Republic of China. During 40 years of development, we all need to email, and we're experts in our area.

In 1999 we created the first Chinese email system. That system was to support a huge Chinese operator, representing the [binaries? 01:36:10] and we are also the most broadly used email system in China. As I mentioned, we have over 600 million people using Coremail and we are also the first choice for our government institution and enterprises. And in 2012, as the strategic partner of CNNIC, Coremail has participated in IDN email… And [deliberant? 01:36:43]. And now we have the honor to be the first IDN email system in commercial use.

So what is IDN email? IDN emails are emails with International Domain Names. That mean they contain non-English characters, such as Chinese characters, French, Korean, Japanese and so on. IDN email has a

significant influence in China and in many parts of the world. Ordinary people might not know English as well as you do so it is difficult for them to remember many of the English characters, however, if you have an email address like this one; it contains the person's name, company's name, written in Chinese. I bet the Chinese can remember it at the first glance. And this email address can also be written in different languages.

So internationalized email addresses are kinds of personal ids, with national and ASCII characteristics. Since IDN emails contain non-English characters, we should do a lot to reconstruct our system. Sorry, this one is what we have done on our demo system; you can see the receiver's email address is written in non-English characters. I guess… I wonder if anyone from Japan...? Yes, I see you. I think it is easy to remember the last email address that is [inaudible 01:38:49] from Toshiba. So you can see that it is very easy to remember the name and also the company's name.

The first problem I want to show you during the implementation of IDN email, that is the protocol and standards. It's well known that the origin of email technology project code does not support the UTF-8, so therefore the International Internet Technology Organization, the IETF, has formulas that were standard to support the sending and receiving of UTF-8 email. You can see that those protocols can make the different languages compatible with each other.

The RFC 6531 can set-up email with an UTF-8 heading and RFC 6532 can direct UTF-8 coding. And RFC 6856 can support UTF-8s through IMAP and POP3. As I mentioned, Coremail has a huge community of users of

email in China, so it impossible for us to make a one-stage replacement. We have done a lot to restore our system to support our standards. They mainly rely on the email id; as you know the email id is everywhere in email systems, so you really are [here inaudible 01:40:48]. Almost every part of the system should be reconstructed. Also, we may come across many problems, but the many problems rely on the system compatibility of our different languages.

And the second one is how to deal with IDN email and non-IDN and email. Since not every one has IDN email so we offer a double or a double email account solution to IDN email users. An IDN email user has two accounts. One account with an IDN email address, another account deals with English email address. The main account will be used between communications of those who support IDN email. Rather, if the user cannot support IDN email then the alias account will still work.

Let's take a closer look at the sending and receiving procedure. Receiving email is another big problem because the system can classify the email with different standards with the help of RFC 6531, the IDN email can be easily received. However, sending email is quite complex because the system has to make a judgment whether the receiver of the system can support IDN email or not. If yes, the email will be encoded by the standard RFC 6532 and sent by the RFC 6531. If not, the email will be encoded by RFC 2045 with the alias and sent by RFC 2821.

So here comes another question – how to know whom supports IDN email. We have our solution. It is easy for those who use the same system because the system-written code can contain IDN information, so the system can tell whether it is IDN email or not. But if a receiver

and sender use a different system, that will be much more difficult because of Chinese special network involvement. System-written code will often be shut and so Coremail offer another cloud-service solution. As a [inaudible 01:43:55] we are the most broadly used email system in China, so we can gather the system-written code and information as much as we can.

So based on our cloud-service [inaudible 01:44:12], we can update the IDN list so we can make a judgment whether the receiver or sender are using IDN email or not. We also suggest all the IDN email supporting systems [inaudible 01:44:19] can release and show their Domain records. In that case we can judge whether it is IDN email or not. That is three problems we conquered; the area of implementation of IDN email. Some still ask us whether there are huge technical problems or there is a technical barrier. Actually there's not.

Our experience is that we have found that the most important part is how to make a different language compatible with each other. That is the major problem. And the email system supplier, no matter if Chinese or foreign supplier, is welcome to join us in the promotion of IDN email. And now we are technically ready for IDN email with a Chinese email address. We now have customers like the Chinese Academy of Science and CNNIC, and should we apply IDN email on our cloud email service platform, Coremail, which has already millions of users, and later we plan to deploy our customers like government customers and large enterprises customers to make a better promotion of IDN email.

Thank you very much. Are there any questions?

EBERHARD LISSE: Well, let me make a little comment from the Chair. I agree with you; this is very complex.

MELODY XING: Yeah. [laughs]

EBERHARD LISSE: Any questions from us? From the floor? Unfortunately my native language is German, there is only few IDN characters which we can do away with through [inaudible 01:46:45], but if you use a totally different script?

SAM SALIF: Thank you. I'm [Sam Salif? 01:46:53] from the Telecommunicator of Sudan. I was wondering how your email system can deal with the other known systems like for example Hotmail or Gmail, when you type a Chinese letter in, how will they understand it? And I see there is RFCs approved, did the browse have to update the [identification? 01:47:24] to accept the Chinese letters? Thanks.

MELODY XING: Yeah, email system should support UTF-8. In that case they can receive an IDN email. If not, it cannot be… The email address will be shown with an alias account, as we mentioned. It is the traditional English email address if the system supports IDN email; the email address will show for example a Chinese character, a Japanese character and other

characters. So for example Gmail they should support the UTF-8 before… The system of Gmail should support UTF-8 before they can make it work. But as we know, so far Gmail still cannot support the UTF-8.

So you really are waiting a long time for the promotional IDN. You really need all efforts for the email system supplier, because if my email system supplier can support UTF-8, the internationalized email can finally get a [chip? 01:49:17].

EBERHARD LISSE: All right. Thank you very much.

MELODY XING: Thank you.

EBERHARD LISSE: It was quite interesting, this is what many of us who only speak and use English email have no concept of or not idea, as big as a problem as it is in your country.

MELODY XING: Yeah, for example my mother who does not know English, I apply an email for her and she said 'I don't know English character, I don't know what is the meaning of this character?' So for her it is impossible for her to remember the email address. But if I can use a Chinese email address, because it is her mother language so it is easy for her to remember it. So I think in China and many parts of the world there are

still many people that don't know English and so it has a significant influence.

EBERHARD LISSE: But to be honest, mothers tend to have different problems that we have.

MELODY XING: [laughs] Yeah.

EBERHARD LISSE: All right. Thank you very much. I think this was a very good presentation.

MELODY XING: Thank you. [applause]

EBERHARD LISSE: A small housekeeping announcement. We have got about 80 tickets. I've done a headcount; it will probably work out well. You go, you file out this way, grab a ticket and then file out that way and follow the lady over there to the restaurant.

[Tape change – ccNSO-tech-2-08apr13-en.mp3]

SPEAKER: For the transcript, this will be Part 2 of Tech Day on April 8th 2013.

EBERHARD LISSE: Okay, so, welcome for the postprandial sufferers, and of course I'll punish the ones that are here for the ones that come late. You will suffer a little bit from the stragglers walking in every ten minutes. But welcome. Dave Kipling is from NCSS, NCC Group, one of the Escrow providers and he will make us a presentation. And I've asked them to go a little bit more into the technical detail so that we can have a little bit of an idea how this is done and maybe we can take up some ccTLDs who don't have to have Escrow maybe can pick up something.

DAVID KIPLING: Okay, hello and welcome to NCC Group's presentation on Escrow for ccTLDs. I'm David Kipling as Eberhard said. I'm a solutions architect and I'm here just to explain why Data Escrow is essential and how it is achieved. Okay, so the presentation is going to focus on why Data Escrow is required. Okay? Then moving onto what the latest ICANN requirements are. Following that we'll have a look at how data security is managed; how data and integrity is ensured, and then finally how to keep the customer kept informed. Okay?

So, why should you Escrow your data? Obviously with ccTLDs there's no requirement from ICANN for you currently to do this, though there are many good reasons why you should. Data Escrow helps in the scenarios such as loss of data in you Domain. No doubt you have robust DR solutions in place, though Data Escrow gives you and equally importantly

your customers the security of daily copies of the secure data, which is being validated and impartially verified by a third party.

Another reason is about securing your business and your revenue. What would occur if your services were lost and you were unable to restore them? Data Escrow works to make sure the data is always safe, available and as we'll see later, usable. The reputation and trust in your TLD has a huge impact on your business in revenue. Data Escrow itself does not give you a five-nines-uptime solution at all. What Data Escrow does give you is that measure of a last resort. Okay? So the security; the customer's data is secured and verified if required in an emergency.

This instills trust in your current and potential future customers, that you understand the importance, not only of their data, but of their interests. Okay? As many of you are aware, with the intro of the New gTLDs, there is now a requirement for every gTLD to Escrow their data. This shows ICANN's commitment to making the whole industry better equipped to handle disasters such as the RegisterFly scenario. With gTLDs having Data Escrow mandated along with the extra data integrity checks involved, they could be perceived as better prepared for a disaster and obtain a commercial advantage through increased customer trust of the New gTLDs over other TLDs.

Okay, so to summarize, Data Escrow gives only not you peace of mind that your data is safe, available and usable, but also gives your customers, both current and future, the confidence in your services by independently securing and validating your most valuable asset; your customers' data. Okay? So now I'll just go through what the latest ICANN requirements are around the gTLD process. So the New gTLD

process is currently in the draft stage soon to be finalized. There was a new one actually published I think on Thursday, so there's just minor changes being made to that at the moment.

What I want to do is just highlight the processes involved and what's to learn from handling all TLDs, not just gTLDs using this process. The differing Escrow requirements that are currently out there across TLDs, highlights the need for a consistent approach, to allow for a rapid response in an emergency. Also, a simple and straightforward process is needed to ensure the data required to run the Registry is Escrow Daily in a consistent format, to then help within that emergency scenario.

The overall aim of the New gTLD process is that if the backend Registry operator was to fail, that the processes put in place would allow ICANN and the Registry operator to act and where required instruct an emergency backend Registry operator to resume services. One key aspect of the process is to be able to restore without the need for lengthy legal battles around the data, as seen with RegisterFly. Okay? So as you can see, the New gTLD process involves the backend Registry operator depositing encrypted files to the Data Escrow agent. So the deposit files are supposed to include all the information to run the Registry services.

So records such as Domain records, host records, contacts, Registrar records. These are then deposited daily. ICANN initially specified that there should be a full deposit made every Sunday and then the rest of the week should be differential deposits. I think with push from the community it's now been accepted that you can deposit full deposits every day, which is a good sign there of seeing the community giving

feedback to the process, just to make it again that simple process that everyone can easily follow.

