BUENOS AIRES – Tech Day
Monday, November 18, 2013 – 11:00 to 17:00
ICANN – Buenos Aires, Argentina

UNIDENTIFIED MALE: Buenos Aires 48. Monday, November 18. Tech Day, La Pampa

EBERHARD LISSE: Good morning. For those of you who don't know me, my name is Eberhard Lisse. I run dot-na and I'm not going to make the usual joke about day job and night job because I'm going to stop doing my night job at the end of the year. I'm only doing my day job. Somehow we sit in front of the mics. Is that better? Okay.

Today we are having our 22nd Tech Day, I counted it this morning, which is not too bad, which means that in L.A. we will have our 25th meeting and I must say I never expected it to be such a ride. In Durban we noticed – because we were in the same room as the ccNSO meeting – that we had more participants on the tech day than we had on the ccNSO meetings.

We had to change the agenda slightly today from the printed one because our host speaker feels a little bit uncomfortable in English and we were not aware of that that so we haven't arranged for a translator so we pushed him to the afternoon and we'll make a plan to find the translator or to find somebody who can assist him with the speaker – with the English that he doesn't feel uncomfortable.

As you have probably have seen in the agenda we have a mixed bag of things. I'm very happy that we managed to get Jaeson Schultz from Cisco. We had this presentation from Nigel last time in Durban about bitsquatting and we've noticed that Jaeson has done research and Roy Arends mentioned this and then we looked up and found that he'd done research, so I think it's going to be a very interesting speech.

I haven't met Jeremy Rowley yet. Is he in the room? Oh, there you are. So are you Jeremy Rowley? No. You're Jaromir Talir. But where is Jeremy Rowley? Not yet here? Because we're missing his presentation so we'll...

UNIDENTIFIED MALE: [inaudible]

EBERHARD LISSE: We have the presentation now? Okay, we have the presentation. So he will be coming and that's not going to be a problem. If we have problems with speakers not arriving on time we'll just switch it through a little bit.

Jothan Frakes is going to speak about Public Suffix List. In case you don't know what that is, Mozilla project maintains a list of sub-level domains in each ccTLD or in each TLD, and when I looked at the one in front of me, maybe it's not really correct, so we're going to update this. But probably he wants to mention what it does, where it is and get context so that they can update this.

Then Don Hollander has been recently re-appointed as the manager of APTLD and he has some projects in mind, so I thought I'll give him a short time to present.

In the afternoon, Anne-Marie will speak twice about ISO 27000 and about an EPP testing tool that they've developed in Sweden.

Norm will speak about the Secure Domain Foundation. The IETF has asked – and for a spot I haven't really – I don't remember exactly what the topic will be but I think it's my view on this is, whenever the IETF wants to have a spot we'll give them a spot and because they are the people that actually developed the technology, so it's very good to develop a relationship and to then hear what they have to say. Norm is going to speak about secure domain foundation as I said.

Daniel Kalchev, who I haven't seen yet – oh, there he is. Good morning. He has got a protocol where the clients or the applicants can talk to the registrar and to the registry. Dot-co is going to talk about some security incidence that they had and how they handled it.

Then comes the host presentation. Warren Kumari is going to speak about automating DNSSEC Maintenance. And the last thing is the South Africans are going to speak about their Mark Validation System. I haven't read the presentation yet, but I think it's good to have a little bit of competition to places like Trademark Clearinghouses and so on, which is good in the [biosphere].

Then Jay will talk about something interesting as usual.

Hence, without further ado, Erwin Lansing. We have got two floor microscopes, so if there's a discussion we will be able to do this.

ERWIN LANSING: Good morning. I've been involved in the FreeBSD project for a number of years and when I started DK Hostmaster a couple of years ago they were so generous to sponsor some of my time that I could continue working on it and of course to change my focus a bit on DNS. And a year or so – a bit more than a year ago – we start looking into how we do DNS in the base operating system, and there's been a number of changes since and I want to talk about that in a bit.

A bit of history. BIND and FreeBSD go back for a very long time. I did a little bit of archeology and find the first import of NAMED back about 20 years ago which was in FreeBSD 2.0 and the libc resolver functions actually go back even further to 4.3BSD or even further to – a lot of people seem to have gone to details since.

I think actually the resolve functions came from BSD into BIND and they were back and forth into FreeBSD in the lwres library. So until FreeBSD 9 series which is the production release we have right now, we got a whole BIND distribution in the base system. That means the command line tools: nslookup, host, dig, nsupdate.

As the Daemon, both the recursor and the authoritative – what you need for an operating system is actually just the recursor part. Thus, the DNSSEC key handling and just everything under the kitchen sink is there.

So we start looking at doing things differently. When we did the actual commit of removing BIND towards of course the Internet rumor mill that went wild, so I just put up the slide saying this was for technical reasons, not political. We're still very fond of ISC. They're big sponsor of the FreeBSD infrastructure, especially the high bandwidth download stuff we have, so we still have before ISC. It's got nothing to do with that. It's technical.

So what are the technical reasons? Like I said, we just need the resolver library and command line tools. So we want a smaller code base. It makes things much easier. One of the things that make it easier hopefully is less security advisories. Just by having less code, hopefully we get less security issues.

Again, in [inaudible] a BIND that's very insecure, it might have been BIND 4 and BIND 8. I don't think BIND 9 is much more insecure than any of the other servers, but it's very highly [abused], highly scrutinized so probably more bugs are found. I'm not saying BIND is more insecure. It could be something else.

Then we have the support life cycle. ISC support is very detailed but it doesn't overlap the FreeBSD support cycle which means we either have to upgrade in our release branches or backport to fix it ourselves. I'm just going to say this is not specific for BINDS, some of the projects just support the top of the tree so that would be even worse for us to do so. And then we got the big elephant in the room that's BIND 10. Due to several of the new dependencies, BIND 10 has grown. It's impossible for us for to import that directly into the base system. And we, so far,

have no idea what's ISC's plan for BIND 9 and how long they're going to support it.

Then slightly going back to the large code base – people have used BIND directly from the bases and for all kinds of services and setups and all kinds of options and it's just too hard to support from the base system and what it uses to support as a third party package for the port system.

So if we don't want BIND, what do you want? Well, we want to have a DNSSEC-aware resolver library. When we start looking into that it turns out if you want to do a DNSSEC validation, you want to do caching. You do not want to go through the whole chain of the tree and do the validation again for every look up even if the recursor is very close to you. We need at least the host and J-command line tools and we want a liberal license which would be BSD or similar.

So what's around today? There's a lot of projects around. Several new ones popped up about two years ago. We got BIND. We got Knot, Powerdns, djbdns, dnsmasq, ldns, and Unbound, and there are probably even more.

Going back to the list, just looking at the license, there's actually only two projects available including BIND. All the other ones are GPL or djb, and not so far only has utilities and not recursor Daemon yet.

For FreeBSD 10 which is supposed to be released soon, hopefully before the end of the year, we replaced BIND by Unbound and of course ldns library and tools. We really wanted to have a complete replacement for dig. We couldn't find it, so we have to manage with drill. We did find a tiny host wrap around ldns, so we do have host. We did look into

importing the Knot tools but that would mean also importing the Knot library and then we go back to having a small code base so that didn't make much sense to have two full libraries into the system.

If we install a local Daemon-only, that means if you want to do any services outside of the local machine, just install the third-party package. This can be Unbound, this can be BIND, this can be any of the other projects. You could do anything you want with all the options you want. We just don't support them directly out of the base system anymore. This makes a very simple setup. This also means if we want to change things again for FreeBSD 11 in a year or two, we can do so very easily.

Of course, we turn on DNSSEC validation by default. We also do SSH Fingerprint checking, so if you use it as a fingerprinting at DNSSEC, you're no longer prompted by a new host. And like I said, it would be easy to replace again FreeBSD 11.

So what do we want for FreeBSD 11? A full wish list here. Again, a DNSSEC-aware resolver library, a caching Daemon, we want the command line tools, it should have a liberal license, it should have a low footprint, it should be fast, it should be thread safe, it should use some of the compartmentalization projects we are working on. I always forget the different between sandbox and compartmentalization, which I cannot pronounce even. A project called Capsicum which is co-sponsored by the FreeBSD Foundation and Google and work going on at the University of Cambridge.

And we want to have a standard API for reporting DNS and especially DNSSEC information back to the user. When we did the import we had a very interesting case with SSH Fingerprinting and Open SSH. If you have a mismatch between SSH Fingerprinting DNS and the local cache key that's on disk, by the time – if Open SSH prompts the users to ask what it should do, it has actually forgotten about the DNSSEC data so it cannot tell the user if the DNS data was validated or not, which doesn't leave the user many options.

Those kinds of things get even further with project like DANE. We got some work going on in IETF which some of us are following and if any of you is involved with this kind of stuff let us know, we're very happy to work with you because we got an offer and a challenge for you then. Our workforce is really happy to work with you to get this working and get it integrated and implemented into a working operating system and get you supply as many users as possible, and the challenge is unlike IETF we don't just want running code, we want running code in production and we want it in 18 months from now.

OLAF KOLKMAN: Do you want me to speak to that?

ERWIN LANSING: Yes, actually I do. I'm glad you came in.

EBERHARD LISSE: We must speak onto the microscope – microphone – because we have a remote audience and they will not hear. So please also identify yourselves.

OLAF KOLKMAN: Olaf Kolkman, NLnet Labs. I'm partly responsible for Unbound. There is an effort to produce running code. VeriSign is doing that in collaboration with us and there is a plan to have a bunch of hackathons early next year. I haven't got all the information in the back of the pocket at this moment, but there's actually work going on and that will be open sourced and made available. The idea is to do the loop of implementation and update of that API. So work in that area is going on.

ERWIN LANSING: I'll find you later.

EBERHARD LISSE: Any of the other names of producers like Knot, EuroDNS, [inaudible]. D they want to speak about any of these? Any questions?

ROY ARENDS: My name is Roy Arends. I work for Nominet. I really like the idea of secure by default, having FreeBSD at one point in time, having validation on by default. My question is how do you prime for the root key since the moment you deploy FreeBSD in production, it might not be the same time as that same key is going to be deployed? So, how do you bootstrap?

ERWIN LANSING: That's a very good question. Well, if you start a new system, it's of course included in the installation. So if you don't trust that CD, you can't trust the root key either. If the root key is updated, it will be updated with the FreeBSD update command line tool which has the secure – which does all the security signing, etc. for everything else including the root key.

ROY ARENDS: Okay. Thank you. Good presentation by the way.

ERWING LANSING: Thank you.

EBERHARD LISSE: Any other questions? Alright, good. Then we go to the next one. That would be Jaromir Talir, if I'm not mistaken.

JAROMIR TALIR: Hello. Good morning. My name is Jarormi Talir. I work for CZ.NIC and one of the major topics for our work and during the last two years was validation of contact data and FRED or in registry. We are using FRED as a registry tool so my presentation will cover all the aspects regarding contact data validation that we did during the last two years and that we also implemented in our registry solution and made it available for the others.

So, first, I will introduce FRED for those that don't know the system and then I will speak about – whether validate contact data on the registry site or not. I will talk about some approaches how other registries do validation and how we do validation. Then I will mention some – what we do to keep the contacts data valid.

We at CZ.NIC created our own registry solution and are working on this system for maybe five years. The system is called FRED and it's open source registry software which is now used also in few other countries. The most recent registry using FRED is Albania this year, and there are I think two others that will probably use it until the end of the year.

The system runs on Linux and we have binary packages for Ubuntu and Fedora. Starting up the system is a matter of few minutes but I know that there are some other things like customization and migration that can take much more time. The system is full-featured. There's all the things that you can imagine when you – of our registry operator. There is a website that you can find some information about the system.

As a registry, we maintain a lot of contacts, a lot of registrant's data and this registry is not intended to be for anonymous entities. We try to do the best to ensure that the contacts in the registry are valid and there are many reasons to have the data valid. There are legal issues that we must solve that we are asked by some law enforcement agencies to give them information about the owner of the domain, who is behind. There are also some technical issues that if you want to fight against spam and things like this that you'll need to know who is the real or at least technical contact of the domain.

There are some other reasons that the user can benefit from having correct data in the registry like that we always notify the registrant about some important event like the domain is going to expire and things like this. So then it's better to have the contact data valid. Otherwise, the registrant can easily lose his domain.

So there are many advantages to have contact data valid, but of course there is some cost associated with this validation. It's not straightforward, but we are convinced that the advantage of this – to have contact data in registry valid – is bigger than the cost that we must spend on validating this data.

If you are interested in registry data validation, I would like to draw your attention to Security and Stability Advisory Committee document 58 that was released in March this year. It quite summarizes all the things that we are doing over the last two years about domain data or contact data validation and the content of the document there – again, the reasons for validation is similar to what I've described on the last slide.

Then the [SSAC] suggest some taxonomy for validation. They mentioned three different kind of validation. Syntactical validation, that you can check whether e-mail format is valid, whether phone number is valid, and things like this. Then the operational validation where you're trying to find out whether you can reach the contact either by checking whether the address exist in some registry or whether the domain or the e-mail exists and is active or, for example, by sending some information to that address to find out whether it's reachable. The last one is identity validation where you really try to check some ID card of

that contact, which is similar like certificate authorities do for validation of physical presence.

At the end of this document, describe some data of domains and contacts that are subject to validation like as you might expect its e-mail, phone number, fax numbers, physical address, and these.

We work in a registry-registrar model, and in this model it's the registrars that are responsible for putting the data of contact into the registry. We have a contract just for registrars and according to this contract they should provide sufficient effort validation of data. As you can expect, it is hard to enforce this – to fulfill this requirement and it's really problematic to try to suspend the registrar agreement and to fight with them that they do not do this sufficient effort. So, it's not easy.