Again, once that data's been produced by the Registry operator, it's encrypted and then securely transferred over to the Data Escrow agent. The Data Escrow agent then securely stores that away in multiple locations and then alongside that we have the verification steps. So once we've got that data we just want to make sure that it's a valid deposit and it is usable data. So that's where the Data Escrow agent will then decrypt that data, validate it, make sure it adheres to a set number of rules, which I'll go through in a minute, and then obviously notifies both ICANN and the Registry operator of the outcome.

Sorry. Okay. So now going onto data security. Within this process the utmost concern is the security of that data. So you want to Escrow your data securely and effectively. You want to make sure you've stored the files that you were expecting to receive. You want to make sure that that data is safe during transit. How do we then go about this? We have a [10 cc? 00:09:27] we use across our Escrow solutions transfer protocol, such as SFTP, FTP, HTPS. Again, we've found the most prominent in this sort of scenario, where you're doing a daily deposit, is SFTP. Again, normally coupled with SSH keys for password lists authentication. Again, just because this is generally just a batch process that's run in the background.

So again, this gives us this reliable connection then, the secure authentication with the client to make sure it is the person we're expecting to connect to. Then we move onto the actual deposit itself. Like we saw in the diagram, the data is initially encrypted, so we do that

generally using PGP, so we'll provide a public key to allow the depositor to encrypt the data, we'll securely hold the private key ourselves to use then to decrypt that data.

Finally, we have digital signatures on the files, so that allows us to make sure that that file we've received hasn't been tampered with during transit and also that it was generated by the intended person. Okay? So what we've got here overlaid is the use of these keys within the process. So the Registry backend operator initially, like I mentioned, encrypts it using the public key, which has been provided. Then we have the through four authentication; the SSH keys being involved at both ends, and then finally the data being decrypted by the Data Escrow agent, to allow them then to validate the integrity of that file.

So moving onto data integrity. Once we now have the file over at the Data Escrow agent, what we want to do with that file once it's decrypted is we want to make sure that that file is a valid deposit. So there are a number of checks we can do; I've already mentioned the digital signatures, but on top of that there are schemer checks that can be done. Just to make sure the format of the deposit is correct, but most importantly on top of that comes the verification of the content. So these are the processes involved in really making sure that that data is valid and usable.

The current extent of verification steps required by ICANN are the ones displayed. I think initially we have the schemer check; so we're checking the format of it. Then we move onto actually interrogating the content; looking at counts, following the counts. We'll also look at linking records, so if there's references within that data, which should link to

other objects within that data, we want to make sure they're all present in that deposit as well. Okay?

So there are currently two formats of the deposits that can be made for the gTLD process. We have an XML only format and then we have an XML manifest followed by CSV files, which actually hold the data of those objects. This obviously poses challenges for the Escrow agent, having the two formats, but again this was the feedback from the community as to when handling large amounts of records, that this was one of the preferred formats.

If I just briefly go through – with the XML format, on the left-hand side you have the overall manifest there where you're getting high-level information such as an identifier for the deposit, the type of deposit – this is whether it's a full deposit containing all of the objects for the Registry or whether it's a differential deposit since the last deposit was made. Again, there'll be a menu in there telling you what type of objects are actually contained within that deposit.

And then finally we move onto the actual real content of the records. So in a differential we'll have a delete section, but within a full deposit all you're going to get is a content section. Within here, the specifications then allow you to add the different types of objects to represent your Registries data. So as you can see in the example, we've got some Domain records and host records in the grey section on the left. What I've done is expand one of the Domain records to show you for the XML format, that data is actually represented within the main file, whereas if I move onto the CSV format, instead what you get is actually elements in there just to define how that object is represented in the CSV files.

So you'll see here a contact on the right-hand side. It takes a number of CSV files of different formats, to actually represent one contact record. Okay? So the file structures, as you can see, then expand out onto the right, showing you what fields should be included in this CSV file. So this allows, through the specification, to state this is exactly what the types of the fields should be within the CSV files. So again, both formats have the space set of fields for all the records; your Domain records, your contact records, etc., that must be supplied.

So these must be supplied just so an emergency backend Registry operator would then be able to resume the services if required. Again, there's checks in there around the types that I've mentioned; both on the XML and CSV formats. Again, these have to be adhered to, otherwise the deposit will fail the validation. Another interesting concept is the use of custom schemers. If the backend Registry operator identifies there are a number of extra fields that they believe should be still Escrowed, there's an opportunity for them to generate custom schemers themselves to make sure this data is present within the deposits and also of the correct type.

Okay, so now we've gone through the schemer checks that are done around the actual structure of the deposit. What I want to do now is go through the verification rules and where their importance really shows. As you can see with the example we've got here of the Domain record, you can just go into all the elements in the deposit and you can look to just put 'the quick brown fox jumps over the lazy dog'. Okay? So it's not really useful in an emergency. You're not going to be able to restore this to a valuable record.

So where the validation comes in is you can see with the blue underlining on the right-hand side, that some of the fields are already identified 'invalid'. This is just through the XSD types; so say down the bottom we've got some dates, etc., so you can see the failures there. But when some of the checks really come in is say in the following scenario where the deposit looks like a valid deposit, it's just when you get to say the invalid reference that we have for the contact – so it's a contact for this Domain record, and in the deposit we have this text, which does pass the validation of the type of field, so it matches the pattern of the value that should be in that element. But again, it doesn't actually match to a contact record. Okay?

So to work around this, what the verification rules do is then go through that file and check that that contact actually does exist, proving that that data is actually usable. So just moving through the extended verification, which is actually run through, like I've mentioned, we've got the counts, so in each deposit there's a header record which details the number of records contained, the number of contacts, the number of host records, etc. These are then compared with the number of records within that deposit, for a full deposit. If it's a differential, a process of actually getting the last full deposit and then applying those differentials before we can actually get the counts to compare, is run through. But this just serves to make sure that you've actually got the full set of data to that correct deposit.

The link contacts – so the contacts linking to Domain Names are present – so that's the example I've just gone through; making sure that any contacts linked to Domain records are present as contact records.

We've got Registrar links to objects, so any changes they've made to any of those records, any of the contacts etc. are all detailed and present. The next one, actually the specification that came out on Thursday, that's actually just been removed because there were some queries around that process.

But moving on, we've got the policy element. I think I worked with Gustavo just to clarify some aspects on this, so there's been an update within the recent draft, so the policy element just allows a Registry to state extra mandatory fields in that deposit. So if your business rules for someone to restore your Registry, they would need this field to be mandatory, that can be defined within the deposit to make sure that check is done. Okay?

Right. So how to achieve these rules has been quite a challenge. Again, where we could be facing millions of records, some of these XML files uncompressed are gigabytes of data, so this means that the use of normal, say XPath queries or Dom technologies start to fail due to the intense resource usage. So you're talking there not only memory but also processing power. With some of these checks like the linking records, you are having to do say a scan through the record to find the initial reference and then you're having to search all the way through the file to actually find whether it's present.

So to work our way around these solutions, we've actually looked to do a custom in-memory solution. So what we tried to do it, based on those rules in place, is scan through that file initially, created indexes of the information, to allow us to at the end just do comparisons with our in-memory indexes, so therefore not requiring to load the whole, say,

gigabyte file, into memory to actually do this processing – we can just stream through, pull into the structures we require and then use them to give us a really fast result as to whether it passes or fails the validation.

How to progress from here? The extent of the verification rules, a bit like I mentioned, they're in draft. Again, there were some changes made to them on Thursday, Friday. They're still open for comment to go through. Any ideas people have to help with that are just being taken on board. When the deposits are made and validated, like I mentioned, you have the reports that are then sent out to both the backend Registry operator and ICANN. Going forward, heuristic analysis can be done on these, just to try and give us that early warning light of whether there's a problem.

So if suddenly the counts have dropped down significantly, at least then there's some kind of a warning light that is then triggered to help assess whether there's a problem. Finally, we've got the IETF Forum, so this is where a lot of the discussions around the draft specifications happen. It's worth having a look on there at the discussions, just to have your input. It's a really good place where aspects where the CSV format that I've mentioned, that originated on the Forum. Again, there's been quite few comments around that and they're actually looking to merge the drafts into this consistent approach.

The acceptance of the full deposits instead of the differentials, there were discussions on the forum there that finally led to that being accepted and included in the draft. So what I want to do now is just go onto how NCC's process solution. NCC focus on putting the customer

first. We're trying to aggregate all the data, the reports, the billing, to simplify the process for the backend Registry operators. The notification is not only be email; if there's a problem we've got processes to actually contact you.

A quick, clean portal to help you with the resolution of the failures. I will just quickly move onto just some example screens from the portal. So what we've tried to do here is have that traffic light set-up where you've got a clear red, amber, green to show you that red: there's something that needs to be done now, there's an immediate action that you need to look at. Then you've got sort of a warning where there's been a problem but it doesn't require urgent action, and then finally green; that everything's okay at the moment, there are no actions to go forward.

So we've used that throughout and then with the aggregation on top of that, just allows us to have this quick, clean system, which allows you to then drill down to the important information. So if you're managing multiple Registries, this allows you to just go on and see high-level, is there something that I need to action or is everything running smoothly? We also contain the detailed logs just on those verification steps, so you can see where it failed if it did, or if it passed etc.

So in conclusion, Data Escrow is crucial not only for the recovery process for a Registry operator, but also for the trust and peace of mind of your customers. Data Escrow shows you have an impartial third party confirming your data is secure and usable. This generates trust and can lead to increased business and revenue prospects. We discussed the latest ICANN requirements and what processes are involved to perform

the Data Escrow for the New gTLDs, and how that affects you, the ccTLDs, in the changing marketplace.

We then discussed the processes involved in securing the data and also how we can make sure that the data is fit for purpose. Obviously, don't forget the forums that we've mentioned. It definitely is worth going on there just to see the discussions that are going on and how that could affect you in the future. And then finally we had the review of the NCC solution, where we aim to make Data Escrow a simple managed by exeption process, giving you all the benefits of Data Escrow with minimal effort.

I hope I've given you a good insight into the process of Registry Data Escrow and what business value it can bring you and your customers, especially with the introduction of the New gTLDs. As you can see, you can find us at booth 16 or again we've got… I'll have to update the telephone number there because that's actually incorrect, but we've got the email address and also our website; nccgroup, where you can also see the other services provided by NCC.

EBERHARD LISSE: Thank you very much. That was a good presentation actually. We've got time for one question? Still everybody is tired after the lunch. [laughter] No, I must say you have answered some things, which we never considered. I pick up by whole database and ship it with a CP out, and the integrity is tested only on the other side, but because I don't anticipate anybody needing to take my Registry over, but if some people in government have designs on it. But as I said, it's still something that

the process has been designed by ICANN for the gTLDs, but if you want to do Escrow, why [did they? 00:30:30] invent a wheel?

And these validation issues I wasn't really thinking about. Have we got anything from remote? Okay, thank you very much. There is one question here.

HUTTI TUTTI: Hello, my name's [Hutti Tutti? 00:30:53], I'm from [Vietnam? 00:30:55], Registry of [.vn? 00:30:57]. I'm not a technical person so maybe my question will be deemed to be not relevant, but I would like to know if [Veitnem, .vn? 00:31:10] Registry, we are not the EPP based yet, so should we do and what can we do with Data Escrow?

DAVID KIPLING: Yes, so the Data Escrow process involves currently for the gTLD process, to try and make things consistent, it is around a schemer of set data. That data can come from various systems, various areas and can still be Data Escrowed and verified. The aim would be to try and make it as consistent as possible but all Registry data can be Escrowed.

EBERHARD LISSE: EPP means data input, but it also requires a certain format, so your question is probably if your format is not identical to what ICANN says… Basically, it must look like what the Escrow format is needed to reconstruct your Registry in case of failure, and then you send that. You don't necessarily have to use the ICANN format. If you can get your data

into ICANN format, you can use an ICANN accredited provider – if you want to. Or you do it yourself and you don't have to.

The point here is that you want to have an independent of your operation if something happens; another hurricane goes through and everything blows away, that you can, with relatively simple methods, even on a different hardware, on a different software, reconstruct your system, in a relatively… Okay, thank you very much.

DAVID KIPLING: Yeah, I think the key point as well from that is is it having that impartial third party who actually has the data and verifies the data. So again, that's instilling the trust in your customers that they know their data is safe. But yeah…

EBERHARD LISSE: Thank you very much.

DAVID KIPLING: All right. Thanks very much. [applause]

LUIS DIEGO ESPINOZA: Good afternoon. I want to share with you some of our experience with the use of TPM in DNSSEC signing in .cr. The TPM is a chip included in most of the actual hardware by Dell by example, the laptop computers, and it uses it to encrypt some information within the equipment. This chip is very interesting for us. We are taking into account the possibility to use it for signing and we do it… We really do it and we do it well. And

this presentation is really to show you some results about the use of TMP for signing. Okay?

What the motivation was for us to use the TPM? Well, the first thing I must comment to implement the DNSSEC, the country code TLD, using a small TLD, less than 15,000 Domains, and I think the Domain Name resolution for our country is as critical as any order of contra-services. Then there must be trust, because we're making a resolution for government entities, for financial entities, industries, education, then it's not something you can take lightly. You must take it seriously. And with DNS resolution normally we take it away. Then with the DNSSEC we want to take these kinds of measures.

The trust of the [inaudible 00:36:08] on the follow best practice and [inaudible 00:3:13] procedures. This comes from a certificate authority, from BKA system, [Word? 00:36:21]. And the other thing we need to find some solution is because we try, the first time we tried to create some DNSSEC keys using the random on the server, and it takes a very long time. Then somebody told me there's some kind of chip in the computer that can bring some type of random number generator, on hardware, that can facilitate this creation of keys, in the safe way, because it's in hardware.

Then I started researching this chip and the first intend for this TPM chip was the RNG included in the email, in it, because it's hardware based; the random number generator on this chip is hardware based and that is a good practice, that is a good thing to generate keys. Generally, there will be more. I've found that there are some implementation of PKCS

#11, the token where it's basically the RSA keys for DNSSEC signing, that use the implementation for this chip, TPM chip, not only RNG.

So, this software, this is open-source software [processors? 00:37:56] for .net, it's created by IDN, provide [inaudible 00:38:02], PKCS #11 device, that is encrypted using the key within the TMP chip. The TPM chip is included in the existing Dell servers of ccTLDs, we don't need to buy any new thing, then we tried that for free. That's the thing we start working. Then Richard Lamb, from ICANN, liked the idea and put the things together and put it working, really, because there's not a lot of [recommendation? 00:38:45] how to create these PKCS #11, talking using these [processes of recommendation? 00:38:50].

There were some tricky things but at the end we can make it work. Thank you. Well, about TPM, this is encrypted hardware, this compliance FIPS-140, level two. It's not only level one, it's level two. It's supported by open-source software. This bit in fact, in the production environment, is more or less one RSA 1024 bits signature per second. But [currently? 00:39:29] the chip can bring ten signature by second and effectively half a building hardware RNG, random number generator, and the PKCS #11 interface simplifies the migration to real HSM or to complete HSM. Then in this way we want to go for the solution, we work a little bit. And for the ICANN meeting in Costa Rica we want to launch the DNSSEC signing of .cr, using this chip.

Well, this is a framework of the way we can use the TPM, using the TrouSerS libraries and the Opencryptoki. Opencryptoki are libraries for PKCS #11, and the thing is, within the chip there's only one key, there's the storage root key, the SRK. The process software, what it does is

create a [built in? 00:40:46] token in the hard-drive, but wrapped by a hardware key. In this way the PKCS #11 token is protected by the hardware key.

This hardware key requires for example to reset, or if you fail three times with the password. The only way is having physical access to the server, because the only way to reset the TPM is from the BIOS, from BIOS interface. This is in force to have physical access to the server to reset the key or to reboot or to clean the key, and use it again. In this way, because we have the server at our office, not collocated, we can play with that, we can play a lot with that and we were thinking about if we have the servers collocated in some [dot? 00:41:54] center or rented or in the Cloud, and that is not possible to play with the TPM in that kind of server, because you don't have access to the BIOS of the server.

Basically, that is the framework of the PKCS token, using the TMP chip to protect the keys. Well, after this instruction the rest of implementation is just the same as the implementation of generating a key using software; the same implementation. We use the same comment, we use the same DNSSEC tools, because using the [inaudible 00:42:45] libraries, it creates the instruction of PKCS #11 to work with the TPM. And initially, in this environment, in the future if we want to change to HSM, it's very easy to change it because you only replace the driver of the Opencryptoki, and instead of using the TPM you will use the HSM.

Well, our first implementation of the use of TPM wasn't in this environment, using the real zone files. We initialize the TPM on BIOS, we'll show you the screen where it's in BIOS, where we initialized the TPM. The first try was unsuccessful because in this implementation you

have as many slots as you want, it's like a software implementation, but the only slot that you protect with the TPM key is the slot zero. Then we tried many times and all the slots enabled is one, two, three, but zero was very difficult to initialize because the procedure to initialize this will be followed exactly in this specific order.

If you miss something you screw up. Then after a lot of work and at some point we were near to giving up because it doesn't work for weeks. I think it was Richard Lamb that sent me an email some day and he said he kept putting in work, then it was possible to initialize it with some tricks and the procedure was very specific. And this shows you, this in PKCS #11 to how we can list the slots, and this slot zero this is like… You can see in the token flags there's an RNG login required to initialize. The slot zero is protected by the TPM chip. You can see after the slot zero at the end it says 'TPM'. But this little one it says 'soft' this is a software slot.

Well, okay. Yes. In November 2011 we put some production environment. After initializing the TPM, [usersoft? 00:45:35] created some signing key and key signing key for small zone, in this case we used the zone sa.cr, it's a very small zone. We signed it with DNSSEC, using the binding tools. The key ASCII and zone signing key was generated inside the server. The PKCS #11 backup is [per annum? 00:46:02] modified but by Richard Lamb, that can be used easily to manipulate the keys, using the TPM crypter.

The Opencryptoki was configured to use the TMPD, this driver for using the TPM, to protect the keys. After a week we resigned all the list of zones in the Domain, each hour – it has a procedure, each hour the zone

is reloaded and resigned with no errors. Then we decided to sign on some zone and on Top-Level .cr for the TLD. And we let that on a production environment.

EBERHARD LISSE: You are five minutes over.

LUIS DIEGO ESPINOZA: Okay. Yes, this key management is ideal but it's not what is in production right now. In the PowerPoint this is supposed to appear the things in some order to show what happened, but the idea with the [hold is? 00:47:17] is we can create the keys with all Dell laptops that have the chip inside, offline, using a ceremony or something like that, and then you can move the keys created by that TPM to us, to the production server, on a USB key, for example, following some ceremony, and you can load it again on the server and re-wrap it by the TPM on the server. This TPM is very different.

In this way, this is a good follow of some good practices in this environment. [clears throat] This is a string of one of the servers, the TPM on BIOS, you can see the meaning of security, if the server has it there will be TPM security, a TPM security tap. First you turn on the TPM, you need to reboot the machine and after getting inside the TPM again you should choose the second one; activation, and activate the TPM.

After you do this you can put it with Linux and use the TrouSerS, the TPMD to manipulate the TPM. The TPM too has some TPM [initialize?

00:48:49], TPM token, in the way you can create these. The PKCS #11's token. Well, some results, some numbers. The eight different zones, all the zones are buyable, signed by this mechanism; .fi.cr, .go.cr, .ac.cr, etc. There's nearly 400 signatures each hour. I know this is a small number of signatures because it's only a few Domains that sign it, but the thing is, the system has been working now for 12,000 hours; that means 4.8 million signs, and no errors from when we start using the system, one year and four months.

Right now, the process takes 15 minutes to sign all the zones. First, we sign all the sub-zones then generate the error sets included in the parent zone, and sign the parent zone. It's a little bit slow, yes we know, but for that size of Domain and that size of keys to sign, around six to sign, it's enough. We were thinking about why is it so slow? Then Richard Lamb did a presentation in Prague about the funny device he is building with the TPM, and he can obtain from this device I think five signatures per second, something like that. Then I tried something very simple, it was running the signature in parallel, and I ran two processes of signing and it was surprising that the time of signing decreased.

I don't know, I don't have an explanation for that, but the time of signing decreased, in this way. If they normally give you around seconds to read all the signs, it takes 15 minutes. If I start four signs in parallel it takes something like five minutes, something like that. It looks like this library can be optimized in some way to provide faster access to the TPM signing. Some conclusions. For a small zone or at least for a few signed records, it could be a big zone but you don't have to need to sign too much records, it is possible to do the TPM in production envronment.

It's a very low cost; free, the hardware is free, and it's easy to access and crypt the hardware. One of the things, one of the problems we have when we start working with this, we are trying to look for some Smartcard reader, and some Smartcards to use that kind of hardware for encryption, and sometimes it's very difficult to find that kind of thing. You can only get it online and sometimes it's not that easy. But this hardware was very easy to find because it's good in the servers. In my opinion it's very reliable, according to this proof. It's fast enough because it's hardware again, and it provides the initial phase to migrate to an HSM environment after that.

The TPM is another cryptographic accelerative, this is important. It is not a cryptographic accelerator, but this is a good example to put some technology there to work for us. Okay. This is an [inaudible 00:53:02] two by one. [laughs] Then this is my half part. Do you have any questions?

EBERHARD LISSE: No, we don't allow any questions now. We have got seven minutes [laughs] for the second presentation, are there any questions? If there's any implementation questions, maybe take it offline and take to Diego around this week.