For the ICANN, there is new Registrar Accreditation Agreement that you may be aware about released this year. For ICANN accredited registrars, this document prescribes some mandatory requirements for doing this kind of validation. It is hardly applicable for our registrars because we have about 50 registrars in dot-cz and there's only one ICANN accredited registrar. So this will not fulfill our problem, and it is easier to follow the ICANN way and try to put into the contract this kind of prescription or this kind of requirement because probably there would be the straight opposition from the registrars to do that. And so we decided that maybe the registry as a central point where the old contact data at the end are [filled] is a better place for validation. That's why we started to do that.

Last year, in Costa Rica, there was a Tech Day and there was two presentations of some approaches of other registries that are doing contact validation. There was a presentation from Turkey registry and they mentioned that registrants send some documents themselves to do the registry via the portal. This is one approach.

The good approach for the validation is in Estonia but they are quite far ahead of us because they have EID infrastructure, and so for locals they can take advantage of this EID infrastructure. For foreigners, they have a different approach. They using the information associated with bank accounts, and the registrars when they receive money from the owner of the domain, they get some document from the bank, from the accountant. They embed it into EPP commands and send it to the registry. So this is other approach which, for us, because we don't have any centralized registry for our government – or we have but it's not accessible by commercial companies to check against this immediately. So this is also probably not the way how we can do that.

So we decided to implement two complementary approaches for validation of contact data. One is sort of voluntary and automated from the view of the registry and this is when the owner of the domain or registrant will decide to voluntary validate its data and on the website they will ask for the validation. We will send some codes to his contacts and then we will accept all these codes back and mark the contacts as sort of valid. But of course, you will pass this way if you are probably sure that you have the contact data valid. So we must have some other way for the contacts that are not active and the other way is proactive seeking of some suspicious contacts in the registry. Then this procedure is much more manual than the previous one.

So now I'll talk in more detail about these two approaches. That voluntary approach works in the way that there's a website that the registrants will enter handle of the contact and we will send the code to him by e-mail and SMS mechanism. And after collecting these two codes, on the other website, the contact receives the first grade of validation. And at this moment, we send a third code to his postal address via snail mail. After receiving the code from this channel, we will mark that contact as second grade of validation. So each contact has a new state that says how well its contact is in the registry.

For SMS and letter delivery, there are some companies in Czech Republic that provides that interfaces for these tools, so we just use these third-party services to do this communication. It's customizable by some shell script so it's easily modifiable to some other companies that have probably similar systems. We try to participate with registrars on this validation – on this voluntary validation. So if a registrar will send his clients to this validation website, they can personalize the website by some logo or some information about the registrar. The level of this validation is visible for all registrars via EPP so they can know how these contact data are validated or valid. And it also can be seen in the WHOIS.

If the contact decides to change the data, of course they lose this validation and they must do the whole process again from the beginning. Because it's voluntary, there must be some marketing of this to promote the effort and one way is again through registrants. If they have a valid contact data, they have an option to hide address in WHOIS because if before – or if the contact is not valid the contacts can hide all the e-mails, telephone, and things like this in WHOIS but they can't hide

name and address. Right now, if the contact passed this validation procedure they can hide the address to close it from displaying in WHOIS.

The other option that we are thinking about is to do some marketing campaign and put some small gifts to these validated contacts like flash drive and things like this. The other side, there's also marketing to registrars to support this voluntary validation and this is through our co-marketing program which was quite successful in marketing DNSSEC. In the same manner, the level of participation of registry, how many money we will put into this co-marketing campaign is now dependent on the number of validated contacts. So this is how we're trying to promote, that the registrars should make some efforts to have the contacts valid.

Now, to the second approach, not voluntary but selective. Then in this approach, we try to randomly select a group of contacts from registry and do some automatic checking over this contact data, similar to what is in the SAC document syntactical correctness, checking whether e-mail exist, has some [MX] records. We have some publicly available registries of streets and addresses so we pass the contact data through this procedure. If there is some failures, then we go to manual checking by our client center that they will try to find out more information about the contact and maybe contact him directly.

The list of this individual automatically checks is quite extensible, so everybody can write simple module to plug in other check for the contact data. We do that similar way, like we are checking periodically our technical information. We are checking name servers through a lot

of tests like whether domain is reachable and there are different implementations [inaudible] name servers and things like this. So [we're] trying to do this similar attitude to our contact data to get them valid.

The part of this manual check is also written request to contact – if there is some possibility or if the client center thinks that the data are not correct, they have to write a request to the registrant to fix this problem. According to our rules, if the registrants will not fix these problems, we can suspend registration or delete the domain or all domains of this contact.

The last topic of my presentation is how to keep the contact data valid, which is not easy. But we think that one of the attitude that may help in this is trying to get rid of duplicities in contact data and registry. We did some survey or some statistic in our registry and we found out that about 15% of data are full duplicates. So we think that this is really too much and if the people have more contacts, they will probably not do many efforts to keep all of these contacts valid. If they change e-mail they will not do that in all of these contacts.

Of course, you may ask, "Why is that?" We don't know. I can think that maybe it's a way how registrars implement their interfaces, that maybe some of registrars don't have a possibility to associate existing contacts to the domain. They just force the new registrants to fill complete form of the data, or maybe it's a laziness of the users that they don't want to search whether they have some information in the registry. They just fill the same data again. So, duplicity is the problem.

So we implemented in our registry a new operation, merge contact, which try to merge the same contact into one and change all linked objects domains name servers to this new contact. Of course, during this merge operation the registrars and the registrants are informed about changes in their objects and we will run this procedure for all of our data in regular intervals. So we will try to keep our database small and compact but without duplicities, if it's really possible. Because we don't want to force it on the input site. It's not probably easy, and anyway we would have to somehow take care of all the data that we have in registry. So this is a new operation in the registry to merge contact.

And the last slide, some conclusion. We think that the registry, as a central point, is a suitable place for doing this validation of contact data. The registry software can help with that and our solution FRED has some features to support this validation. Although the users of FRED need some integration, some customization for this SMS and letter sending facilities – in the just-released version, there are that voluntary validation and merging functions. It's FRED 2.16 and that selective validation will be available in the system later this year.

So that's all about the work in validation in dot-cz, and maybe if you have any questions, I'm happy to answer those.

EBERHARD LISSE: Okay. Thank you very much. It's also a very interesting presentation. I recall we had in Brussels a few years ago a presentation how they do that. Nigel Roberts who is in the room has – technically I had a

presentation on how they are trying to validate addresses. So this is an obvious thing that is elucidating a lot of work in different venues, different registry systems.

In Namibia, we had quite a difficult problem. We had one South African registrar who didn't want to play by the rules, so we pointed it out repeatedly until they cancelled the Registrar Agreement and went to [inaudible] to an even more rogue registrar. Yup. They just flat out refused to abide by the rules and when we're done, gave them notice, went to all sorts of shady things to try to affect us, which all went to naught.

We have at the moment one registrar who is recently new to the game but he got EPP working and he hasn't figured out this out really yet. So we have to apply consistent pressure for them to not only validate their clients but just put in the proper addresses. They've just put their own addresses at the admin contact and it is a difficult thing to educate registrars because you don't have much leverage on them. The only real leverage you have is by terminating the agreement which is counter-productive because, first of all, you lose a client that register domains names and you need to find somebody to take over their names and that causes all sorts of other problems.

If you sort of lower the prices for registrars to validate their clients like you say with some incentives, that I think is a very good way. We have in our country no legal obligation to do this, but we think it's good to be a good citizen. We had one registrar who was involved with several internet pharmacies and it's just not good to develop a reputation that you host lots of rogue operators.

Nigel had his hand up. We have time for questions. I think this is an interesting topic, so please feel free to contribute.

NIGEL HICKSON: Thank you. I'm not going to talk about the validation work that we reported on the Tech Day in Brussels. That presentation is available and my colleague Mark who is behind all of that is available on e-mail.

I wanted to pick up on something that Eberhard said just now about particular rogue registrars. We had our first ever registrar failure about three or four months ago. Generally speaking, that was for financial reasons and we gave them a commercially, shall we say, a lot latitude hoping that they would pay the bills. But being in the small country, we found that – very quickly we realized that we weren't the only company that was having problems with. So we did exactly the same. We terminated the Registrar Agreement.

That had in a small ccTLD has really knock-on effects that you don't expect. You don't actually want to do this because end up with a lot of work. And the work that we ended up with was to do with the quality of the data. We terminated the Registrar Agreement, we transferred all the domains into a special registrar which we called [inaudible] instead of the previous registrar's name.

EBERHARD LISSE: We call it escrow.

NIGEL HICKSON: Yeah. We have one of those as well for different reasons. Then we found that about half of the domains, many of which were very from the earliest days of the registry, because this guy had taken over one of the very early registrars stuff – and they were wrong. They were all not just in his name but in the names of companies that didn't exist anymore. And terminating those domain names and so on, it causes real repercussions because there are real businesses in the islands who are using these domain names, they don't know that their technical people have got incorrect information. We turn them off, we're the bad guys. So we have to find a real diplomatic way to do it, so it's a lot of work and we don't want to have to do that. So, proactive validation is a very good thing.

EBERHARD LISSE: And to go back to what we did, we put them also into an escrow thing. We usually delete – we suspend after non-payment and then we delete after 45 to 60 days and in escrow we let them run for a full year before we look at deleting. And the ones that are still active, they come and squeal. We don't facilitate easy transfer to a registrar on an expedited manner so that they don't lose any web access. And that's only the ones that do this because they don't know any better. It's not that there is anything, any intent behind it.

If you have a registry like that with 800,000 registry, you will have a number of clients who do this for illegal, for criminal purposes. Whether you are legal law enforcement or not, you are somehow involved. You get visits from the local police. They say, "Look, we know these guys are

doing something. What can we do?" Then we come in, all sorts of problems.

So if you can motivate your registrars by way of saying, "Okay, if you validate it, we'll give you a discount." You don't have to give the discount to your client, not necessarily. So they would start figuring out, if they have got 10,000 domains, they put a little bit of effort and score some extra margin on that if they can automate that, force their clients to do that. It's actually financial incentive and that probably will work better than even trying to give their clients.

Don Hollander in the back. Please don't forget to identify yourselves to the microphone because we have remote participation who cannot see who is speaking.

DON HOLLANDER: Thank you. My name is Don Hollander from APTLD. I've been watching FRED on and off for a number of years and I'm just delighted to see how well it's grown.

Two questions about the philosophy as to when you indicated you do reviews or validations. You said you take a random sample that you do the automated validation, so why don't you just validate the whole list?

And then you said that you review things again automatically– you were looking at duplicate entries. You do that from time to time. And my question is why wouldn't you do that every night just to stay right on top of it?

JAROMIR TALIR: Yeah. Okay, thank you. So the first question – that random selection is just because this procedure is connected with some manual effort, so we just select, let's say, from the group of thousands of contacts and after this whole group of validation will pass through the all manual processes of our client center and the number of people in our client center is quite limited, then we will go to other let's say thousands of contacts. That's why I said it's a predictable procedure and that we will not do it all at once, if it answers your question.

The second was about the fighting against duplicities. Yes, that's exactly how we do that. I'm not sure if you we do it every night or once a week, but we are thinking too to find out the right period. It's quite customizable, so this is not the problem. We will do it as much as possible to make database clean.

EBERHARD LISSE: It's quite clear for us with 3500 names, it's easier to go manually but it doesn't scale very well if it's a manual effort to go to a registry with 800,000 names. Let's say 15 million, like in Germany.

JAY DALEY: Hi, Jaromir. Sorry, [inaudible] next speaker. You said that...

EBERHARD LISSE: This is Jay Daley speaking.

JAY DALEY: Yes, Jay Daley from dot-nz. You said that you are looking to remove duplications because you worry that duplications means that data get out of date or get inconsistent. Have you run statistical tests to prove that? Because I worry that removing duplications is quite an interference with the way that registrant or registrars might choose to hold their data, and I would need quite strong proof before getting involved in that.

JAROMIR TALIR: You mean to prove that – let's say that these duplicates diverse into time?

JAY DALEY: Yes, that's right. And that they become incorrect.

JAROMIR TALIR: No. We didn't do such a survey or such a statistical exploration. Let's say the common [inaudible] or we think it's the problem for the people to maintain, let's say [tenths] of the same data. And the one thing it's important that we are doing that – we are checking for full duplicates. That means inside one registrar. I don't know if you can bring some use this case that the registrar wants to have the same contact but several of the same data in the registry.

JAY DALEY: The issue is the nature of the registrar database. So if they have a different object for each of these contacts and you then combine those

internally in the registry but they don't combine those in their own database, then you'll get a side effect where they are unaware that changing one will change another and it's problematic. To me, the state of the data that we have represents the nature of their database. If I want to clean that, then I need to ask them to address their database and clean it that way rather than try to push into them.

JAROMIR TALIR: Well, one thing I must say is that we definitely coordinate this effort with registrars and we said how we will do that during the regular registrar meetings. No one did complain like this, so maybe they don't have these scenarios implemented.

EBERHARD LISSE: We find in our registry it's mainly ignorance or incompetence. They just don't know and if you tell them, "Do you know what normalizing data means?" they just say, "Huh?" Yeah. I mean when I started this 15 years ago on my skill, I didn't know what normalizing data means too so I also had every contact as one individual object. We don't really enforce it by agreement but we want to have as much data normalized as possible.

But what Jay says is correct. We cannot just now tell them, "We take one and delete all the others and move them over." The registrars have to do that because the registrars must not only do it in our database, they must also do it in their own database.

Jay is quite right. I never considered this really. There may be clients who want to have two objects or three objects for different databases

with slightly different data. But if the data is identical, then it's not necessary to have three objects, I think. But in any case, it's better that it's not done from the top down. It's better that it's done from the bottom up so that the registrars are aware what's happening.

[UNIDENTIFIED MALE]: Hi, I'm [inaudible]from CIRA, dot-ca. I'm just wondering if you expose the merge contact function via EPP, so it's something that [RARs] proactively do?

JAROMIR TALIR: Not yet, but we are thinking about it. Yeah.

[CHRISTIAN HASSELMAN]: Hi [Christian Hasselman] with dot-nil. My question is how many registrants that you get validated with this process so far?

JAROMIR TALIR: So far, we have right now we have about 700,000 domain or registrants or owners or technical contacts and registry contacts. And out of these we have not much, a few or something between 10,000 and 20,000 is validated.

EBERHARD LISSE: Okay. Thank you very much, that was very interesting. Let me just see who is next because that's about it. Okay now we have Jaeson Schultz

who is from Cisco. Okay. As you remember we had a presentation about bitsquatting by Nigel last time. Do you want to stand? Okay, sure. Sure.

He came across this it is on registry. The idea is basically not to register something that sounds like Microsoft but which one is the next bit wise character next to an "i" for example and then substitute the "i" with maybe "z". For human it doesn't make sense, but if there is random chance that a bit gets damaged and that bit gets damaged, the DNS will resolve to the Microsoft set.

So if you put a phishing site up that looks very similar, you might attract significant traffic because the Internet is growing, growing, growing, growing, growing and then get a little bit of advertising off so that is probably something where there's a little market. Jaeson has done deep research into this, so we are very happy to have him here. Welcome.

JAESON SCHULTZ: Thank you. Muy buenos días. My name is Jaeson Schultz. I work at Cisco as a Threat Research Engineer. For the past year or so, I've been doing research into bitsquatting.

A bit of an introduction. For those of you who attended the last meeting, there was an intro about bitsquatting talk, and very informative. I've done a bit further research that goes a bit deeper.

Bitsquatting is a form of cybersquatting. All of you probably know what cybersquatting is. Bitsquatting happens when a bit in memory will flip from a 1 to a 0 or vice versa. Anytime you write something in memory

and then the next time to read it, it's different, that's what's called a memory error.

So bitsquatting originally was coined – the term bitsquatting was coined by a man by the name of Artem Dinaburg at DEFCON in 2011, and he inspired me to do a little bit further research into the topic because I thought it was fascinating.

So this is for those of you who should – everyone should recognize this. This is the ASCII table and this is actually what makes a lot of the bitsquats possible. It's pretty fascinating, actually, the history of the ASCII table. It actually was created in the late '50s and early '60s before any personal computers even existed. In the 7-bit character set actually there's a few different – there's competing entries. You had the telegraph operators who wanted additional control codes and then you had other people who are interested in say including lowercase letters, and so there was a bit of a fight over what was finally going to be included. This is actually what it ended up being.

You can see in purple we've got the letter "r" and in red are all the one bit variants of that character, so just to give a bit of an illustration there.

So why do memory errors happen in the first place? Well, there are several reasons. Cosmic rays. There's cosmic rays passing through all of us right now at this very second. Heat – if you're in some place warm and you're operating your iPhone outside of their normal operating temperature, you can have bit flips. Most devices don't contain error correcting memory, so they'll passing any bit errors right through to whatever application is using that memory.

There's an interesting paper that came out earlier about detecting nuclear explosions using DNS because it'll spike in the incidence of bit error request. You'll be looking for domain that's just one bit different.

Then finally, defects in manufacturing. There's been cases where they have some sort of a radiation source which contaminates the memory chips themselves.

Bit errors have been happening for a long time. This is a photo of Chernobyl. There's an interesting story that I ran across during my investigation about this man who was working on a computer near a railroad station that the computer is designed to help route train cars and things like these. They were suffering from an intermittent bug and this guy was investigating. It turns out that every time they would get a shipment of cattle from Ukraine or Western Russia, these bit errors would happen in the computer, and the computer would crash.

So he suspected that maybe something was going on, had one of his army friends come out with a Geiger counter, and sure enough the cattle that was coming from there was radioactive. The idea was that the Russians were trying to make use of the livestock that was in the contaminated areas, mixing it in with the rest of the food supply so that, yeah, it would raise the overall level of radiation, but overall, it should still be safe maybe. After finding this out, he emigrated out of Russia.

EBERHARD LISSE: He became a vegetarian.

JAESON SCHULTZ: Right. Right. I thought this quote from him was great: "Cows were alive..." but it was definitely enough to flip bits in this computer. So bit errors have been happening for a long time and definitely as early as '86 and before.

In the RFCs for DNS they talk about the valid characters and the domain name. Here they are. One of the things that was missing from some of the original research that Dinaburg did was the dot. The dot character actually has a bit error that can happen where – you can see the difference here. You got a dot and by flipping one, you have an "n". What's the significance of that?

When "n" changes to a dot, you can actually end up with a different domain name. I've got a couple of examples here: windowsupdate.com. If you change the "n" and windowsupdate to a dot, all of a sudden you end up with the second-level domain dowsupdate.com. Similar with this symantecliveupdate.com.

I've got some examples. We registered some of these domains. Here's a bunch of people looking for Windows updates from me instead of Microsoft. Here's some more liveupdate.syma.tecliveupdate, Symantec traffic we were getting.

Bit flips are bidirectional. So you can also have a dot that changes to an "n". This is most significant when you have a third-level domain name and a second-level domain name. You can take the dot that separates the third and the second level and change it to an "n" and register it, which we did.

I have a few examples here: s.ytimg.com. That's actually a YouTube content delivery network. We registered snytimg.com. It's actually been a very popular bitsquat domain. We get requests for JavaScript and all kinds of stuff. So if somebody was malicious, they could wreak some serious havoc.

There's a few other examples here: mail.google.com, mailngoogle. There's many, many of these. I've got some examples here. Here's one for – which one is this? This is YouTube. Okay. This is snytimg. I'm getting traffic from YouTube. You can see in the logs the refer says youtube.com but it's coming to my server instead.

Another example is state.ny.us. We registered statenny.us. We were getting traffic for – in this example, from the office of Mental Health.

Okay, let's see. Here's another example here from Google Mail. So somebody out there using their mail, all of a sudden they're sending request to us for data to the completely wrong domain.

In Dinaburg's original research – this is a graph from his original presentation and it talks about the volume, and you can see that by far the most popular bitsquat domains happen in the context of HTTP applications. So I thought about this a little bit more and it turns out slash (/) characters can also have an effect, can flip a bit, and become something different.

You can see I've got an example. The slash, by changing one bit, it becomes the letter "o". Well, how can we use this? If the "o" flips to be a slash, then it's possible that you can end up having a different second-level domain. I've got a few examples here targeting dot-mil which is

not a publicly available domain, neither is dot-edu. But if you take the "o" and flip it to be a slash, that's going to cut short the domain. So you can see I've got bop.peostri.army.mil and you can register bop.pe. There's several examples here for you.

We registered some of these and there's one. This request actually you'll notice that it's asking for Apple Touch Icons. So both Android devices and a lot of mobile devices actually use these Apple Touch Icons to display icons on the Home screen. What this is is somebody had put this domain on their Home screen and then every time they click it, you get a request for an Apple Touch Icon. This is an example of that, but it's misdirected to me instead.

Once again, bidirectionally, you can have a slash that flips to be an "o". Now, normally this would be invalid syntax for the most part. Unless there is not third-level domain, you could end up with something like oslashdot, otwitter. The two slashes at the beginning are technically what you're supposed to have. Unfortunately, in a browser they try to help you out and correct any mistakes that might be present in the page. So if you end up with http:/ (single slash), the browser will helpfully add another slash or interpret what you really wanted and send you to the wrong domain.

Here's an example. You can see – I basically have a link here that points to http://slashdot.org, but in the HTML, I've changed that second slash to an "o" and when I hover over the link you can see the location where it's going to go is a totally different domain. Here's an example. Again, some Apple Touch Icons that people requested from oslashdot. I've

since given them – most of these domains have been returned to who we believe is the rightful owner.

More bad syntax. Let's see. We have one more type of bit flip with the letter "c". The letter "c" can change a single bit and become a pound (#) character which is syntax in HTML for an anchor tag. Why is this important?

If you end up with a "c' in your domain name, it can cut short the domain and become a different second-level domain. I have some examples here: pki.nrc.gov. That's the Nuclear Regulatory Commission. Bechtel, emergency.cdc – some fairly important websites or domains that actually have some pretty serious bitsquat availability.

Again, bad syntax is okay. If I end up with – in this case, I've got a link to the U.S. Coast Guard site, and if you actually change that "c" in "uscg" to be an anchor tag, you can see – where am I going to go? I'm going to go to cgportal2.us instead of the U.S. Coast Guard.

This is another interesting example because you've actually got a dot at the end of com, and then a "cn". Even if the "c" transforms into a pound character after the dot, it still is a valid URL. Here's where I'm going to go and it's going to take me there. So I thought that was pretty interesting.

This is just inside the domains themselves. The top-level domains also have issues. For the most of the gTLDs, there's only two: ".pro" and ".coop" with corresponding URL delimiter types like ".pr" and ".co" if you remove the "o" in that case. But there are several ccTLD bitsquats.

I'll just let you look at this for just one second. Some ccTLDs have no bitsquats – Netherlands, Paraguay, Uruguay – not a single bitsquat available there. There are some that have ten. Ivory Coast has ten different possibilities where "ci" can change to a valid ccTLD. U.S. has two, U.K. only has one. So you can see from this list that several ccTLDs have valid bitsquats in other country-level domains.

Here's a domain that we registered: kremlin.ru has a ccTLD bitsquat in dot-re (Reunion Island). And this was a request that received for a page at kremlin.ru and it came to kremlin.re instead. I've got an example of the page that they were looking for and at the top of the logs where somebody is asking me for something from kremlin.ru.

Here's some more traffic. Europa.eu has a bitsquat and ".mu". So we registered that and wow, here's a bunch of MX requests. I didn't accept mail for these but it's interesting to get the request nonetheless.

This is an example of SIP traffic. Voice over IP coming from a German domain that we registered. I've got another example of that, too, for bund.ee. So bund.ee is the German federal government website. Bund.ee we registered, bund.dm. Some of these requests are actually coming from IP addresses located inside the German federal government themselves. Yeah.

UNIDENTIFIED MALE: Probably [inaudible] cell phone.

JAESON SCHULTZ: Okay. The theme of this conference here – there's a lot of talk about the new gTLDs which are coming. I took a look at some of those to see what bitsquats are possible there. I've got a list of some of these here. These are n-based bitsquats where the "n" can change to a dot. You can see there are several. Some of the more interesting ones, .helsinki. If you were to register one of these, any domain that's registered under one of these gTLDs could potentially send you traffic. If the bit error happens in the right spot and that "n" in the top-level domain changes to a dot then you get everything. Of course, it has to happen in the right place. Not all bit errors always happen in the right place.

There's more gTLD issues with other letters, too. Like the "o" changing into a slash. You've got .boo, .bio, .cooking, several. The other one here, the "c" turning into a pound character, you can see a list of several proposed New gTLDs which have issues with that as well. It's possible for someone to bitsquat an entire gTLD.

This is an interesting case. The dot-uk ccTD has only one bitsquat for the top-level domain and it's in Tokelau (.tk). Normally, Nominet restricts people to – you can't register whatever you want .uk. It's got to be under some second-level domain, .co, .net, and so on. But if you look, there's actually several of these that are available, some of them not anymore. I've worked with the registrar to try to take care some of the most important ones, namely for example, mod.tk, ministryofdefense.uk. You go register mod.tk and who knows what you might get, right? As far as I know, some of these others may still be available.

How do you prevent bitsquatting in the first place? Originally, there was a few different ideas proposed. One is to install Error Correcting memory, so these bit errors are not passed through the applications. The other was to go out and register all the bitsquat variants of your domain.

Neither of these is very practical. Unless the whole world installs ECC memory, there are still going to be requests that are going to be bad. And sometimes it's cost-prohibitive to actually register all these domains, or someone else has already registered a bitsquat variant of your domain. And how do you deal with that?

So there's a few other mitigations that we came up with – four actually. The first one is to combat bitsquatting techniques that rely on the third-level domain name. It's possible – and I ran into an example of this – where if you divide your traffic among several different third-level domains, each one becomes less valuable as far as receiving bitsquats.

One example that I found from that was at Amazon. Amazon has a domain that they use cloudfront. Normally, the "clo" in cloudfront, if that "o" changes to a slash gives you a great domain in Chile. But what they do is they have something that recompiles the page and creates a new third level, and it's only there for maybe a few weeks or a month. So the odds of you being able to register that domain and actually receive something valuable go way, way, way down.

Another mitigation – because most of these bitsquats happen in the context of web applications – is to use relative URLs instead of absolute URLs. The less a domain name appears inside of a webpage, the less

chances there are that it's going to have a bit problem in memory. Most webpages are stored exactly as they're typed in memory.

You can use a base href and use relative links. Here's an example from Facebook, and you can see there's not a single relative URL in the entire page. They seem to want to go out of their way to use absolute URLs, which is probably why Dinaburg ended up finding out so many bit errors related to Facebook.

Another mitigation is using capital letters. Capital letters and lowercase letters are equivalent. But here, if you'll permit me, I just want to go back to our view of the ASCII table here at the beginning. If you'll notice, these capital letters don't have the same bit errors as the lowercase letters. For example, they're missing – you would actually have to change two bits in some of the capital letters in order to effectively have a bitsquat.