LUIS DIEGO ESPINOZA: Okay, this is a very different topic, but okay. The Security Incident Response required created a Working Group within the ccNSO, to create a contact request for implementation for Incident Response. Then this is the people in the Working Group… Okay, the Computer Security

Incident Response is a service like in that diagram, and it has many, many servers. This is like a first of many orders, this institution of working with the Incident Response. The idea behind this is to create a repository of contacts of TLDs that can help these organizations in case of emergency or in case of a security incident, to contact the appropriate persons in each TLD in case if needed.

This is the common services of our CSIRT. You can see there are some Reactive Services, Proactive Services, Security Quality Management Services. Within the Reactive Services there are some incident handling, and within the incident handling are the points of contact. Then in this point is where this Contact Repository will work. This Incident Handling service is describing some of the services. And in the services of interaction, they need that point of contact, and to have this we need that Contact Repository. Then this Working Group is working on a Contact Repository in this level.

The Contact Repository Implementation Working Group was commanded to explore factors to implement, maintain and operate the repository, provide some funding models and governance model. We split the proposals into sections. One is the database, the system by itself, it's like a directory service system, it's like an LDAP or something like that; a special agenda or a special address book. Then the other part of the Contact Repository is up-to-date each contact information.

And for this we need something like a service provided by Contact Centers. They must call the clients and check the information is valid, prove the email is still valid and the phone is still valid. These kinds of things must be proved with certain frequency. And in this chart we have

the relationship of the many components of the Contact Repository, and in this area is the response entities to be the third, first, any order response team, that within the possible contacts to search they will have a link to access the Contact Service Center or the Directory Service in the Contact Repository for TLDs; specifically for ccTLDs.

And this information is collected from the ccTLD Registries or ICANN or ccNSO and in this way the idea is to create this Contact Repository to provide information in case of incidents. The Directory Services is very simple in the way you need to search contacts, create contacts, update contacts or delete contacts, and we are proposing to use something like LDAP or that kind of protocol to access the contacts.

Okay, maintain and operate and keep the contacts updated requires 24/7 operation, email response management and Web chat, session recording and transcript mailing, self-service knowledge-base, analytics and quality system, telephony infrastructure and IVR. This is the common theme required for a Contact Repository Contact Center. And this is the proposal of how frequently there will be tests of each of the methods of contact for each ccTLD.

The thing is, if we have 200 TLDs in the database, we'll have 400 contacts to contact by each TLD. And this is more like a distribution to have all the contact methods proved around three-month cycles. By now, the proposal is to create a Standing Committee within ccNSO that will report to ccNSO Council and users and the Standing Committee would be supporting by ccNSO Secretariat and with experts from SSR Department.

The Standing Committee will be responsible to manage the Service Provider and Agreement Compliance. The relations with subscribers and maintain the use cases of the repository. The proposed funding models could be one of those but if there's more then we are open to listing more proposals. A uniform subscription and set-up fee – this means a flat fee for each ccTLD. That could be expensive and prohibitive for some small ccTLDs, then this model is maybe not the best. Cross-ccTLD funding – the ccTLDs use more of this system to pay more and in this way they are covering the others that pay less.

ICANN funding or a mix of the before mentioned matters. Next steps. Send the Request of Information to potential providers. The idea is to have input, feedback from potential providers for more specification, adjustment specification and to have an idea of the possible costs. Right now we don't have a real idea of the possible costs of this service. With the feedback, the idea is to compare this Request of Information to a Request for Proposal, sent to the ccNSO Council, then once approved, send it to the potential bidders to have some bids. And analyze and recommend the best offer for the ccNSO Council. That's it. Questions?

EBERHARD LISSE: One question. Bring the microphone.

UM: [inaudible 01:01:25]. I have a question. What's wrong with WHOIS?

LUIS DIEGO ESPINOZA: What's wrong with…?

UM: WHOIS.

LUIS DIEGO ESPINOZA: WHOIS. The extra protocol of WHOIS is very limited in the form of… Well, it has some limitations about the indication, of the [retation? 01:01:48] and the scheme of contacts what give us, because it comes from the past Working Group, is bigger than can be held by the WHOIS normal schemers. We need to maybe modify the schemers to provide this and the other thing about the… I never think [laughs] about… I never think about WHOIS, I don't think about that, but well, the only thing I'm thinking is this database is not for access for anybody, this is security access.

UM: If you want to contact a TLD you just look it up on WHOIS, right? There's a technical contact, there's an administrative contact and they're supposed to be responsible.

LUIS DIEGO ESPINOZA: Yes, but the deal with this is it's supposed to… It's more on the side of Security Incident Response than… The security contact maybe is not the same as the technical contact and based on some… Well, and the other thing is many of this information will not be used in an electronic way. Maybe some information will be a telephone number or a fax or something like that. Then sometimes we think only in the electronic

world, but in this case, in case of disaster by example or in case of serious emergency, the email could stop working and this leads to…

It's necessary to have the chance to make a call or something like that, then it's not written on a stone. The protocol to use it is not written on a stone. Where we think it could be useful is because it's more oriented to like an [inaudible 01:03:50] book, but it will depend on the bidders; what they offer.

UM: The one problem with the context is that IANA is managing this context and they are not really, for example there are some places where they don't want to change the technical contact or the admin contact or the Registrant, so this is a separate issue where a third party is unrelated. If I want to change my telephone number, I may or may not get IANA to do that. If I want to change the name of the contact, I may or may not get IANA to approve that. Sometimes it doesn't work.

So the idea probably is to have a separate entity that is not dependent on a third party, where I can put the ones that I think today is managing my emergencies, and tomorrow is somebody else. I can handle it, I am in control of the data. I have not got to go for approval from a third party. That makes using a technical contact from the ccTLD, WHOIS database is a little bit more difficult.

LUIS DIEGO ESPINOZA: Yes, all right.

EBERHARD LISSE: Good. We are running a little bit late but we are doing find otherwise. Our usual host presentation, Jiankang Yao will give us the usual host presentation and of course talk a little bit about Internationalized Email, because that's the big thing around here.

JIANKANG YAO: Hello, my name's Jiankang Yao from CNNIC. I'd like to firstly give a CNNIC brief, then we'll talk about Internationalized Email. This morning we have already talked about email implementation from our partner, Coremail. So firstly, a CNNIC brief. CNNIC is the main Asia-running [policy? 01:05:47] China. We also allocate IPv4, IPv6 address, so we also do some research; DNS software and hardware. We also [inaudible 01:06:00] IETF activity, also CCSA activity or China Standard Association. We also do some research related to security: DNS monitoring and analysis. Also we do some anti-abuse of Domain Names; so we check which name is for [inaudible 01:06:25].

We also released the China Internet Development Report. Also we have some consultation services for Internet companies. For example, in today's opening ceremony, how many Internet users in China? It is actually the data comes from CNNIC. We also [partner? 01:06:56] ISOC, ICANN, IETF, APAC and CDNC is the Chinese Domain Name Consortium. We also with other companies, have a joint project and joint labs.

So this is our CNNIC management architecture. This is the Ministry of Industry and Information Technology – MIT. So this is our business instruction. We also have a Steering Committee we have oversee the CNNIC. CNNIC also has a Chinese Computer Networking and Information

[inaudible 01:07:48] Chinese Academy of Science. So these are the Administrative Governance, so our Human Resources belong to the Chinese Academy of Science. Our business belongs to MIT.

Also we have some Domain Name information relates to abuse management from the National Internet Information Office. So here is the [inaudible 01:08:19] .CN registration volume until last September, it's 6 million – now we are around 7 million. So for the China registration is around 300,000. So everyone knows that China now has a very strict Domain Name Registration Policy. So first, we have a Customer Register, or registration of Domain Names online, so then customers [send me? 01:09:05] the data, such as ID information via Web pages, fax, email or by post.

So for example their company's information or their personal ID information, so then our register will check the information to confirm that their ID is true of false. So then CNNIC will recheck the data, then we have finished the Domain Name checking. So currently before [2010? 01:09:50] in January, there were only around 18 personal data, Domain data which was true, but now we have almost 99, or more than 99 Domain data which is true.

WHOIS information [gets far more China? 01:10:08] true data. So this is our operation-infrastructure of our service platform. These are our Internet… We have a main, primary Operations Center in our Chinese Academy of Science. We also have a remote, secondary Operation Center in Chengdu City, Sichuan Province. We also have a local, secondary Operation Center in [inaudible 01:10:42] County. We also

have a lot of servers in different ISPs. Most of the major ISPs we have servers there.

So we also deploy our sever globally, they are in North America, Europe and also Moscow. There are also a lot in China and in Tokyo. In the future we'll also be in Singapore. So we also have SLA operation requirements. There is some data. This is WHOIS data and we can check it from a CNNIC website. So for DNSSEC operation, DNSSEC deployment; we plan to deploy this technology this year. We also have simulation system where for one and a half years. So our technology I think is ready and mature enough to deploy. So these are some designs on DNSSEC deployment.

We also drew some Root Mirrors; I-root, F-root, L-root. We also put some machines to them. CNNIC also have a New gTLD for [inaudible 01:12:39] .company or .network. We also have some swap server, such as DENIC and KRNIC. We also provide some recursive name servers for Google it is 8.8.8.8. Now we have 1.2.4.8, so this is also very easy to remember. If you get some check DNS's resources from… These are public recursive name servers.

So we also have anti-abuse for some .cn names. We also build our organization Anti-Phishing Alliance of China, founded by CNNIC. So we have to check a lot of Domain Names. Also, other .cn names are checked by these systems, we found a lot of Domain Names websites, we always report them to our government; our government will block these websites. So there are many, many websites that do phishing. For example, China have a bigger company, a bigger [inaudible 01:14:25]

called ICBC, so a lot of phishing [banks? 01:14:32] learn something from ICBC of phishing of customers.

So we also have phishing…. [charge them, we also dewire? 01:14:57] some phishing websites generator technology. We have designed a system which will automatically detect whether websites are phishing or non-phishing. So we also do some DNS data mining and phishing reports; phishing URL automatic generation technology. We also do some national IP address location… Now, CNNIC locate around 30% of IP addresses in China, so this is the number of IPv4 address.

So we also report some Chinese Internet Data Analysis Platform. For example we report how many Internet users are in China, how many persons use blogs, how many users use Weibo, which is actually similar to Twitter. So we also do a lot of analysis. We also give some reports to the Internet company. So these are the [inaudible 01:16:49] data; there are for example more than 10 million TLDs, active websites, around 2 million. So 1 million are .cn websites.

W also do some research about IPv6 Application Pilot Center. We have a platform somewhere for one to two weeks, [inaudible 01:17:28] IPv6 technologies, they can access our platform to use our Internet IPv6 service. So if you are interested you can log in at www.6pilot.cn. We also have some EPP Registry systems, so there are our clients, there are our Domain Name Resolutions, there is our Registry Service. So if you are interested you can log in at www.cyberspace.cn; where you have Domain Name users.