So by strategically using capital letters – there we go. So there are no bit error variants in capital letters, zero through nine. So if you're using capital letters, it's not possible for any bits to flip that result in a zero through nine. Neither is the dot a bit error of the capital letter "n". Neither "o," neither capital "c". If you strategically use capital letters inside of your links in things like these, you'll actually reduce the incidence of possible bitsquats.

Finally, the best mitigation in my opinion is using a response policy zone. I don't know. I'm sure most of you are familiar with RPZ. It's been around since BIND 9 – I forget the exact version. But there's patches you can apply to earlier versions of BIND which will allow you to do this. But

basically what you do is you set up a local zone on your on your own DNS resolver. When it receives a request, you can actually decide what to do with it. You can actually say to your DNS server, “If I receive a request that looks kind of like it might have been a bit error. Maybe I’ll redirect someone to wall garden that says, “Hey, are you really sure this is where you want to go?” Or just den y the request outright.

One of the problems with this is potential FPs. For example, PayPal has a bitsquat domain RayPal which is actually a legitimate domain. This guy sells music jingle or something. So, if you're going to deploy an RPZ, you need to be careful of any potential FPs.

There’s a script that I wrote and it’s available at this link here. Essentially, you drop in a list of fully qualified domain names. It spits out a list of all the bit error variants that are possible. If you want to, there’s a switch you can use to actually format that as an RPZ that you can just drop in to your DNS server.

Bitsquatting – the more devices that we’re adding to the Internet, I think the Internet of everything is coming very quickly. More and more devices, TVs and all kinds of stuff is getting connected. So actually the potential for bitsquatting is even greater than it was before. It also affects protected top-level domains like .gov, .mil, and so on. But it doesn’t necessarily – to protect yourself, you don’t need to mass register domains. We have several mitigations here, too, that I think are much better. So, thank you.

EBERHARD LISSE: Very cool. I remember in '91 when I registered dot-na, I was getting requests – we used [UCP], so everything went up over the link. I was getting requests for dot-nz and I had no clue why because I used the German keyboard where the "y" and the "z" are transposed. I didn't know that the "n" and the "a" and the "z" are next to each other on the keyboard. But that's not bitsquatting. But we were getting requests that were obviously directed towards a navy.mil request and somehow the "v" got cut off. We got a lot of requests and it irritated me so much that I figured out who was responsible and sent an e-mail to the admiral responsible. He was very upset and I think postings to Alaska followed out of that. But he was very upset that the small country or small operator was affected by this and he start kicking some [honey] to get this fixed.

I have Jaeson and Nigel but we take the remote question first.

UNIDENTIFIED FEMALE: Chris Davis from Secure Domain Foundation is saying, "I wonder if the implementation of zero times [20k] randomization in DNS would mitigate this?"

JAESON SCHULTZ: Could you repeat that? If the what?

UNIDENTIFIED FEMALE: Yeah. I'm not sure I'm reading it right. "I wonder if the implementation of zero times [20] randomization in DNS would mitigate this?"

JAESON SCHULTZ: No. I don't think so. I'm not sure if I understand the question exactly. Zero times 20. Roy?

EBERHARD LISSE: On this question? Norm Ritchie is there, who is going to present on the Secure Domain Foundation, so you can maybe discuss with him later further that.

JAESON SCHULTZ: I would love to.

EBERHARD LISSE: But Roy has an answer to the question.

ROY ARENDS: Roy Arends of Nominet. So what I think what Christopher Davis is referring to is what we call the 0x20 hack. What the 0x20 hack does, it adds a little bit of entropy to a DNS message by randomly uppercasing or lowercasing the domain name.

JAESON SCHULTZ: Right.

ROY ARENDS: So if you do that, you basically add entropy. The longer the domain name, the more bits of entropy you add. Now, in this case, of course if have less lowercase than uppercase, then you invite more bit flips basically, so I don't think it will influence that.

JAESON SCHULTZ: Yeah. I've search engine several requests where people are doing that – mixing the case in their DNS request. There was a presentation by Duane Wessels of Verisign where he studied the incidence of bitsquats, and it turns out that the checksums in UDP are effective at preventing at least network based bit flips.

The bit flips tend to happen in memory, for example, inside the machine. So unless the webpage itself is coded to have a mix of upper and lowercase then most likely, the bit errors are going to happen inside the stored HTML content in memory. So it's actually the incidence of bit errors happening through DNS itself is much less than through HTTP.

EBERHARD LISSE: Nigel?

NIGEL HICKSON: I hesitate to do this because I know my good friend Roy is here. But you mentioned the second and third level in particular of dot-uk. That may be the case today. But there's been a long-standing that's been discussed last year or two for allowing direct registrations under the second level. Potentially if that happens, and I'm sure we'll find out if it's going to happen relatively shortly.

JAESON SCHULTZ: Right.

NIGEL HICKSON: Then how do you see that would affect with dot-tk being such an interesting ccTLD as you know.

JAESON SCHULTZ: Yeah. Anything under dot-uk could be registered under dot-tk and you could take advantage of. But it's more powerful if you can actually register one of the official second-level domains because think of bit errors for everything that's registered there versus maybe one-off domains dot-uk will be something dot-tk. So it will still be there, but you won't get as much traffic from that. But I would love to see direct registrations. The other thing I would like to see is zone file from dot-uk.

ROY ARENDS: Hi, this is Roy Arends again of Nominet. Nominet happens to be responsible, if you will, for the UK namespace. It's very easy to get a zone file for dot-uk. You can [inaudible] if you like to. But I think what you're referring to is the zone file for co.uk, and that's the one we don't hand out. But that's that. Dot-tk is much larger than dot-uk.

JAESON SCHULTZ: I was unaware, really.

ROY ARENDS: So I think dot-tk allows or has allowed for a long time free registrations, and so I'm looking forward in registrations under dot=uk in the future if that ever happens, and under dot-uk to bitsquat any of the tk.

JAESON SCHULTZ: Tk, right. These are bidirectional. Absolutely.

EBERHARD LISSE: Okay, Patricio has his hand up.

PATRICIO POBLETE: Yeah. Talking about dot-uk and dot-tk. I'm Patricio Poblete from Chile. The "t" and the "u" are close enough on the keyboard for mutation like that to be typo rather than a bit flip. Because anything in your research allow you to tell which is which, what are real bit flips and not just typos?

JAESON SCHULTZ: That's a great question. Actually, it's nearly impossible to tell whether something is a result of a typo or whether it's the result of a bit flip. Unless when you're looking at the keyboard, obviously like the letter "r" and the number 2 are pretty far apart. The odds of you fat-fingering that are pretty low. But some of the best cybersquat, bitsquat domains, for example, the German federal government bun.de, I registered bund.ee which happens to both be a typosquat and a bitsquat.

So effectively, an attacker would double their chances, right? Those were actually probably the most valuable bitsquats that you could have, right? But it is hard to tell. Sometimes you can see, for example, in an HTTP request the referrer, and if the referrer is the real site that's pretty much guaranteed to be a bitsquat because someone hasn't fat-fingered the original domain that's coming to you from a referrer or from the real site. Those you know come from bit errors. Good question.

JAY DALEY: Hi, Jay Daley from dot-nz. I spoke to Artem Dinaburg about typos, Patricio, and he came up with a long list of URLs that are only generated by computers. So it couldn't possibly have been mistyped.

We're opening up as well dot-nz so that you can register directly under dot-nz instead of dot-co.nz. The issue for us are the bit flips of .co. So if somebody registered ao.nz – in fact, there's eight of them there –then they could get everything under what was dot-co or dot-nz going with that. So they can get 500,000 domains, only one hit. In the UK, that's probably about [inaudible].

JAESON SCHULTZ: Dot-uk is not the only one. Brazil has several.

JAY DALEY: Yeah. But if you opened up dot-uk and someone registered ao.uk then they're potentially 9 million domains that they get bit flip of in one go. So that's the bigger threat for flipping up is the second-level is now becoming [inaudible].

JAESON SCHULTZ: Now, we've been experimenting at Cisco with response policy zones that protect our own users. Of course, you can only protect the people who use your DNS resolver. But eventually, we would like to be able to publish an RPZ that protects just the most popular domains in general.

In fact, I have another project underway where I'm analyzing the Day in the Life of the Internet data from the DNS-OARC. They've collected three days worth of DNS traffic. I'm kind of wading through it. It's several terabytes, so it's taking some time to process. But there's a level at which a domain is popular enough that it will actually be valuable to generate bitsquats and then a level below that where you might waiting years before you actually get a bitsquat request. For some attackers, that might be worth it depending on who you're going after – a military site or something. You'd be content to wait for a year to get your foot on the door.

So hopefully I'll be able to follow this up with some other interesting research. Artem Dinaburg has a blog, too, where he talks about – he's looking at the DNS request that were coming in. Not only were the names being affected but the query types. He was showing a bunch of unknown query types. So really, if you think about it, the bit errors can happen anywhere in memory. Probably 99% of the time, it has no effect or it crashes the program or whatever. But when you have enough devices connected to the Internet, you're going to see a substantial amount of traffic.

One of the other things I've been tracking is solar activity. We're near the solar maximum, so we're getting lots of this coronal mass ejections and stuff directed at the earth, and I've been watching my domains and you can see little spikes when the radiation starts hitting the earth. It's pretty fascinating. I would like to hopefully put together a timeline of the events matched with the incidence of bitsquat request so that we could see how well it matches up. But that's some of the future research that I've been working on.

EBERHARD LISSE: I'm really lost to cut the pleasant discussion off, but we're running a little bit late.

JAESON SCHULTZ: Apologies.

EBERHARD LISSE: No, no, no, no, no, no, no, no, no, no, no, no, no, no. The discussion is very interesting but we're running a little bit late.

Jeremy Rowley, is he here? He is the next one. And then we have Jothan and Don Hollander. They're a little bit short so we will get this. Unfortunately, I don't think we have sponsored lunch, so I can't make you suffer by [inaudible] a little bit longer. But thank you very much. It was one of the most interesting things I've heard for a long while. We run [inaudible] tools, I can go and look what names are possible with dot-na and just block them. If they're not registered already then they cannot be registered under the good citizen program.

JAESON SCHULTZ: I encourage you to download the script that I have and just drop in some domains that you're interested in. You can even use the fully qualified domain name. I've got a list. Hopefully, I'll be able to connect with some of you who are changing your second-level domain policies. I actually had to go through each ccTLD and map out what are the official second-levels that you can register under so that I could actually shorten these FQDNs to the proper size and compute all the bit errors there.

So it's a very cool script, and as far as I know, it works really well. But it will have to be maintained as people are adding or removing second-level domains that are possible.

EBERHARD LISSE: Jaeson, if you want, we can take this offline because I would really would like to just – I'll give you a little bit more leeway anyway because I find it so interesting, but we're now actually running in a little bit of [inaudible]. Thank you so much for coming.

JEREMY ROWLEY: It's okay. My presentation is shorter.

EBERHARD LISSE: What's your thing here? There is something called Rowley here.

JEREMY ROWLEY: Yeah. It's probably [inaudible]. I'm Jeremy Rowley and I'm from DigiCert. We're a Certificate Authority. I'm going to be talking a little bit about some of the new developments in Certificate Authorities and what's going on as far as browsers and building trust online and what's going on there.

For a lot of people, security seems like a big divide, right? There's all these different players and some don't seem to interact very well. You have the browsers on one side and they're going about doing their business making connections and then you have the certificate authorities that are going about trying to make sure that you have the identity of websites vetted properly and that you've got a good [inaudible] connection going on, there's security there. And you have your server operators that just want to serve you the information and provide things. So a lot of times, there's a disconnect that goes on in the security world on what should be done.

There's some initiatives that we've been doing together with browsers and server operators to try to improve that that I thought would be of interest to you, guys. For example, one of the initiatives that we're rolling out is to try and push OCSP stapling heavily with server operators and get that implemented. That is a better revocation mechanism because it's faster. You don't have to do an out-of-band check with the CA. Everything goes through in that handshake with the server.

Some browsers, for example, Mozilla are talking about using mustStaple to require OCSP stapling from now on. That would be the only type of revocation checking you would do, so that you'd get rid of your out-of-bands check with CAs.

We have had some customers start to implement now Certificate Authority Authorization and that's where you put a DNS record in that says, "This is my Certificate of Authority." Only they're allowed to issue certificates. It's used to help prevent fraud because the CA will check that and say, "DigiCert is my Certificate of Authority." I shouldn't allow – or if I'm DigiCert and I see that GoDaddy Certificate of Authority, I will know not to issue that cert. It's really easy to implement. I think they're working on a BIND update so that it's even easier. But we've helped a few people with that and started to roll out.

Then Certificate of Transparency was a big announcement since the last ICANN meeting. Google has said that they're going to be deploying Certificate of Transparency for EV Certificates coming up shortly, and everybody will need to start including proofs either in their certificate or in their OCSP response. Certificate of Transparency works by shining a light into the Certificate of Authority world.

Right now you have about 60 or so or 20 major CAs that are issuing certificates out there that are all equally trusted. So if DigiCert issued a certificate and you might not know that that certificate exists. Certificate of Transparency requires all CAs to log that certificate in a merkle tree database so that you can go out and find those certificates and say, "Yes, I know all the scope." So if I'm Google I know all the CAs that have issued the certificate for Google.com, and whether or not that's authorized.

So that helps prevent mis-issuance and it also helps prevent government man-in-the-middle attacks, which in light of the NSA is pretty important right now. Because what happens is the browser will

talk to the different logs and make sure they're consistent. So that way, if the government puts a bubble around that player, it will say, “Hey, look, I’ve got this fake log that I’m reading. I’m in a bubble.” The browser doesn’t recognize it and it can warn the user that there’s a government attack going on.