So we also do some research on DNS server software, similar to BIND 10. So our name is Zebra, so we already have such a, something which is [inaudible 01:18:45]. So we also passed the BIND 10 project, these are some photos. So we also have some hardware, some anti-DDoS device. So this is our public recursive DNS services development. This is our IP address for an Internet user. So we also have somewhere where we monitor DNS and also do some DNS analysis. We have an Internet technology research in IETF we have many DNS-related and Internet, Domain-related technology. Also, IPv6 also have some Internet [scenes? 01:20:04].

So in China, as a nation, we also do some research related to it. So now I've finished, [inaudible 01:20:16], now we're on the main part. Okay. We will focus on Chinese Email Address – Internationalized Email Address – progress. So what is a Chinese Email Address? Now, for example my name is Jiankang Yao, it is a Chinese name. we have Chinese English email address and now we can have a Chinese email address where both the email addresses will point to the same actual, physical data.

So these are Chinese emails. Why should we have Chinese email addresses? And email address is an essential element of a business card. For example, everyone has a business card where one part is Chinese and another part is English. But both names point to the same person, so Chinese email address and an English email address is a similar thing. So this is a… The standard already published, International Standard, so now if you upgrade your email server, you can use an email address like this one.

This shows our progress. IETF, so our Working Group was designed in 2006, in February 2012 they were fully published. This much or another four [inaudible 01:22:16] published now your email address internationalized has finished [inaudible 01:22:25] publish, so it's a major [inaudible phrase 01:22:33] is RFC 6055 and RFC 6531. [Why in CNNIC inaudible 01:22:44]. I am also the cause of these RFCs.

So many companies such as Google, Qualcomm, Sun and NTIT are also involved and helped participate in this project. So these are many of these. So these are RFC document plans, we now have all of these partners. We have already finished. In the future, after the deployment of Internationalized Email Addresses, maybe we can have more of these, otherwise there are these. So in China we have published the IETF Civil Standard to push the Civil Standard for CEA in China.

So [eisenger? 01:23:37] Chinese Email Address will have a brighter future, now we have more than 1,900 New gTLD applications, as many Chinese New gTLDs is the topic, so New Chinese gTLDs for Chinese Email Addresses. So Chinese Email Addresses will promote the registration and the use of Chinese Domain Names. So these are… We can maybe… We are convenient for Internet users, especially other people surfing for the Chinese, for their Internet.

With so many users, according to our report, there are many interested in Chinese Email Addresses, because many companies, especially small enterprises are interested, because the [inaudible 01:24:50] business, or offices, are exchanging information using email. So now we are ready for Chinese email commercial deployment. Our partner, Coremail, will release its commercial product this month for they will have [press?

01:25:17] this month. So Internationalized Email Addresses, Chinese Email Addresses may trigger new applications, may provide a new business opportunity also maybe push GDP growth because we must upgrade our server and maybe buy some new hardware.

So this is the CNNIC position. CNNIC now have [inaudible 01:25:52] IETF standards, so we will push the new Chinese applications for Chinese. We'd like some cooperation with some other Internet companies. We have cooperated with Coremail and in the future we will work with companies such as QQ.com, who have 400 million users, so we have talked to them. In the future we hope to have a lot more cooperation, so last June we had an event. First the Internationalized Email Address sending event, so CNNIC, TWNIC, HKIRC, SGNIC, .ASIA, MONIC, AFILIAS and VERISIGN joined together to have the demo sending of EAI address. So these are Chinese Email Addresses sending, these are the replies and these are future… Singapore, Hong Kong.

So these are also [inaudible 01:27:12] in the opening ceremony, so CNNIC is ready for the Chinese Email Address System, for example, Yao Jiankang is my name, [inaudible 01:27:19] means CNNIC.china, [inaudible 01:27:33], now we have the user's email address. So our Coremail partners, the first Internationalized Email Address commercial product to be released, so this is in April 2006 in Beijing. So that's all. Thank you very much. Any questions?

EBERHARD LISSE: Thank you very much. Any questions? Come one, now, there's one.

UM: I will speak Chinese. Translator? Okay. [Speaks Chinese].

JIANKANG YAO: Okay, I have a translator to translate his question. His question is, in Chinese, does it have Chinese data, and English data. He asks how to translate Chinese data to English data, because he can't understand how it only supports English data. So this is a software problem. It is not a protocol problem, so in the future the software can automatically change or adapt a change; change from Chinese data to English data, so this problem, also in IETF already, [inaudible 01:30:01], so in future as more Chinese Email Addresses are used, I don't think it will be a problem.

EBERHARD LISSE: I've just got one problem. How many Staff does CNNIC employ?

JIANKANG YAO: Oh, this year around 260, plus some research students.

JEAN GLATU: [Jean Glatu? 01:30:39] from [Sira? 01:30:40]. So I've got a question about Zebra. It was in your DNS server software slide. So is that an open-source DNS software that you built?

JIANKANG YAO: We have DNS software built, but it is not open-source DNS software, it is commercial. Let me… Yeah. See, this is Zebra, this is a commercial

product, not an open-source product. We also sell this product to our customers, some Internet companies.

EBERHARD LISSE: Okay, thank you very much. Next presentation will be Frederico Neves from .br, talking a little bit about security issues they've encountered and how they dealt with them.

FREDERICO NEVES: Thank you very much. Good afternoon. As Eberhard said, my name is Frederico Neves, I work at the .br Registry and I will talk a little bit about, actually it's not the .br Registry in general, but I will talk specifically about Registrant credentials security at .br and in any other public service on the Internet, in general. So to give you a little bit of information about what we are dealing with here in our case. Our Registration system provides services for the .br Registry, direct customers and customers through Registrars, using our EPP interface.

And we are the NIR for Brazil too, providing ASN number and IP address block, IPv4 and IPv6. So the size of the problem is around 3.1 million Domain Names, around 2,000 ASN numbers and 1,000 IP blocks. 2.7 million of these objects are administered directly through the Registry, on a kind of Regsitrar interface to direct customers. So we have 1.4 million accounts taking care of all these objects, and 500,000 Domain Names, around 60% of all Domain Names are administered through 60 Registrars; mostly Brazilian companies or large ISPs in the country.

So during this morning we had some numbers from the history of the track record of security in Registries or Registrars, and other services in the Internet; I provide a list here so I will not enter into too much detail here, MarkMonitor will probably talk a little bit about their experience. So diving directly on the credentials and storage. Here I have some ways of doing that and I will list all of them here actually. Historically you could use the Clear Text to store the credentials of your users, and these days even some old financial institutions still store credentials this way.

Sometimes when you request password recovery you receive your password, so it means that they store it as a Clear Text. So the second option is Cryptographic Hash and it's definitely a better option but these days, with the current CPUs that we have, this is not too strong, because even with current Notebooks, the ones that you guys have here, we can do more that 200 million hashes per second and so you can pre-compute a database of password hashes for most size passwords and just look up some of these that are [inaudible 01:36:29] and it would be trivial to recover passwords if you compromise the credential storage.

The other ways is Salted Hash; it's much better but even in the case of LinkedIn compromise with the 6 million accounts compromised, we have this with another technology that is available these days, the GPUs, and we have some data telling us that 90% of all these passwords that were stored with Salt could be compromised, because these clusters of GPUs, this number is available here; 350 billion hashes per second you can do with this; this kind of specialized hardware, so it's still possible to break if you get your hands on this credentials storage.

The other ways; Salted Adaptative Hash, is that it's the way that some modern operating systems store their passwords these days. These hash functions are much more… They use many more interactions or the algorithm uses too much memory or… But it's very secure, but in some sense of secure, but to use only this for public Internet service could be a threat for you, because with a really small amount of requests per second you could seize them in a kind of [denial? 01:38:20] service.

So you need to take care in selecting one of these kinds of adaptative hashes. And the last option is to use some kind of adaptative hash and up to that, encrypt the information using symmetrical keys. So another point to tell you regarding the promotion of good passwords, because even if with salted hashes, simple salted hashes, if you use mid-sized passwords or mid-sized passphrases, it's unfeasible to break that, but for most sizes like less than ten characters, it's trivial. That's why they could break only 90% of the LinkedIn hashes.

Another option is actually using Two Factor Authentication. That is something that you know, your password or your passphrase with your identification and something that you have; a token with the one-time passwords. These days, currently, we already have an unencumbered available technology at the IETF. It's the HOTP and the TOTP in this to RFCs. Basically it's a one-time password schemer based on HMAC and a short secret.

The HOTP is based on a sequential number and the TOTP is also a sequential number but it's based on a temporal series, and you only need to store the state of this sequential number, so… And this is what we are using now in .br, we call it an ASM, it's a kind of Authentication

Security Model because we do total decoupling of the frontend systems through the storage of the credentials. Even for passwords and even for the Two-Factor Authentication.

So after you do the provisioning of the Two-Factor Authentication or the password to the end user, the frontend system doesn't take contact with this information any longer, and the state is storing normal RDBMS, secret [basal? 01:41:23] ones, and they don't need to be, obviously, the same database that you use for your frontend systems. In our case, the short secrets are derived using an HMAC and a Master Secret and the id, in this case, of the Two-Factor Authentication.

The password credential is protected using a Symmetric Key. This is a… Basically what we do is, we do an Adaptative Hash and later on we use a Symmetric Key to restore these on the database. Both of these keys – the Master Secret for the Two-Factor Authentication and the Symmetric Key – are protected using a shared secret chain and a schemer, and this is required for the activation of this ASM hardware.

And besides that, obviously you need to do Rate Limiting. And for any of these authentication operations, you need to do some kind of Rate Limiting by source address, by source block, by the identification of the user and basically we use the Token Bucket Algorithm, that is very popular for Rate Limiting Networks, and this state is stored in a key-validated database e-memory called Redis; it's a very fast one.

And for the Two-Factor Authentication we are basically leveraging the fact that these days most people have Smartphones, so we are using an App, and it's a Google App; it's called Google Authenticator, at least for

an Android and iOS, and for the Windows Phone there is a port called Authenticator that is based on the open-source software of Google – it's quite popular.

The provisioning of the shared secret – we leverage the fact that the user is already authenticated in the system and we have a secure channel with him over HTPS and we provision the shared secret using a QR Code; these applications can read this and at the time of the provisioning, the user reads the shared secret in their device and it already provides the first authentication code for us, and then we know that it is ready, provisioning the handset or any kind of device that can support this technology.

Then, afterwards we provision them with a sequence of OTPs, that is the last resort if they lose their handset or if they need to reset their account, and this is a list of OTPs that they can print and store in a safe place. And it's basically a short presentation; only talking about this and regarding the ASM software, we will provide this… Actually, the implementation that we did is written in Go and we will provide this, basically, for the Registrars in Brazil to try to cover the 500,000 Domain Names that we are not dealing directly with.