Another thing that we’re rolling out and encouraging many of the new registries to roll out is key pinning. You can use that to limit which CAs can issue for it. For example, if you're running a dot secure, you could pin to DigiCert and say, “This is the [intermedia] site we’ll trust.” You could even pin to multiple [intermedia] so you could use [inaudible] and DigiCert or whoever and say, “These are the [intermedia] sites I trust,” and then those are the only ones that can issue off that domain. It narrows the scope. Sometimes you see reports there’s 600 different CAs out there issuing certs. Well, if you want to narrow the scope you can use pinning to get down to whatever number you feel is reasonable.

There is a potential bricking problem there, so it’s not for everybody but it is especially useful for registries because they’re often a little more sophisticated and won’t have this big of a problem.

So new developments and discussions that are going on online right now both in the CAB forum and other places is that there’s 1024 bit search are now deprecated. You can’t get them anymore for publicly-trusted certs. You have to go to 2048 as a minimum. Some of us are offering [496] as well.

Microsoft also announced that SHA1 is out. You’re going to have to be using SHA2 from now on, so you’re going to see – if you’re using SHA1

certs be prepared to get an e-mail from your CA saying, "Gotta switch." It's going to take effect in 2016 for code signing and 2017 for SSL.

Internal names are also now out. If you have an internal name that use publicly-trusted certs, you're going to get an e-mail from your Certificate of Authority telling you how to fix that, and that is of great relevance to everybody here who knows about the 120-day rule where Certificate of Authorities have to revoke any cert issued to a newly delegated gTLD within 120 days. So that's been taking effect and we've been revoking those certificates within the 120 days and ruling out so that these new gTLDs can be delegated.

We're also encouraging SSL be on [inaudible] and shopping carts. If you're always on SSL then all of your communications are protected and you might as well not let anybody listen to what you're saying.

As I mentioned, CAA is being deployed by CAs and so is CT. Google should in fact come out with a browser soon that recognizes CT proofs. It will be a dead browser, but we're working on that with them, partnering with them on that.

Some projects that are going on is we're right now looking at how to expand EV certificates. EV certificates are the high validation certificates. You know who the person is to be verified. All of their information with an independent third-party database, you know their name, you know where they're registered, you know their telephone number, you know all this information about them. But right now it's limited to primarily U.S. ones who are easily verified – who are verified with DNB or other reliable sources.

So what we're looking at now is how to expand that to sole proprietorship and then, eventually, hopefully everybody will be able to get an EV certificate as long as you're a legitimate business or a legitimate website owner.

There's also Performance Working Group. We've heard people complain a little bit about that moving to 2048 has slowed things down because of the bigger certificate size, so we're looking at how to compact certificates, make them faster and smaller, especially with revocation. Some people have fairly bloated revocation site – OCSP or response sizes and we'll look at making that smaller.

We've also working on some new requirements for new baseline requirements for code-signing certificates. So right now there's been the number of – the amount of malware that's actually signed by code is actually on rise, especially coming from certain countries, primarily Brazil and China, and so we're looking at ways to reduce that malware and that way code-signing can be a more effective deterrent than it is now. Right now it's kind of everybody for themselves on what process that they follow when issuing code-signing certificates, so that's going to be an exciting new change. It will be the first standards there.

We're looking at moving certificate lifecycles. Right now you can get a ten-year cert. Some people have ten-year certs. That's supposed to go down to five-year certs by 2015. But right now we're doing a push with the SHA2 initiative to see if we can get it down to 39 months sooner, and so all certificates will only be valid for 39 months. That will help prevent key compromise and problems and problems with domain change – if the owner of the domain changes, the old owner doesn't

necessarily revoke their certificate and so you have a problem where the old owner might be able to do a man-in-the-middle attack on the new owner and that's never very good.

Again, I mentioned the OCSP stapling push. We've worked with NGINX. They've got it theirs. I think Microsoft had it enabled for a long time and pretty much all of the major – most current version of server operating software has it available. You just have to enable it. It is not enabled by default. We're trying to get that enabled by default so it's easier to use.

Now, I talked a little bit about the collaborative nature. These are some of the things that are – everybody's working on together to improve the Internet. For example, CAs have talked about that we're looking at better issuance practices. We've got baseline requirements now for SSL certificates. We're looking at improving those. We're looking at improving EV. We've got better standards for network security and we're looking at improving implementing new technology like certificate of transparency in CA.

Browsers – right now looking at how to enforce those good practices, how to get people onto better versions of TLS and a lot of people are still using SSL3, which is not very secure so if we can get on TLS 1.2 that will be great. They're setting higher standards and deploying new technology. We've got server software providers that are looking at how to deploy TLS 1.2 better, enabling OCSP stapling now and they're providing other enhancements.

And you have the ICANN community which is doing a great job, I love the new WHOIS update that they're doing. That will provide better

information for CAs and for everybody else to know who the domain owner is. Now one of the things that we commented on that is it would be nice if there is a push change mechanism so that CAs can receive notice if that domain owner changes. That way, we can revoke the certificate and it won't be that man-in-the-middle problem I described earlier.

So the path ahead is we're going to continue looking at improved online security through new technologies and better standards. Many of us are active participants in IETF in watching and see what happens there. There's been some bumps in the road with Legacy devices and software, and that's causing some problems with a lot of the changes with some things. Especially in Japan, a lot of stuff can't use SHA2, so it's causing a bit of concern there. A lot of people don't – there's a resistance to change because they just don't see the need to move perhaps to 2048 or SHA2.

There's always the chicken-and-egg problem with digital certificates and new technology. We keep coming into do that, whether or not to go with early voluntary adoption and see what happens or go with mandatory and push everybody towards better security practice. DigiCert right now is trying to break that chicken-and-egg problem by going right now to SHA2 and also implementing CT for our certs. So look for that. If you want to test it out, you can use your certs there.

Then we're looking at improving transparency in the certificate process, making sure everybody knows what's going on. We want more visibility in certificate issuance and enhanced information on what CAs are doing and how we're working to improve online trust.

So if you have any questions – sorry I talked so fast in the beginning – but if you have any questions or you want clarification, you can contact me there. There's additional resource available at the www.cabforum.org and www.casecurity.org. CAB forum is primarily where a lot of this work is going on. Unfortunately, it is not – you can join as a interested party but that doesn't give you voting rights, so it is kind of a closed group. They do have an open mailing list though if you'd like to follow a discussion. If you'd like to be involved in working groups, you can join by signing an IPR. Sign in at the CAB forums. IPR would do the [RANDZ] policy and participating there.

So, does anybody have any questions?

EBERHARD LISSE: Come on. There must be a question. Okay. Oh, Jay has one.

JAY DALEY: Does DANE scare you?

JEREMY ROWLEY: No. DANE is an interesting proposal because there's two implementations that are great, that we like a lot, of course, because it requires publicly trusted certificates. The reason DANE is not on the side is primarily because DANE isn't being implemented anywhere at this time. Warren is going to correct me perhaps. But as far as I know, no browsers are planning on supporting it yet.

EBERHARD LISSE: Okay. Thank you very much. I don't really want to cut this short, but we're running a little bit late. If you want us to go into lunch, we can do that. Then Warren is the next. I think Warren goes first and then Russ. I'm going to take only those two questions.

WARREN KUMARI: Thanks. Warren Kumari, Google and also co-chair or the DANE Working Group. Yes, not many browsers have actually committed to it. There is some more interest now after the Snowden stuff.

There are some plugins for Firefox. There's also Bloodhound which is a Mozilla clone. It is also being implemented now between MTAs so [inaudible] has it in – it's not quite unstable yet or maybe as soon as the current beta gets released it will be unstable. And [XM] is also adding support, so it's getting a fair bit of traction in the mail.

JEREMY ROWLEY: I want to clarify that I think DANE it's a good idea. There are a lot that CAs – especially if it's combined with CAs because CAs do a lot of the checking that browsers don't.

EBERHARD LISSE: Okay. Russ Mundy from Sparta.

RUSS MUNDY: Yeah. Just a quick question. Sparta – Parsons now. New name. Company bought us, whatever. That Bloodhound, as Warren mentioned, does do

DANE right now and full DNSSEC. I just wanted to identify that where you go get it, it is public and is open source. DNSSECtools.org, and publicly available, downloadable – easily installable if you're using a Mac.

EBERHARD LISSE: Alright, any other questions? Good. Thank you very much. I don't really like to cut the discussion short, but in the end it worked out well. Thank you very much for coming and for willing to present. Jothan Frakes will talk to us about the Public Suffix List.

JOTHAN FRAKES: Hello, everyone. Thank you for the opportunity to speak with you today. I'm here – I think many of you know me. I have the privilege and the opportunity to know many of you as friends, many of you as colleagues. You may have seen me present this before. I'm here as a volunteer with the Mozilla Foundation. I caught very early on working with ccTLDs that there was a disconnect often with what browsers thought were TLDs or not, and I found an opportunity to volunteer and jump in with the Mozilla Foundation with an initiative they had called The Public Suffix List.

It's a publicly available list. It's available at PublicSuffix.org. It's essentially a list of known suffixes. It's hosted by Mozilla and it's modified and updated by community volunteers like myself, the majority of which come from the developer community and not from our community. I saw an opportunity to volunteer on this and to

advocate within this community to basically keep your listings up-to-date so they work correctly in a variety of different places.

Essentially, what it is where the IANA list would show an individual top-level delegation, The Public Suffix List simply is a bit more elegant and that it goes and descends into your authoritative stub zones or sub domains.

It was designed to aid browsers to help with cookie delegation, actually, and then it developed over time into much more where search engines use it and furthermore it became used by application developers and libraries. Because really there's not, aside from the IANA list, any form of list other than balkanized efforts to have some sort of a comprehensive list like this. And in fact, as Jaeson pointed out, it'll probably a very much help in his effort with bitsquatting.

So it's used within the browser market. So a variety of browsers use this and collaborate on it, but it also trickles out into different – there's libraries available in the majority of the popular development languages out on the Internet and it's also used in operating systems. Hosting companies use it. Statistics companies, anti-spams, security services, SERP providers use this, form validation – there's just so many uses that this trickles out to.

So for you in the community who see that browser may not be behaving correctly or there's something that doesn't look correct, I would simply encourage you to go and take a look at this. You can make a change in one place. It doesn't trickle necessarily real time but it trickles out rather quickly into the development community because it's a list that's

updated very frequently and it's kept fresh and maintained through submissions from the community.

They do validate this information. They make attempts to validate this information with the authoritative registry, so if you would take the time to go and validate the information that's available at these lists, make sure it coincides with what you believe to have your stub zones so that it accurately reflects what is used within your particular zone. I know each of you have potentially different approaches. I'm not saying good or bad. I'm just saying that they exist. This allows the software to be aligned with sophistication that you offer within your particular zones.

So the developer community has a huge disconnect with this community and I'm here to just sort of help and ease those together. I would encourage each of you just to be very brief so I can get Don up here and let you all onto your lunches. You can view The Public Suffix at PublicSuffix.org. You can review the list. I would request that you do please look at your entries and then submit changes and also actively monitor to see if the listing for your contact information is accurate so they would be able to reach you to validate if changes are submitted.

Changes can be submitted by the NIC. We prefer that that happens because it's much easier to validate. However, they do often come from the community. And if we could have a way to contact you – I know many of you, so if I see a change come in I know I can reach out to you and validate that they're accurate but we do want to make sure that the information is accurate there.

I have links built into the presentation. You can find out more at PublicSuffix.org. I thank you for your time.

EBERHARD LISSE: Thank you very much. Before I take some questions, all presentations will go on the meeting website. So it's not a problem getting access to the link. Any remote questions? Any questions from the floor?

JOTHAN FRAKES: If I could add something. Part of my efforts with volunteering with Mozilla was to do away with the white list that they maintained for IDN at the second level or third and that's been done away with in favor of essentially a default of working prior to displaying a punycode. That's a big win for this community here. So we've been able to accomplish that as well.

DMITRY KOHMANYUK: Just two quick comments and maybe a suggestion. I used the public list, it does work. Well, two comments. While the turnaround time for my request was close to two months, which I think is a bit sub-optimal so I would suggest to maybe bring it down to a week or so. And the second I would probably request some kind of notification service so if you are the NIC representative or maybe just anybody who can subscribe to an updates for the TLDs of your choice, so if somebody else would add some suffixes that you don't want to be there or like then you can just be notified and corrected. So I guess that's it then. Thank you for the good work.

JOTHAN FRAKES: Thank you and thanks for keeping your entry listed. I'll bring that feedback back to the team. Yeah. It is of course going to take some time. It's consisting primarily of volunteers and folks putting it into their time schedules, but we try to be as agile as we can.

WARREN KUMARI: Warren Kumari, Google. I want to mention there is work on IETF to just publish all of these in the DNS itself. That way it's authoritative, it's only in a single place. You don't need to wait for somebody else to sort of figure out if your authoritative if you can publish at a specific place by definition. And so Andrew Sullivan has a draft, Asserting DNS Policy Realm Boundaries: The SOPA Resource Record (draft-sullivan-domain-policy-authority), and there's going to be buff in London I believe. There's already been a fair bit of discussion on this about this but it seems like a much more scalable.

JOTHAN FRAKES: It certainly does just like DNSSEC. And so each zone would manage their own stubs or some list of stubs and that would be wonderful to have that because there's no question to the authority of it being directly from the top. In the meantime, this effort would supplement that and help universal acceptance of these domains. Thank you, Warren.

JAY DALEY: I was also going to ask for a notification service and if possible using the technical contact in the IANA database.

JOTHAN FRAKES: Well, that was one of the – I had it on a slide that I passed. It's often a request fall into the bit bucket or the ether, but I'm sure that you get a lot of e-mail to those e-mail addresses.

JAY DALEY: Well I still prefer it wasn't separated from that. And the other thing just to point out to anybody who's done this is the syntax as a public suffix list is weird. Okay. There is a silent asterisk, and so the first time you look at your entry, you'll think it's wrong and it's not. It's right. And then the fourth or the fifth time you might understand it. And every time you look at it again a year later you'll think you've done it wrong. It's the weirdest syntax.

JOTHAN FRAKES: Certainly. I certainly appreciate that.

JAESON SCHULTZ: This is Jaeson Schultz. Thank you for that link, by the way. That would've saved me probably two days worth of effort hunting down all of the official second-level domains and so on. I'm just looking at the list. I noticed you have a lot of gTLDs which haven't been approved yet.

JOTHAN FRAKES: Yes. And this helps in the display and behavior in many software applications. So when info museum travel some of the TLDs that have a longer string – there were folks who had TLD validation that was based upon the length of the string and really didn't have any elegance to it, and so that empowers the elegance in advance of when the delegation occurs but post contracting is what we're able to work out with ICANN. So we're trailing within 20 to 30 days of contracting. We have a good block of those names added and it aids with the variety of different places.

JAESON SCHULTZ: So some of these gTLDs are list that are ones that are scheduled to be approved.

JOTHAN FRAKES: That's correct. Yeah. If they've gone through this – so ICANN maintains a list and I have the link inside of here that explains where this originates but we follow the contracting of TLDs. We're on that mailing list and as those trickle in, we keep track and we try to keep it within two to three weeks of when those come in, we get those submitted to suffix list as well. It's been able to allow, for example, Firefox work with all the [dot-schubacas] and all those right out of the box. Chrome – you have to change some language settings but it worked just perfectly right out of the box with the new TLDs and that's not been a trivial effort on any of our parts, so thanks for noticing that.

JAESON SCHULTZ: Yeah. No, that's great. And one of the things I didn't touch on in my talk was IDNs, the dash and the M, so there's actually any IDN at least on the left hand side. There are no bitsquats in the TLDs part but the second-level domains you change one of the dashes to an M and you've got whatever you want.

JOTHAN FRAKES: Oh boy.

EBERHARD LISSE: Alright. Thank you very much. I'm a little bit lost to cutting discussion short as I said. But I would like to get Don Hollander another chance to do his presentation before lunch. And in any case, if there are questions, Jothan is around. This can always be done offline anyway.

DON HOLLANDER: I didn't know that I'm going to have slides. Look at that, how efficient is the system.

My name is Don Hollander and I'm the new and old general manager of APTLD. So I used to be the general manager of APTLD and then I wasn't the general manager of APTLD, and now I am the general manager of APTLD, and at some point I won't be the general manager of APTLD. So this is really just to introduce myself. I used to be something of a geek but more of the applications level than the network level. So that's APTLD.

This is the slide that I'll talk to in about two minutes tomorrow so you can save it for then. But one of the things that APTLD will be doing in the New Year is we'll have three meetings. One of the meetings is geeky and that's scheduled for February in Bangkok as part of APRICOT. If you have any ideas as to topics that we might cover that I don't know about, I would very much welcome suggestions of issues and opportunities in the tech field that ccTLDs certainly in Asia-Pacific might be interested in hearing, so with that you can have lunch.

EBERHARD LISSE: Just a little housekeeping. We seem to have been able to arrange for a translator. We haven't got the headsets yet, so each of you must pick them up on your way in. I will make a plan because it's going to be at 4:00. So after 4:00 you should be able to get that sorted.

Alright. We'll be here then at 2:00.

UNIDENTIFIED MALE: November 18, ICANN 48, La Pampa, 2:00 PM.

EBERHARD LISSE: So good afternoon. Welcome to the afternoon session. Of course, as usual the participation is [inaudible] much less than before. But usually traditionally one punishes the ones that come for the ones that don't come.

The next presenter is Anne-Marie Eklund Lwinder – complicated name – from dot-se. She will talk about Information Security Management and

Certification. I like this idea about certification like ISO 9000 or whatever it's called and this one. I haven't got much insight into this myself, but I find even if you don't certify the principles that are used that you have reproducible processes that, you know, if something happens you have a process. You know how it works, you can figure out what happened and why and so on. It's quite a good thing. And the bigger ccTLDs like Sweden with how many domain names? More than a million?

ANNE-MARIE EKLUND LWINDER: There are about 1.3 million.

EBERHARD LISSE: 1.3 million, they need this also probably because of regulatory requirements. So without further ado.

ANNE-MARIE EKLUND LWINDER: No, actually not. We're doing it because we are so good.

EBERHARD LISSE: That's, of course, evident.

ANNE-MARIE EKLUND LWINDER: So good afternoon. Welcome back from lunch. I'm Anne-Marie Eklund Lwinder, and it is a complicated name. I work as a security manager at .se, the Swedish registry, and I will very, very briefly take you through what we are doing and I will talk a bit about Risk and

Information Security Management in a ccTLD and explain a bit about 27001 – what, how, and why you should do it.

This will be kind of a crash course, but I don't expect you to be able to read everything I have on my slides. But I put them up because I would like to give you the opportunity to get back to them later and say, "What was she talking about anyway?" So we are a foundation and a nonprofit organization running not only dot-se but we also run dot-nu. This is a little bit about us in figures from last year.

I used to bring these up when I want to remember what are our responsibilities. And that is from the RFC 1591. "We are the trustees for the delegated domain and have the duty to serve the community." And trust, as you all know, is something that you have to deserve. It's nothing that you just gain or get from anyone.

So how do you do that? Well, one way of doing it is try to work in a good way, efficient and safe and secure. And information security is an everyday work. That is a matter of skills and knowledge. It's not about luck. I mean, we can be lucky sometimes to get out of problems, but still you need to do this very properly to gain availability, information integrity, confidentiality, and not to mention traceability.

So the tasks of a TLD registry, this was put together by Wim Degezelle from the CENTR secretariat for another occasion, but I thought it was kind of a good say to look at what a TLD registry is doing from a day-to-day work. And the domain name resolution service. We have the registration service. We have directory providing WHOIS information.

And we have traditional business service as anyone else, like billing, customer support, and sometimes dispute resolution.

The TLD part, the domain name resolution part, that is the key function, and that is something that we are really good at. We are doing it extremely well. We have been working with it for a long, long time, and we know exactly what works and what doesn't work. And we are very accurate when it comes to infrastructure and building that in a good way.

So the other ones – I mean, the regular business security needs – the only one that we have that deserves really special attention, that we give really special attention, that is the domain resolution service. But we shouldn't forget about the other parts, the regular business class security needs because they are important to. They are important to be able to show publicly that you are taking care of the customers in a good way.

Bruce Schneier used to say that the Internet is too complex to secure. One of the reasons is that it is too complex to understand. And that is also true if you consider taking it into one chunk, but don't do that. I mean, if you just take things apart and you take them one-by-one, step-by-step, I think you are on top of things.

The methodology I use to work with is know your business needs. I mean, if I don't know what my business needs, I don't know what to do. And then you have to identify assets and threats, secure and harden, monitor and detect, respond if that is required, and learn and improve.

That is, you know, the old PDCA process, and then you start all over again.

This is a cite from a book "Signposts in Cyberspace," and the reason I took this in is I will once again point out that operators of the DNS are responding to all the threats they see, but the rest of the organization, the rest of the business, also have their threats. And we tend to forget that part sometimes.

So while a ccTLD are a relatively small infrastructure provider, but it's used by most Internet users. So that is our responsibility to make sure that it is usable all the time. And current practice on high level security resilience are in place to ensure the DNS part of the function that we have.

And we all have all these good active platforms to exchange experiences and knowledge, which is good, but we also have human resources. We have the information to take care of, other kind of information than DNS. We have hardware and software, and we have premises that we have to protect. And all these are part of an information security management system to be able to go through all the different details.

Sometimes when I start working with this, it's a long time ago I decided to go with ISO 27000 not to certify in the first place but to have, as [our chair] mentioned, to have this framework to work with to make sure that I don't forget anything. Because sometimes – and I know I'm not alone – I tend to prefer working with things that appreciate more or that are more fun, and I tend to put on hold those things that are not so

fun. So to avoid that, a framework that takes in all the parts is very good actually.

So why should an organization go through an ISO 27000 certification? Firstly, as I mentioned, it's about trust. It's a way to prove and show others that you are taking information security seriously. And it also makes you work with a continuous information security improvement cycle, which as I said makes sure that you don't forget or skip something that is not as fun as DNS, for instance. And the working processes are very structured, and it also builds up your image as a serious TLD.

And there are two parts of the standard. It's 27001, which is about what you should do, and 27002, how you can implement the controls. So there are two different parts that match each other.

Before you start, I recommend that you get full support from the CEO and the management team because if you don't have them behind you, it's very, very hard work to make this happen and come through. And you should have a good reason for certification, and a good reason is we feel responsible for our customers and their domains.

And work together with your colleagues. You can't do this alone. You can't stay in your room, try to work out this information security and management system on paper and just believe that it will happen. It's not how it works.

And build your own. Don't take anyone else's because if you build it suitable for you, means a small organization needs a small ISMS. Very easy.

And keep it simple. Do not complicate things. You don't need that. Be very practical. And choose the right auditor. Make sure that you have an auditor that not like to punish people but to help you to come through this process. I mean, they can be very, very supportive, but they can be very "audit-ish."

The things you shouldn't do is, again, don't do it on your own. And don't lose yourself in the risk analysis swamp. That is easily, easily done because risk analysis is something that the management team mostly have heard about, and it's easy to get caught in that. You do the risk analysis, but you don't follow up later and you don't have the [mitigatings] that you need to either eliminate risks or just manage them in any way. And don't have too many Key Performance Indicators (KPIs) because it's too hard to measure and follow up. Use a small number and change them from time to time just to make sure that you know where you're standing.

So the fundamentals of the standard is that it is a model for establishing, implementing, operating, monitoring, reviewing, maintaining and improving an Information Security Management System (ISMS). It is, indeed, a strategic decision to get on this track. And it should be scaled in accordance with the organizational needs, as I said. And the requirements are separated in eight different chapters, and there are some annexes to it too. I will not go through that, I promise.

And here's a sort of PDCA cycle for the ISO 27001 implementation. I mean, you can use this wheel for a number of things. But when you're Planning, you establish the management system. You have the risk

analysis. You define the scope, which is very important. You have to restrict the scope of the certification. Define GAP and risk mitigating control objectives.

So the next step would be Do, and that is when you implement the management system. And here something happened in Adobe, I think. So because it has cut off the text on the Check and Act. I have to mend that later. I'm sorry for that.

There are some very strong documentation requirements, and the documents must be not only written but identified, reviewed, approved, which you should have versions, revisions, and you should make sure that you have distribution that is correct and updated that people are knowing that you have an information security management system, that they are aware of the content of the system. And you should have control of records.

So the most important documents that you put up is the scope, as I told you, the statement of applicability, which is also known as SOA. I know that the DNS family will get confused when you start to talk about SOA, but this is another SOA. And you should have all the related documents in a good order.

This is our scope. This is what we put up as the scope for our ISMS and the certification. And that is the administration and technical operation of the national domain name registry for the Swedish top level domain dot-se.

So we haven't certified the entire organization. We picked the part that we think is most important in regards to our customers. We have a lot

of other things that we are doing like running projects. We have conferences. We have education. But that is not part of the certification yet. It will become that because I found out that it was much harder to restrict the scope than to take in the entire organization.

And management is responsible for a number of things, and this is I believe the most important commitments. You have to have at least four meetings a year about this, and there should be minutes from them. Everything has to be traceable so the auditor can come to you and say, "Can you show me the minutes or the protocols from the meetings that the management team has had about this?"

And resource management, provision of resources, training, awareness, and skills – so you have to have education of staff. You can choose any way you like to do it, but you have to do it and you have to be able to prove that you've done it. So if you have an educational program, you have points, what you've gone through or you're having a test or whatever. Just show that you have done it.

I can do internal audits myself, and I do that regularly. They should be in planned intervals too. And the goal is to determine whether controls are effective, if they are maintained, and so on. And we have an auditing process that we are following. Yeah, I think I'm running out of time. This is a bit of a small detail from chapter six about the internal audits.

The management review, that should take place at least once a year where you go through everything you have like the feedback you got from earlier, the audit, and vulnerabilities that you have seen lately, changes. And the output from that is, again, minutes and

documentation. Documentation is a big part of this, of course. So continuous improvements. You have to audit and improve the system itself.

The certification is valid for three years, but you need to do an annual update to make sure and be able to keep the certificate, which looks like this in our case. And then you can brag about it on your website. You can put this logo on and let people know that you are working hard with information security.

So that is very, very, very briefly about 27001. And my recommendation is start working on it. The goal itself is not –I repeat, it's not – to get certified. The goal is to have an information security management system that works for you. Thank you.

EBERHARD LISSE: Thank you very much. As I said, for small organizations, I don't know whether there's a cost involved in getting certified whether who is the actual authority signing this. But we discussed it I think already in Mexico on the Tech Day. We had a presentation by a risk management company, what in more detail needs to be done to sort of anticipate and mitigate risks – hot standby, warm standby, cold standby – these kind of things.

It's not so much that you have a certificate but that you have got a process in place where you start thinking about what you're doing, how it can affect your operation if there is a failure, and how you plan about it. Any questions? I know that the Czech registry dot-cz has done this. It's another registry that has done it. Denmark, dot-dk, has done it, and

dot-nl have done it. What's your experience? And dot-eu. What's your experience? Have you got any contribution?