EBERHARD LISSE: Okay, thank you very much. Any questions? Going, going, gone. Oh, there is one. Luis has got one.

LUIS DIEGO ESPINOZA: A quick question about this Token App. Do you need to synchronize something with the server, right?

FREDERICO NEVES: Synchronize time.

LUIS DIEGO ESPINOZA: Only time?

FREDERICO NEVES: Only time, yes, because actually it's used for the TOTP, not for the HOTP, so the time needs to be synchronized. Some synchronization is needed but the time synchronization of the GSM that works is good enough.

EBERHARD LISSE: Which language did you say it was written in?

FREDERICO NEVER: Oh, the SM? Go. Go, 'G O', yes. It's a Google promoted language, yes.

EBERHARD LISSE: Okay, yet another one. Not another Google language but yet another language to learn. Yes. All right. Thank you very much. The next one will be Jaap Akkerhuis. [applause] Do you want to go on, yes?

JAAP AKKERHUIS: No, sorry, it's so I know where I am. So I know where I am. So that's… So I don't have to…

EBERHARD LISSE: No, no.

JAAP AKKERHUIS: So I know what the next slide is. [background chatter]

EBERHARD LISSE: Hang on, we're just downloading the presentation. It will take a few seconds. You just start talking, I'll just…

JAAP AKKERHUIS: Okay. Let me first… I'm Jaap Akkerhuis from the Netherlands. I used to work for [SEDN.nl? 01:48:45], but moved to NLnet Labs and we are interested in security, especially name servers – some of you use our name servers – and I've been asked to give a quick overview on what this is, about RRL, and the background of it. Okay. Well, this is an overview, and I notice that actually points three and four are the other way around, so don't… But it is better.

First of all I'll try to explain what amplification is, what reflection is and then what are the types of things you can do to be [friends? 01:49:42] to harm you. [mumbles] So, everybody knows here how DNS works, so I won't really go into the details of this picture. The only thing it says is you've got questions and you have answers, and that's where the gap is, because the problem is that a small question can cause big answers, you

know? If you ask a politician whether he is lying, you get an enormous story and not really what you are really asking for, but anyway… And the reason why… And the other thing is, the questions – I don't know if you know UDP protocol? UDP has no built-in notion of any authentications, but TCP you do have to have because you first do the handshake; 'oh, I want to talk to you!', 'oh okay, give me your address!' stuff like that, and then you do the actual data exchange.

So there is some form of authentication there, with UDP there isn't. And that makes DNS a target because there are so many DNS servers out there. I mean, you can abuse a UDP and DNS is a big user of UDP protocol, so… What is the amplification? Well, the amplification is how big it can be. ANY Queries, okay, you can really take and make them really big. ANY Queries give you some information about this name; it doesn't matter what it is or how much that you have, and normally it comes out of caches and stuff like that.

And if you start to play with it you can see that easily every byte that you get back is 18 times, for every byte you get 18 times back, which is kind of a lot. So you sent a Query of 100 bytes, you'd get a lot of kilobytes back. But even if you don't give an answer, you can get a lot of amplification factors. If you say 'I don't know this Domain next to me' and you have DNSSEC switched on, you can get it. With NSEC you get about effectively 18 times a bigger answer for every byte, and if you do NSEC3, that's even more depressing because you get 25 times amplification factor if you do it properly.

So these are not small figures. So [coughs] reflection and spoofing is actually making the answer go to somebody else then [inaudible

01:52:35], so just lie about who to talk to. I ask you a question and say 'oh, but the answer, [inaudible 01:52:45] really lost the answer' and so everybody [inaudible 01:52:49] an answer. And it's very easy – you replace the original address with the one of the victim so that the answer goes to the victim. That's the case; this is all what it is and nothing else.

The two ways of getting these reflections and amplifications done, the easy part is just sending the questions to open and manage resolvers. Now, there are also open/close resolvers out there, but people actually know what they're doing; like Google, some other people, open DNS, stuff like that. They're actually checking what they're doing, but there are an awful lot of open and manage resolvers out there. People are installing whatever the latest version of – the not latest version – of whatever the name server is, which concentrates the machine of whatever. Most of these things are just open for the guests of the world.

Somebody did a little research two weeks ago and found within a week, 21 million open resolvers. So there is a lot of instruments to hit somebody, and so what you want to do is… This is not an easy way you really want to do it, instead of being hard to other people, and it saves you bandwidths, I mean people are misusing bandwidths. Also, you get a bad rap about being able to help people attack other people, and so it saves a lot of headaches.

And for people who want to read about how to do that, it's a recommendation, RFC 5358 will tell you all about the gory details and [coughs] also, you might want to have a look at SSAC or [RR8? 01:54:55],

which is a higher-level description about what to do with open resolvers. This is the easy part. Now we get to the more difficult part because authoritative resolvers have a slightly different way of operating. By nature they actually have a wide audience and by nature they actually should answer all the questions they get, because otherwise you don't get an answer! And the game plan for the TLD and authoritative, is to get an answer about what's happening.

And it also serves a lot of many recursive name servers depending on these answers. So you have to be slightly more careful about doing something with the traffic and the traffic… Things that can happen is – and this can be hard – the little traffic you see is one of our main servers; this is a TLD we serve, and you see there about ten times the rate, and somebody is attacking somebody else, in a different country – we're just innocent bystanders here and so this costs a lot of money.

Also you might actually find a lot of problems like being sued because you are instrumental in doing attacks to the harmless victims, and there are a lot of reasons why you'd want to do something about this spoofing and amplification. [coughs] And as I said, with authoritative server principle is that you don't really want to help the attacker [coughing] you really want… And that's why you have to keep answering some queries to well-behaved clients, to well-behaved… Because if you just stop answering queries at all, the attacker has it's call and the victim is off the net.

The other thing is, if you do all these mitigation efforts, you really don't want to have a lot of influence on your normal servers. It's not very nice to see the performance of your server degrading, even when you're not

under attack, but especially when someone else is under attack. Now, a couple of ideas people have about doing that. At first, they can do something to try to limit the amount of queries. If somebody is asking the same question all over again, just ignore them; so you can DNS firewalls, take some heuristics and stop answering their queries.

Now, there's a problem with that. There are actually legal reasons why you have a lot of the same queries. I mean, considering that people have behind [inaudible 01:58:11] nets and so actually, a lot of different people are using the same address and you don't really know whether it's the same person of a different one. Another thing is suppress ANY Queries; ANY Queries are actually only used… They were initially meant for doing deviant purposes, so applications which are using ANY Queries, look at what they deserve, but it's not really nice and you never know which applications are completely dependent on it.

So the other thing is, what's known in the world as DNS dampening. DNS dampening is basically saying 'I keep a test hold and if I give the same answer multiple times, I just stop answering for a while until it goes down again, and then I start again'. Well, the problem there is of course if you just keep on with the attack the factory will be the same as doing a query limiting; you won't answer at all. One of the reasons for doing that is for legal reasons. [Donna Acker? 01:59:30], who actually did some of this stuff, does that because of these reasons. He doesn't want to have any chance of being short about being instrumental in an attack.

And one of the most sophisticated manners is to do Response Rate Limiting. You actually look at the responses you give but you kind of

limit it to an acceptable rate and that later part is [inflamed? 02:00:07] by [FIXI and FERNL? inaudible 02:00:06]. And at first implemented for BIND and now also available for NSD [Norsk? 02:00:20] and probably some others. So basically, you drop answers that exceed a certain [trestle? 02:00:29] limit. If it's getting really high-end, just don't do the answers.

And how to deal with the false positive; meaning that you actually are permitting the victim even more by not answering at all. This should be a fallback mechanism. Well-behaved clients actually know how to deal with TCP, and so what you do, as we said, TCP has some form of authentication, so the address-proving cannot really take place there, at least a last problem, and it allows to give some service to the victim so you don't really kill him completely.

And the nice thing is, it performs reasonable well, and the impact isn't that big. We did some measurements on it and so we basically set-up a bad guy, a reflector and a good guy, and then we looked at basically what was happening. So we sent queries and look at what's happening there, and change the account, the middle box, [inaudible 02:01:58] about doing the Rate Limiting.

Well, we [usually? 02:02:03] attack, and basically we… Suddenly we went ten times, I believe, quicker and more answers and then we switched on the Response Rate Limiting – and this is for an ANY attack – and you see that we go down from traffic effect of 50. So suddenly the traffic goes down to the victim, which is nice. So this is actually very effective in this case and here you see what happens with false positives if you just let in… If you let all two of the TCP, you don't really see a lot

of change in traffic, but if you do one out of three, we do it again of TCP, you see the amount of traffic going really down in the direction of the victim. It looks like success, great.

Now, if you do more and more different queries and still do a reflection, then it looks like more and more to little requests, and you see the test hold, if 100% of the questions are for reflection Domains, it really looks like official questions and Response Rate Limiting is not happening at all anymore. Not effective. So yes, it works for the short-term, but it doesn't work for the long-term.

And these are more… You can do the same with the false positives and see what the effects are. And this is all written down in a big report and with [inaudible 02:04:18], all the details and all the parameters. But in the end, this is just a game of Whac-A-Mole, you know? Just the moment – although it's highly automated – it just doesn't do a principal solution. And furthermore, if you get really sophisticated, you have an army of Botnets and an army of reflective servers. I mean, you can't really blow somebody fat very easily out of the water, and that's why it's again, a Whac-A-Mole game.

If you get smarter about stopping attacks, the bad guy is getting smarter as well and will probably out-smart you fairly quickly. And as I said, this is really in [inaudible 02:05:10] has to do it UDP. There are other protocols which have the same problems. Up to two weeks ago, one of the biggest attacks was actually not on UDP, that was a reflection, but SNMP reflections, which is having somebody send an awful lot of SNMP answers to the victim.

Well, DNS is actually the low-hanging fruit of the day. Meaning that it really can get batted down. Well, go and pick something else to be as equal, that is before. What you really want to do, the real solution is that you should prevent spoofing. You should implement source validations, so you really should, in your network, only allow packets getting out of your network which belong to you. There's no point in sending packets out where you never get the answers for, because don't let fake packets leave your network.

Now, a couple of documents about it. The problem's really old. I think BCP 38 just passed its ten-year publication date, but there are more publications about it. And why do we want to do this? Well, for the good of the Internet, that's why we want to do this, and also for your own reputation; that you're not aiding and betting at the bad guys. As I said, here's a list of other, further reading. And the Firewall, DNS Firewall; www.bortzmeyer.org from AFNIC has some examples of how to do that.