ANNE-MARIE EKLUND LWINDER: I mean, the cost is a fixed fee actually. I think it's restricted. ISO have stipulated that there's a fee connected to the certification, but it's fixed so it's not that expensive. We do more things that are more expensive than that.

EBERHARD LISSE: Alright. Thank you very much. When we planned this, we decided to give Anne-Marie a break because she has got two presentations, and I don't really like people to speak two presentations in a row. It's also a bit of a strain on the voice. Therefore, the IETF is going to give us a short update on what they want to say.

JARI ARKKO: My name is Jari Arkko. I'm from the IETF and am the chair, and I prepared a small update on what has been going on at the IETF in the last meeting. We have Russ Housley, chair of the IAB. And this is maybe not exactly on topic on the kinds of things that you normally deal with in the country code world but, I think, very interesting nonetheless.

So what I'm going to do is go over some high-level topics that have been hot topics a week ago in our meeting in Vancouver and then talk about one of those topics in more detail, which is the pervasive monitoring surveillance issue. In this week, you have probably already seen and will see some sessions where people talk about Internet governance and

how the Snowden revelations affect that. This is not about that. This is about the technical aspects of that and is there something that us engineers can do about it.

So the topics that, at least from my perspective, were the highly interesting ones in addition to pervasive monitory was HTTP 2.0, which is a relatively big redesign of the protocol we all know and love, and more efficient, hopefully more secure. TLS

1.3, also a similar redesign, fairly big change again. It's going to be more efficient. We're trying to cut down on the setup time, for instance. And it's going to be more secure – remove the weaker algorithms, adopt some new procedures that protect identities better and so forth.

We also talked about WebRTC quite a lot during the IETF week. WebRTC is this technology where you can instead of having an application for a voice over IP thing you can do this in the browser and without any plugins.

And we also talked about the evolution of transport protocols, looking at things like TCP and TLS and HTTP together. Can we optimize that even further beyond the things that are done by HTTP 2.0 and TLS 1.3?

So, I mean, in many cases what we do at the IETF is we tweak the details and worry about [big fills] here and there, but I at least felt that during the last meeting we were really dealing with a number of pretty fundamental issues that will have a long-lasting effect on the Internet. We're going to have to live with the new version of HTTP for a long time, and hopefully out of the pervasive monitory discussion we'll get some new security features into the Internet protocol stacks. So I think

it was a very, very good week and interesting discussions and heated discussions at times, of course, but that's [inaudible]. We need to get the issues out.

But on to the pervasive monitoring topic, I just wanted to set the stage a little bit first. What is this? Is IETF now a forum for political discussion? We're not. We don't really. I mean, some of us might have opinions about things, but the IETF is not the forum to voice those political opinions or opinions about what is being done or not being done and so forth.

And I'd also like to remind people that if you read the newspapers, you see a particular picture. If you actually perceive the world, and many of us are involved in networking around the world, you should realize that it's actually a wider problem around the world than what is being perceived in the newspapers.

But that being said, we at the IETF should understand what dangers in general are facing Internet traffic. And if there's something that we can do or should do about improving the Internet security and Internet technology in general, then we should really go about that. So that's the general scope.

And we've had a discussion in Vancouver as well as on list before and in various other discussions about how should we perceive the surveillance or pervasive monitoring. And the general consensus has been on this is that this should be considered as an attack from the perspective of Internet protocols, or at least it is something that is indistinguishable from attacks.

So we don't really know what the information is used for. It could be used for good or bad. Consider thieves stealing passwords. And then there's this consideration that vulnerabilities or security problems tend to democratize over time. So if today an intelligence agency can do this, then maybe tomorrow a criminal group can do it and the day after tomorrow your neighborhood kids can do it. So if there's a vulnerability, we actually need to deal with that in some manner. So we do consider that as an attack.

And just a little bit of technical information. Obviously, we don't have information about exactly what has been going on with the NSA, and in some sense it doesn't even matter necessarily so much because it's just one example. But the likely attacks that sre possible, first of all obviously, if you have unprotected communications, you may be vulnerable to people listening in. Well, yeah, but the Internet is largely about unprotected communications today, so maybe there's a lesson there.

Direct access to the peer. So even if you have perfect communications security when you're talking to someone, if they leak your conversation then that might be an issue.

Direct access to keys. This might have been the case with Lavabit that was going on in the U.S. awhile ago. Third parties could be used to make an attack. For instance, fake certificates and routing information and sort of inserting yourself as the man in the middle without anyone knowing.

Implementation backdoors. And this is really huge actually too. So different pieces of software or even hardware might have backdoors or vulnerabilities that the rest of us don't really know about, and one example is random number generators. If they're not as random as they're supposed to be, then the [inaudible] space for cracking something is smaller.

Or we could have vulnerable standards. One example of that is the one on the screen from NIST, which they basically publicly confessed that this is not necessarily safe anymore.

So different types of vulnerabilities and attacks, and you can sort of see how to deal with some of this. And it's kind of obvious, but an important observation is that some of this is maybe standards related, some of this is implementation related and so on.

Vulnerable standards. I mentioned the random number generator example already. There's another one which is pretty important given the widespread use, RC4 algorithm when used in TLS. We don't have as warm and fuzzy feeling about that as we used to perhaps, so that may need to go.

There has also been some discussion about on purpose planted weaknesses in various standards, including IETF standards. I think one case was made about IPsec. Yeah, no one really knows whether those should be taken seriously or not, but in my personal perception having been involved in some of those discussions at the time that some of these standards were made and discussing with other people who had been there even more closely following the situation is that may not be.

I don't personally think that's very likely. I think there are issues with algorithms, and this is not to say that IETF standards don't have security problems. They do and they have plenty of other issues to. It's just that this was a case of whether some maliciously planted direction in the standards has caused us problems.

So anyway, with all this background, what can the engineers do? So first of all, it's important to understand that technology alone will not be the sole solution because it only goes so far. I can protect my communications with someone. I can have perfect communications security but, again, if I don't trust the other party, I'm in trouble.

We can prevent some attacks. We can also increase the cost of attacks or surveillance in terms of money that has to be spent or in terms of there being more risk to getting caught. We can shift attacks from a sort of wholesale attack on everyone just in case to more targeted attacks. We can move from passive to active attacks and so forth.

And the other point is that given that there is so much attention on this matter, this is an opportunity. I mean, it's required for us to look at the situation, obviously, and make whatever we can do about it. But at the same time, it's also an opportunity that now that there is so much attention on Internet security around the world, so if we want to actually want to push some new security [inaudible], this is the time. It's not always easy to do this at any given time, so this is maybe the one chance that we have, I don't know, this decade or something. So it's a kind of important opportunity. We should use that wisely.

So some directions for protection. Obviously if you have unprotected communications, you should consider protecting them. If you have vulnerable standards – you're laughing but most of the Internet is unprotected communications. And the vulnerable standards, what we can do about that, we can do more public review. And this is like what the IETF is about. We have a broad base for reviews and lots of people reviewing documents and proposals. We can decommission old algorithms. We can design new solutions.

We can look at the implementation backdoors. And that's not a standardization issue, by the way, but there are ways of going around that. If you have more diversity in the set of implementations on a particular type of application, then that's good. If you have open source, that's good. We can do other kinds of review, even if it's not open source. That's good. Some of us were involved in a proposal to develop a piece of open source hardware that would be a little better verifiable that you can actually trust better to do important security tasks. And so there are many kinds of things that can be done.

So what is the IETF doing around these kinds of efforts? So first of all, we kind of felt that it's important to discuss the topic. This is not something you just sweep under the rug. We need to discuss it openly. We have. We have the PERPASS list, which is one of the primary lists dedicated for this topic. Russ was running technical plenary at the IETF, which talked about only this topic with Bruce Schneier and Stephen Farrell and Brian Carpenter and talking about the topic.

The IAB is going to organize a workshop before the next IETF meeting in London. The Friday before, yes, Russ? Yeah. On this topic we'll be

issuing a call for papers soon on that. And of course, for me at least, the primary thing is that the actual working groups will do most of the work. The IETF are actively engaged in this topic and trying to understand what can they do and make improvements or do more analysis.

And we of course are working on the problem in a general sense trying to understand the threat model and what classes of solutions exist. We have some specific proposals on what to do with TLS and its algorithms, Perfect Forward Secrecy (PFS) feature in TLS is important to be turned on as an example.

We have ongoing efforts that were started before all this came out that actually have an impact on this matter as well, HTTP 2.0, TLS 1.3. And so the important thing is also that the IETF is a place to bring together the different stakeholders to discuss the different solutions. So when you're doing something, let's say improve the security of the Web, it's not just the browser or the server or both of them. It's also proxies and all kinds of security solutions that we have to worry about. So we kind of feel that this is an opportunity for us to discuss this in a broader setting.

And some of the highest interest efforts, for me at least, have been, first of all, outside the IETF. You can just turn on more security whether it's TLS or some other mechanism. But TLS in particular has been very popular in many types of services and already before everything came out. This trend will not accelerate. So you've seen Facebook and Google to some extent at least moving to secured connections or HTTPS, so that is something that will probably accelerate with other services as well.

The we're going to do algorithm clean-up on TLS and IPsec and many other things. And that's an implementation thing and specification things, so there is both work for the implementers as well as the IETF.

Then we're having a huge discussion about what to do with security for the Web, for HTTP 2.0. The current standing proposal is that HTTP 2.0 is only for HTTPS URLs. This basically means that you want to do more HTTP 2.0, then you're going to move more of your Web presence to the secured links. And it's a huge discussion. It's ongoing. It has been going since mid-last week in particular. If you're interested in that topic, please join the mailing list.

And also, this wasn't such a high profile thing during the IETF meeting. I just wanted to bring it up here given the ICANN context. Some people have been asking about DNS. What should we do about DNS? Should we have confidentiality for DNS queries? That's basically the question.

So that's what I was going to say. I have some further reading and pointers that you might want to look at. The only thing that I want to mention in particular is this recording of the plenary session that Russ was running, the technical plenary with Bruce Schneier and others. I think that's some really, really interesting material. If you watch Internet security related video this year, maybe this should be that video. So that's it. And, Russ, do you have?

RUSS HOUSLEY: A few things. So I would like to point out two statements that the IAB made in recent weeks that are related to all of this. They're both available on www.IAB.org.

The first one is a joint statement that the IAB made with the Internet Society, the W3C, and the IEEE. We pointed out that with these accusations that the standards process might have been tampered with by intelligence agencies that the only defense to that is broad review during the development of the standards. So get involved, review the work, because that's your only defense against that kind of an attack.

The second thing is that we wrote a letter to NIST. They apparently have this one standard that was messed with in some surprising way, and they asked what should be done about it. So the IAB sent them a very specific and detailed list of things they ought to do about it and basically suggested a much more open and transparent process where all comments from all reviewers and how they are handled is put online for every standard they develop in the crypto space and the cyber security space going forward.

Anyway, I think that basically we're saying follow the example that has been used in the Internet community for a long, long time.

JARI ARKKO: Yeah, and I don't know if we'll have any time for questions. The basic message is we're looking at this problem from a technical perspective, and the community is hugely engaged in this topic and are actively working to improve the security of the Internet in general, not just because of this particular revelation. So we would be very happy to get even more people involved, so please join the mailing list and engage in the discussion.

EBERHARD LISSE: Okay. Thank you very much. That was a very nice presentation. It's also the first time that we had sort of the adults talking to us, and I think we should continue this. We all know how IETF works. Everybody can go there. And it's a meritocracy, so if you are clueless, they will figure it out, and if you have input to make, they will figure this out too. And so everybody who has input to make, just join the mailing list. I think most of the work is done on the mailing list, so if you don't go to the meetings. I follow one or two of these thing. It's well over my head, but it's always very interesting to read [this stuff]. Any questions? Oh, they're all overwhelmed. Thank you very much for being here and I hope we will do this in the future again.

JARI ARKKO: Thank you.

EBERHARD LISSE: Okay, Anne-Marie will do her second presentation now. Before she starts, I didn't read the presentation, but I remember – maybe some of you remember in Tech Day in Cairo and in Brussels, we had Patrick Mevzek speaking who wrote Perl DRI library where you can [inaudible] very easily tools to access [EPP] server. That's what I use, for example, to see if our server is alive and things like this. And sometimes I hear complaints EPP is not working; I run [the thing] and it works very nicely. But it's always good to see different things in the wild.

ANNE-MARIE EKLUND LWINDER: Okay, thank you. I feel like part of a Monty Python. I want to do something completely different from the last presentation, and it's not only because I like to hear my own. I'm the only one brave enough to want to talk about this.

So the background is, EPP has a number of years. Created in 2000, draft RFC 2004, 2007 it was ready. And then we implemented it internally while we were putting our new production system into work. And by 2009 when we changed our business model, we forced all the registrars to use EPP only and nothing else – no e-mails, no website, no nothing – only EPP.

And in 2012, we got a contract with ICANN for pre-delegation testing. We got a lot of experience from EPP implementations, and we realized that there was no publicly available test tools to help support the industry to make sure that you have it running smoothly.

So the overview, it uses the same XML input file as the pre-delegation testing system (PDT) does. It tests all different parameters: connectivity (IPv4, IPv6), domain transactions (create, update, delete, transfer), contact information, hosts. And it can actually be expanded to do more good stuff and can handle any extension schema.

So this is how it works. It converts a pdtepp.xml file into an intermediate/internal config file. And the test performs through communication with an EPP server and writes to a log file, and that is where we get the report from the outcome.

It's open source. We realize that people will need to use this and implement it themselves. So t's on GitHub and here is where we have our account and the name of the tool.

It's not a one-to-one equal part that the EPP Selftest Tool and the actual EPP testing under PDT is not equal. So running the EPP Selftest Tool successfully doesn't mean that you will pass the PDT test. So I just have to make sure that you understand it's not the same.