Dampening by lutz.donnerhacke.de. The Rate limiting details edition, the [massive one? 02:07:10] to report [inaudible 02:07:11]. And it was also just two weeks ago a website popped up where people were putting in documents on how to do BCP 38 (www.bcp38.info) for various networks and stuff like that. Any questions? And I'm losing my voice. [laughs] Back to the pointers in that case.

UM: Hi, I'm [Yosh? 02:08:01] from Finland Registry. I was wondering, do you believe that this might also be the end of DNSSEC or that we should

maybe make a new version of DNSSEC or something, or in that community, are there discussions about that?

JAAP AKKERHUIS: No, because it's very easy to get traffic without DNSSEC; just get another Botnet to do it, it doesn't really matter. Yes, DNSSEC makes bigger packets than normal, but it's so easy to do this stuff, it doesn't really matter. And the net result of not doing DNSSEC won't help a lot with these attacks. Most of these attacks happen actually on [usual? 02:09:57] servers. It's lies to have [usual?] servers to do DNSSEC, because there's less problems for the bad guy, but it doesn't matter at all in the end.

I mean, [the packets? 02:09:15] are two, three times as big, maybe two, three times a Botnet. Well, this costs $10 for 1,000 servers or clients and then you can day-long hit somebody, if you really want to do that. Ja.

EBERHARD LISSE: Any more questions?

JAAP AKKERHUIS: So one thing you could do as a TLD is ask your [vendors? 02:09:48] about how they might protect it and defenders of your [ANY class? 02:09:52] cloud, or whatever. Or, your own routers on your own [ads? 02:09:56], make sure that you're not by yourself by inside operators or inside jobs or whatever. Just a few, basic ideas about how you deal with your network architecture will help already.

EBERHARD LISSE: I mean, in that context, Ed's presentation was actually quite nice because it systematically stated what you need to and specific examples of this are quite helpful. Thank you very much.

JAAP AKKERHUIS: Okay. [applause]

EBERHARD LISSE: So, next but not least is Janelle McAlister from MarkMonitor. Initially we planned a round-table, but one of the presenters didn't come, not this one but another guy told me they wouldn't come because of the bird flu, to which I said 'what bird flu?' But we'll make out of this intended round-table, we'll just make a presentation.

JANELLE MCALISTER: [background chatter] Okay? Thank you. So like you said, I'm Janelle McAlister, I work with MarkMonitor, and for the last four years I've led our ccTLD Incident Response Team, so anytime there's been a major incident with either a ccTLD Registrar or Registry, I've been working with that vendor to respond to the incident. So the intent of today's presentation is just to share what we have learnt, because there are some common trends on how Registries or Registrars are being attacked, and what things can be set into place to try to alleviate some of these attacks.

So what we've seen is now, more than ever, hackers are noticing that one way to compromise high-profile Domain Names is either to compromise them at the Registry or Registrar level, by updating the DNS on the Domain Name. And they are doing this through an exploitation of either technical vulnerabilities and also through social engineering attacks that we'll get into in more detail.

Once a TLD has been compromise, it's highly likely that they will be a target again. What we're seeing, especially recently, is a second wave of attacks on TLDs that were previously compromised. It appears that it is the same hacker that has hacked them before that is targeting them again, and they are making referent to the fact that they were able to attack a TLD multiple times.

Since 2009 there have been 31 TLDs that have been compromised. This includes both ccTLD Registries and Registrars, that have been compromised. And the way we define, again, a compromised Domain or a compromised system, is anytime the DNS is maliciously updated on a Domain Name. This is also, actually, just Registries or Registrars where we have had a client affected, so this does not include any other third party Registrars that we are not affiliated with.

So here is the breakdown of the numbers, that Cynthia earlier made mention of. Since I started leading our Incident Response Team in 2009, there were four ccTLD instances, followed by three in 2010, we saw a slight jump in 2011 with eight instances, and then over the last six months, between October of 2012 and recent; last month, there's actually been – let's see – 13 ccTLD Registries that have been compromised, in addition to three Registrars that have been

compromised. And then we have seen multiple attempts on Registries over the last six months.

Part of this started in October when there was two higher-profile Registries that were compromised, that received a lot of media attention and what we saw right after that is that the hackers came out in waves and tried to do the same type of attack in any TLD that they could. So it is what seemed to start the wave of attacks, and then again that's also 13 unique attacks. So we've also had a couple of Registries that were attacked multiple times in that six-month period, so it's definitely been on the rise over the last six months.

Once a hacker gains access to either a Registry or a Registrar system, they update the DNS on a higher-profile or high-traffic Domain Name, so that the site redirects to their 'brag page', and I've put an example here of what it tends to look like. Our concern right now is that they currently are just going to a 'brag page', but worst case scenario is that they actually start spoofing the original website to try to capture user data or any information like that. So right now, the spoof page is somewhat harmless; I think it's damaging to both the Registrants brand as well as possibly the TLDs brand, but our concern is as these attacks tend to be coming more sophisticated, eventually this could lead to these hackers spoofing actual Registrars or Registrant's data.

So the ways these attacks are happening are there's three targets right now. Like I said, Registries are the current target, however in the last four years there have been waves where the DNS providers were the target – there's definitely been specific periods where it was Registrars that were the target. In the last six months is tends to be more

Registries – but we see this somewhat fluctuating over time. I think right now there is just an awareness that Registries can be compromised, and that seems to be where the focus is at.

And these attacks are very similar on each group. The primary way that something is compromised is through a SQL injection, an attack on a Web interface. We've also seen social engineering attacks. The most common social engineering attack that we see is any time a real person's name is listed as an admin contact on a Domain owned by a corporation, we see someone contact either a Registry or Registrar and pretend to be that specific person. Which is why I think most of you have come up to me and asked that you not require a real person's name be listen on a high-profile Domain, just because it leads to that type of attack or that type of attempt of an attack.

We also occasionally, with a social engineering attack, we'll see forged documents on high-profile Domains. Occasionally we'll see a brute force attack on passwords and different options like that. One of the things that we're seeing as precursors to an attack, on some of the attempted attacks that we've seen in the last six months, are a wave of multiple requests to modify a high-profile Domain.

And generally we see this with Registries where you can request the Domain update through an online system or an online page, and then it initiates an email sent to the admin or Registrant, that requires an approval. And what we see is that these are initiated over and over again, with the attempt of trying to break what the formula is, to approve that Domain update. And it's very similar with password resets on accounts or Domains, is we see multiple attempts to either recover

that password, in an attempt to try to break whatever formula is used to reset passwords, or approve a password reset.

We've also seen an influx of hackers using social media to provide information to other hackers on how to compromise a Domain Name. Recently we've also heard from them asking for your emergency contact information, because when we do see these, we do try to contact a Registry to let them know that information has been published, so that they can start taking measures to prevent an attack from happening.

So some of the ways that you can… Or some of the measures that you can take to prevent an attack, there's multiple different things. I think the security scanning that Cynthia mentioned earlier today is one step, but I also think that this reaction to this or the responsibility on this is on both the Registry, the Registrar and the Registrant. I think we all have our roles to play.

With the Registry it is having a secure system and offering things like account lock or – we'll get to this but – Registry Lock, having secure security protocols in place between you and your Registrar – for example some Registries will never accept a request directly through a Registrant but must always go through a Registrar. That helps us mitigate the risk of somebody pretending to be a Registrant. We also really like Registries or Registrars that offer Two-Factor Authentication, which .br just mentioned that they rolled out – that was definitely a success.

And one thing also is – and this is common with most Registries, unfortunately there are a few out there that, if you guess a password

multiple times, it does not lock the account, which can lead to a brute force password attack, so we do ask that a Registry does lock an account if there is a request made for the password more than three times. We also have asked for IP restrictions on accounts, to limit who can actually access the account and Domain Names.

And lastly, again, I'll give you more detail about Registry Lock. This has been a big topic for us recently; it's something that our clients are asking for on a regular basis. And the purpose of Registry Lock is the ability to take a high-profile Domain – so we're just looking at, generally, the top 20 high-traffic Domain Names in your TLD –, removing those from an online system so they can't be updated unless there is a manual process or a manual security process set into place.

So for example, it would remove the ability to update the DNS, to transfer the Domain Name, update any contact information, unless someone from the Registrar and someone in the Registry, both review the Domain Name and the update and follow specific security protocol to update the Domain Name. And that prevents both malicious Domain updates but also accidental Domain updates that would be critical if the site should go down.

And this is something that large corporations with high-profile Domains are willing to pay for it – it's not something they are expecting for free – so as a Registry is implemented it's definitely something to keep in mind; the feedback we're hearing from our clients is they would rather pay more to have additional security on their high-profile Domains, rather than having that Domain being treated like any other – because it would be critical to both them and possibly the TLD if it should go down.

And again, Registry Lock is just used in conjunction with our Registrars security protocol, so we would have security protocols set between us and our client and then a security protocol between us and our Registry. Any questions? Okay.

EBERHARD LISSE: I noticed on occasional request from MarkMonitor, which happens to be our most favorite Registrar, as I must say, our biggest also; you guys when you contact us do not sign your emails.

JANELLE MCALISTER: With PGP?

EBERHARD LISSE: Either PGP or with a certificate. None of them come like that. I have rejected things from MarkMonitor because it came from an address that I didn't know and it wasn't signed. I think that's the first thing you should do, if you're contacting, you should either put a PGP signature or you should put in a certificate in your email, so that we know it's a more or less… It's an additional step.

JANELLE MCALISTER: Absolutely, I agree with that. There are some Registries we have implemented that with, and I'll make a note to ensure that we're using that protocol with yours. Some Registries we use email more than others, but I think PGP or some type of security measure is…

EBERHARD LISSE: You don't implement it with the Registry; you implement it on your side. If all your staff that communicate with Registries about this, sign their emails, or put a certificate in, so we know that when they come, our system can check it is from this address to this at the public server, and we know at least we can communicate in a safe manner. You shouldn't also expect from the other side, from us, to sign our messages.

I don't really sign our Registries, but my own email address is signed, and if we know that we're dealing with you, then that's good. But that's really a suggestion I think… You should have all your outgoing emails, whether to Registries that use it or not, signed. Jay was first?

JANELLE MCALISTER: Okay.

JAY DALEY: Thank you for the presentation. In your last slide you said that you would like this Registrant Lock put in place that inserts a manual process? Now, it appears to me that there are two elements there an you are perhaps only focusing on one of them; that what you, I suspect, could well be asking for, is for there to be no way through the Registry for the Domain to be modified. In .nz, for example, we don't provide any form of serviced Registrants, there's no way a…

A compromise for us is the full-scale thing; there's no website on the front where someone can do any form of injection attack to do that. But then inserting something manually, something very, very different. So it seems to me that you are asking for the hardest one that anybody can

provide, which is manual, while actually asking for the simpler one – that there is no access to the Domain Name through a Registry's Web interface – is a more practical thing to ask for.

JANELLE MCALISTER: Yeah, I think what we're asking for is for both the Registrant or even, say if you're using an EPP system over an automated system, any commands from the Registrar would be rejected unless there is a manual process that is… Am I? Okay.

JAY DALEY: You are, but let's just work through the logic in this, all right? The compromises that you have listen there, if your figures are to be believed – and not to be rude but this is difficult data here – is, are compromises of a Registrar's Web-based interface, by a SQL injection, that let's someone then get in and change individual data on a Domain Name System, a Domain Name. Okay?

We don't provide, as an example, a frontend that could be vulnerable in that way. Okay? Right? All our Registrars do… Now, what you are saying is that we need to take on the investment in order to prevent against the failure of one of our Registrars, okay, by providing this service in the case of us? And that is the bit where I think you are making an unreasonable jump.

What the errors are that you've identified, the break-ins you've identified, are through failures in the Web-based interface, so let's fix the Web--based interface in the Registries, okay, but let's not jump to

the next step, which is that Registries like us should have to deal with a broken Registrar by inserting a manual process into it if the Registrar is broken.

JANELLE MCALISTER: Sure, I think based on the current attacks the focus is definitely on Web-based systems, and so that is initially what needs to be fixed. I think the second piece that is somewhat separate is the implementation of Registry Lock, which doesn't necessarily just recognize what's recently happening, but it also gives a Registrant of a high-profile Domain Name that additional sense of security on their Domain Name.

EBERHARD LISSE: What I understood her to mean is that you shouldn't, if your Register – let's say Google; you've got Google haven't you? – if Google wants to make a change, you don't have to bother about Registrar C, you only both about MarkMonitor talking to you. That's what she means.

JAY DALEY: What you could reasonably be asking for is for those Registries that do have Web-based interfaces, for certain Domains to be taken out of that Web-based interface, they're not available to that. That's the next stage to what you've been asking for there, and I think a simpler, more reasonable stage to ask for and I just… And it also has a logical progression as to how you get to that, rather than going directly to the manual processing, which is a huge policy issue for many of us, very complicated and if we get it in any of our lifetimes, congratulations.

EBERHARD LISSE: [inaudible 02:28:12], which we are running, and many small ones, only MarkMonitor can see Google. I think we can even put a Registry Lock on once we have done that, so if they want the changes they have to contact us, and that's why I want the secure email. I think since they have indicated already to us in the past that they're willing to compensate us for manual labor, we don't really complain about additional revenue, but so far it hasn't happened yet. I'm fully with you on what you're saying, that if a client wants to pay me for paying special attention to – as they say in Nigeria – I will do that. Jaap, you were on?

JAAP AKKERHUIS: This is more a matter question. What I'm always confused about in these types of discussions is, what's the real special part about Domain-related attacks? If I read all the causes of most of these things, it's just like yet another Web shop with very bad IT. We've got cracks…

EBERHARD LISSE: Jaap, for Google it's a public relations disaster, for me, I don't give a damn. That's… For them it's a monetary value.

JAAP AKKERHUIS: But the processes, whether you buy a teddy bear or a Domain Name, they are basically the same process, which needs to be handled carefully.

EBERHARD LISSE: If you sell five teddy bears or 10 million teddy bears, it makes a difference. If then there is a drop from five to for, it's not a difference, but from five million to four million, that's the thing. Technically, it's the same, but they don't want the hassle. Microsoft don't want these attacks anymore. That's why they've finally figured out that they must do something to help us about it. That's at least how I see it. It's purely public relations. They don't want to get bothered.

JANELLE MCALISTER: I agree. It definitely reflects on their reputation because there can be the perception that they were compromised, that their security systems were compromised. It's also a matter of they lose the revenue of the Web traffic on the Domain Name, for as long as it's down. And with these instances we've seen Domains go down for a couple of hours, or a couple of days, and that can mean a lot of traffic for these companies that are being targeted.

EBERHARD LISSE: All right. Anybody else? Thank you very much. It was a nice presentation. [applause] So as usual Jay Daley will close and give us some thoughts, or thought-provoking thoughts…

JAY DALEY: Thank you. I am Jay Daley. I'm the Chief Exec of .nz Registry. So we started off this morning with Ed Lewis, giving us a very detailed presentation on best practice in security. I think for all of us that was an interesting presentation because of the breadth of issues that he

covered, for us to think about. And I think that there are a number of us that have begun to realize that when you get such a broad subject area, that you cannot fix that just by clever people and documented processes; we need something beyond that.

We need something that externally audits us, something that is a standard that we can set and achieve, and things that we can be very clear about, whether we are there or not, because to use a very culturally-specific phrase, but one that I hope you'll all get – the devil is in the detail. And so we cannot assume that we have got everything, unless we systematically check every piece of our security, right the way down, following a structured process.

The Microsoft service that was presented to us was a tool in this and I think you heard my views that it was a useful but very small tool in this; many of the problems are greater than that. We then had the SSAC talking to us about the security issues in the Domain Name space from New TLDs. This is clearly something that's going to grow and get bigger over time. The issues with people naming their local networks .home, for example, is a large issue.

The letter from PayPal to the ICANN Board earlier in March has a wonderful phrase in there. it says that it is 'practically impossible to rename a large, active directory'. And so when you have made the decision to call it something.home, you have to live with it forever. I know the answer to that and it doesn't involve Microsoft. Okay. And we can be certain that there will be more to come in this area, as more technology is discovered.

Then we had the interesting presentation from Coremail on… Oh, no, sorry, within the SSAC presentations as well, the… You noticed that SSAC is doing something interesting with formalization of specifications around data, around WHOIS and other things. It's a new thing. It's a level above the IETF in formalizing protocols, and it's a level below ICANN informalizing policy; it's formalizing a space in the middle about what our terminology is, what the components are that we can put together to create the services and our policy.

And I think we will see a lot more of that and I think that that's something that many of us here have the skills to do already in our day jobs and should be contributing to. Then we had the presentation from Coremail. It's great to see that we have Internationally Enabled Email now available, but it has taken a very, very long time to get here and I worry that it will take a very long time to be accepted globally. And I think that many of us recognize that we have a larger role than just being a Registry or just a TLD, and for us to promote the better use of technology in this type of way is an important part of our mission.

Then we had the presentation on Escrow, perhaps not the most interesting one for many ccTLDs, because we generally don't have an Escrow requirement at all, but it's a useful case of understanding that there is good evidence that things do go wrong, and it enables us to ask ourselves how much are we planning for failure – because we know the evidence is there –, that those failures will happen.

So I have… There is no need for us to do Escrow in our Registry, our relationship with the regulator doesn't require it, but it is something worthwhile us thinking about, certainly. Okay, then we had Lewis

talking about the TPMs, and this is interesting. We are… Many of us are many years into DNSSEC presentations now… Sorry, into implementations now, and yet we are still looking for good DNSSEC hardware. There is still not the cheap, simple, wonderful, extremely fast, shiny, nice looking DNS hardware that we all want, so if anybody wants to form a company to make it, then come and talk to me and we'll see what we can do, because this is just driving me up the wall.

Okay, the Emergency Contacts presentation. This is unfortunately something for us all to pay attention to. I think we will all be required to be part of that relatively soon, and it's another example in an evolving landscape, of where collaboration is born out of crisis. We've seen some groups form attempting to create standards and mechanisms for people to share data in emergencies. Some of those have not been that successful, and this is one that I think is such a small task that it will work, and then we can build things on top of that.

Then we had the amazing presentation from CNNIC, because of the huge scope of operations within CNNIC. I think there was everything in there. Everything that I know that anybody does on the Internet, it being done by someone in CNNIC. I was very, very impressed. Some years ago I saw a presentation from CNNIC where they had a picture of one of their computer rooms with armed guards outside it and I have wanted an armed guard outside our computer room ever since.

Yeah. If anyone else again, would like to collaborate, we could share one and send him around the world, you know, just have him there when a visiting dignitary comes. [background chatter and laughter] Yeah. But it does show us just how big our world is. Just… And it also

shows I think that the major benefits in the last few years, come from a dedicated and professional structural approach.

So CNNIC created this Anti-Phishing Alliance as a way of tackling a problem. It wasn't just another project team, it was a structure of individuals, a collaboration, intended for a long-term success in this area, and so it was a very impressive way of dealing with it.

Then we had .br looking at credentials and hashing and Two-Factor Authentication and the wonderful Google Authenticator. This is really getting very complicated now. And none of us, I think, can afford to ignore this. We cannot stand still in this regard. These types of technologies, the use of these technologies is very important, and we simply have to start being engaged with that. I think there is a degree to which ccTLDs who do not get engaged in this, will end up changing – perhaps by outsourcing more, so they have these facilities elsewhere, or perhaps by becoming the major targets of hackers.

The Rate Limiting presentation from Jaap again gave us lots to consider and again told us that we cannot simply ignore this problem. But thankfully it gave us the very simple solution, which is Rate Limiting patches and BCP 38. We've used the Rate Limiting patch when we had to, in an emergency, and the problem went away immediately. So it's great. And if I ever become Emperor of the Internet – and you can vote for me, obviously, for this – then I will mandate BCP 38 on pain of complete disconnection from the Internet or disconnection from Facebook for two hours or something, for anybody that does not do it.

Okay. And finally, MarkMonitor, their presentation was interesting for us to hear that Registries are on the global target list now. For many years people assumed that we were competent and that we knew what we were doing and that we were safe and we were secure. And now they've tested it and found out that maybe we're not quite as good as they thought we were, and so they will start coming stronger for us, more and more.

And it will take an effort; a long time for us to rebuild some of that aura of invincibility for them to stop attacking us. So there are four, I think, real themes that have come out from today's presentation. The first one of course is security – I think there are a number of people who are here… We had a presentation in 2008 when somebody said that 2009 would be the year of security, and somebody else put their hand up and said 'and so will 2010, 2011, 2012…' and so on. You get the picture. Security is the big, big theme. I mean, I'm bored, I'm over it, I want to talk about marketing, but security is the big thing that is going to go for some time.

The other big thing that is coming through is professionalism. It's not good us just being good. We need to be able to prove that we are good, describe how we are good, share how we are good and check that we are good against others. Then the next thing that we are finding is this phrase again, that the devil is in the detail. We have to look at every individual component of what we do, in some details, and analyze it and fix it for us to continue to be good at things.

And the final theme is that we simply cannot stand still. These things are pushing us now more than we are pushing them and after many years of

us just saying 'would you like a Domain Name?' and taking money, we now have to earn our money a little bit harder. So that's it for the closing ceremony. Thank you. [applause] Okay, and we'd like to thank our sponsors, Coremail, for the food, the lunch today, that was very nice of you. And we'd also like to thank the Organizing Committee and Dr. Lisse over here for making it all work.

Thank you.

[applause]