And if you want to learn more of this, there's a readme file on GitHub, and you may also address questions about the tool to my colleague Jan Saell, who is the developer behind it.

And there are a number of other useful tools from us. IDN properties is another useful tool for those who are running TLDs that have IDNs. We have our DNS check, which I frequently use to do a health check of .se every year. And that is a checking tool for checking DNS delegation quality. And we also have a small tool called PacketQ, which is very nice to get deeper into PCAP files to get the details out from it. And there are others who are making very, very good tools like NLnet Labs the Credns and Validns is also another tool that is very useful.

So that would be all, and I hope you find it useful. And if you have any suggestions or questions, don't hesitate to e-mail either me or my colleague Jan Saell. He would be delighted. And questions will be answered on best effort because I don't know anything about EPP.

EBERHARD LISSE: Okay, I have a question. Have you run this already against FRED or against [inaudible]?

ANNE-MARIE EKLUND LWINDER: I don't know. That is one of the questions you could put to Jan maybe. Or do you [inaudible]? No?

EBERHARD LISSE: Okay. That was a short answer. Okay. That's something that I'm at least going to play with because it would be interesting. I'm wondering why wouldn't you try to make this test to sort of conform to PDT testing so that people can use that sort of to get it ready or test if it works then they can go to pre-delegation testing.

ANNE-MARIE EKLUND LWINDER: Yeah, but still the pre-delegation testing is something different from just having a selftest tool of your EPP implementation. And I think that the selftest tool have been further developed than the version that we're using in the environment that we put out for PDT. I think that is the easy explanation to it.

EBERHARD LISSE: Okay. Thank you very much.

ANNE-MARIE EKLUND LWINDER: Thank you.

EBERHARD LISSE: Let me just look who is the next one because I forgot. Okay, the next one is Norm Ritchie. He will speak about the Secure Domain Foundation.

NORM RITCHIE: Hi, I'm Norm Ritchie. Today I'm here on behalf of the Secure Domain Foundation. I wanted to give everyone here an update for those of you that are aware of it and for those of you who are not aware of the SDF. Hopefully, it will be very informative.

So a bit of background on this. The SDF really came from a concept we started back, and it was ICANN in Costa Rica, so some of you that were lucky to be there. Chris Davis, my partner in crime in this, made really a concept.

And his observation – Chris is involved in a lot of cyber security circles – that recidivism is rampant in the domain industry. So the bad guys, they all use domains now. So when they're caught doing something, the domain is taken down, they just merely go get another domain and carry on. So the damage to them is six dollars, which is nothing. And that's very frustrating because it's actually a lot of work to do the research and the analysis on the malicious domains in the first place and build a case only to have them quickly just move somewhere else.

So the concept that was put forward is very simple. If we work together as registries and registrars and hosting companies, share information

about the bad guys, then we could actually prevent them from doing that – a very, very simple concept.

So the feedback for the most part has been pretty positive. Everyone said, "Yes. Great idea. Why don't you guys go and do something about it." So for the last year and a half, Chris and I have continued to work on this. We actually started off with basically a daily feed of malicious domains that we identified. We analyze through emerging threats about anywhere from 100,000 to 200,000 malware samples a day, and from those we extract the domains and we have some indication of what the badness is about the domain, whether it's command and control or some of those [seen in] website buying, credit card numbers or whatever that may be. There are about nine categories.

So that daily feed evolved to an API, and then that API has evolved a bit further. It has actually evolved to more of a reputation system now. So basically from an e-mail address, you can get a score and the reputation of that e-mail address – is it good or bad – allowing a bit more proactive anti-abuse measures to take place. You have some indication before you actually register a domain or set up an account if you're hosting or whatever that has someone else seen that domain, and what badness do they see about it?

The SDF is incorporated. It's a Canadian nonprofit. Yay. Chris and I are both Canadians. Chris is online by the way. If anyone has seen him online, he's there somewhere. And we're going to officially launch this in January of this year coming.

The handsome guy on the left is me, and the other guy is Chris. The reason I wanted to show you this though is that I come from the domain industry. I've been doing this for, God, it seems like forever now. Probably 15 years. Chris is the same. He has been doing cyber security since the early 90s.

And Chris and I happened to be friends. We worked together on a number of initiatives in Canada. But what we found that I know the domain industry but not so much of the cyber security. Chris knows a lot about cyber security and a bit less about the domain industry. But we both realized that the two of them are so tightly intertwined that it's important that we work together.

And through the 18 months, it hasn't just been Chris and I doing everything. We've had some other people that have stepped forward and have been assisting us to develop the SDF and the database, the API. I'll get into that in a second and what that is.

I'd like to recognize them. Facebook has been awesome. They provide us a daily feed of bad domains that they uncover. They're very, very supportive. Internet Identity, the Anti-Phishing Working Group also quite supportive in supplying information. Demand Media who many of you know as eNom has also been a great partner as well. We are actually currently doing a baseline analysis of their entire suite of domains, 18 million of them.

Nominet, I think are always here somewhere, always great in this. Emerging Threats, Trend Micro, Gonzalo who is up later on today, SOCA, and of course CrowdStrike. Chris and I both work for CrowdStrike now.

They're very, very supportive of the SDF, and they support us in doing it. There are others. Bank of America has just expressed interest in this and asking how they can help. And ESET, the antivirus company, they're also stepping forward. And tomorrow we're talking to a large hosting provider.

So what do these people do? Well, they give us bad guy e-mail addresses and that goes into a database. I'll get into the size of the database here in a second. When there's a Terms of Service breached quite often we can, depending on the contractual agreements, we may be able to get additional information about what was behind that account that is not normally available. So that might be what is behind the privacy for a domain registration or what is a browser, the fingerprint that bad person was using.

Malware data and analysis, of course. People may help in investigations, whether that's research or analysis. There may be some type of development that needs to be done. We can't do it all. Or also just on the operational side of running the service.

So the current evolution, we have MD5s, which is how you identify the malware, mapped to the domain name and categorized. There are over 86 million records on this in the database. On the e-mail addresses, we have over 7 million of those. We have a very, very comprehensive set of WHOIS data, over 26 million records there. We think we have, actually, one of the most comprehensive WHOIS databases right now.

We have browser fingerprints. It turns out that browser fingerprints, if you're not aware of that, it's a way of looking at a query in the browser

and asking it to respond on its configuration, such as what [font packs] are installed in what order etc. and what plugins may be there. But it turns out there's a great deal of information sent back, and it's also sent back in a specific order. And that actually gives you a very unique – not totally unique, but 99% uniqueness on that browser fingerprint.

Again, I already talked about the Terms of Service violation logs. For those of you that do cyber security, you might be familiar with Maltego. It's a visualization tool for connecting different things that's used quite a bit by cyber security analysts. We have a transform server for that that ties into the database, and overall we have over 300 million records now.

And what's coming soon, and part of the reason that we want to give you this update at this time, is that we try to work at attribution. So not just, is this domain bad? Who is actually behind that? Who is the real person? Then that goes to law enforcement and take them out of business forever.

Part of that though is knowing something more about the address. Is it a real address, is it fake, etc.? Because in cyber security you're using the same information that a lot of people in the domain industry, and it's WHOIS, etc. The bad guys tend to use the same garbage information over and over again though – very interesting – so you can actually track them that way and correlate them.

But we have been working with Canada Post. I'll get into that in a second. I'm jumping the gun here. We have a postal address validation system that's free to use. We are going to be working on a phone number

validation to follow that and the browser fingerprint, which I talked about, and also registrar rankings. I'll get into each of these. I'm watching the time.

EBERHARD LISSE: Don't worry.

NORM RITCHIE: Okay. Postal address validation. In partnership with Canada Post – as I mentioned, we're both Canadians – it just turns out they've been developing a system for the last eight years on address validation. Obviously, in Canada they have a very good handle on it even to the point where they know that they've actually delivered a piece of mail to that address, so they know it exists because someone has actually gone there and done something.

They also have fuzzy logic that's put in, so like I live on Bush Drive. If I put in mail Bush Road, it would still find it. It wouldn't matter. And I live in this town called Greely which kind of part of Ottawa. If I put in Ottawa, it will still find me. So it has that fuzzy logic built in.

They partnered with somebody else on international address validation. We found that's not up to the stuff that we want it to be. So we're now looking into actually Google Maps API, getting a business license for that so we can do high volume transactions on that. And everyone knows that Google knows everything about everything in the world. So we're in the early stages of adding that into the address validation as well. And this is free. It's free to anybody who is a member.

Oh, I wanted to mention U.S. Oddly, Canada Post does not have a good handle on address validation in the U.S. I forget why. There's some silly reason for it. But UPS, they have a system for validating addresses where they're going to ship in the U.S., and they said we could use it.

So all that was just passed through. If someone wants to validate an address, we'll just pass it through to these partners and you get a score back. It's free. We're looking for beta testers for this, by the way, if anyone wants to speak up. I've already talked to our friends in Prague.

Phone number validation. Right now we just do a simple regular expression check on the phone number. Is it a valid phone number? For instance, if it's in country code one, which is about 14 countries. So I guess Canada, the U.S., and a bunch of the Caribbean islands comprise country code one. But make sure the area codes exist, and that 1.7 million did not exist. They're garbage. That's a very, very simple test. We'll publish the findings on those soon.

We have been talking to some other companies about actually doing some phone number validations, something a bit more hefty. We haven't found anything yet that is cost effective. It's a very costly venture.

JASON SHULTZ: Excuse the interruption. I was just wondering if you do address validation as well.

NORM RITCHIE: Yeah. You mean in connection with the phone? Not combined yet, so yes. So, sorry, the answer is no. But, yes, you are correct. We don't do that.

On browser fingerprints, we have a browser fingerprint algorithm we're using. It is about 94% good. We've turfed that. Sorry. We actually threw them all away. So we started over, and we have a much better one now, and the algorithm actually is reversible. So that let's say that someone updated Java in the browser, you could actually see that that's the only part that changed within the browser, which is kind of a cool algorithm. That code is available if anyone wants to use it.

Registrar ranking, this is something we intend to do. We actually run every day, we look at the breakdown for actually all the registries on bad domains that were identified. Those reports are available to anyone. Just ask for it. You can have it. Well, as long as you are the registry, you can have it. My address is at the front of the presentation. Just send me an e-mail and we'll get that set up.

What we're looking for the ranking though for the registrars is not how many bad domains they have but how responsive they are to abuse complaints because you actually want them to improve. So we have to develop the algorithm for that, but it will actually be based on turnaround times and have they actually responded to an abuse complaint rather than just how many domains are there. Through that what we expect to happen is that the registrars who basically don't care will probably rise to the top of that list.

So a little synopsis of the API. We have validation scoring for postal address, phone, e-mail. We have malicious ranking on e-mail address, IP address, domain, also the address, phone number, and fingerprint. The reply back [can] actually be very simple or like JSON or you can go with XML. You kind of get your choice on what you want.

If you put in a little +details, you'll get all the details. And it looks like that, and you probably can't see that at all on the screen, but I'll give you a sense of what there. This is actually a query on an e-mail address saying that it's malware that was associated with a botnet. That same e-mail has been seen in the forums [stealing] for stolen credit cards. And then it's just a generic abuse as well, and it has more information on what WHOIS data is associated with this. I think there's also a browser fingerprint in there.

So the summary of this is that if you're interested in the address validation, it's free to registries and registrars and hosting companies. Just contact us and away you go. If you're interested in participating on the abuse, either getting data or helping us and provide data, share with your counterparts, come talk to us. And let us know if there's something else that's not here that you need to make this useful for you. And I think that's it.

EBERHARD LISSE: Alright. Thank you very much. Any questions? We are a bit ahead of time, so please feel free to ask. Yes, Jason?

JASON SCHULTZ: I'm just looking at the website and looking at the registrar rankings, and it says coming soon. So maybe can you give us a sense of who are some of the worst registrars that you've found so far, [S domains]? I have experience with a few that I think are pretty shady.

NORM RITCHIE: Well, yeah, I think if people are doing any type of cyber security, you tend to know which ones they are. There are three or four, right? But there's no hard evidence right now against them. I think that if we actually start sending them abuse complaints, they're nonresponsive, then that's going to become more evident.

We'll actually have some type of metric because one of the problems is – we talked to ICANN. Actually, we spent a day at ICANN and talked to all the various departments. And on the compliance side, they need help as well, right? ICANN is very tied, handcuffed with all their contracts and everything. They've got to watch about getting sued because they [inaudible] sitting there, and people want to [inaudible] get their hands around that, but we can help them. Thank you.

EBERHARD LISSE: No questions? Alright. Good. Thank you. There's one from the remote.

UNIDENTIFIED FEMALE: There is one person called [inaudible] who says he would like to get the abuse report, but it's not really a question. And Chris Davis is responding. Okay, sorry.

NORM RITCHIE: Okay, perfect.

UNIDENTIFIED FEMALE: I'm sorry.

NORM RITCHIE: That's great. That's great. That's good.

EBERHARD LISSE: Okay. Alright.

NORM RITCHIE: Thank you.

EBERHARD LISSE: Thank you very much. Let me just see who the next one is. Daniel Kalchev will now speak about the REGRR protocol. Afterwards is Gonzalo Romero about dot-co.

DANIEL KALCHEV: Okay. So this is a presentation for protocol that we have been developing. The reason we decided to develop our own protocol is I would say mainly because we first had the registration procedure, then the registration system, and then finally came to the conclusion that we cannot do it without the protocol.

We looked at what was available as, let's say, various ideas and implementations. We did not find any of those compatible with our system and registration process, so we decided to implement our own protocol. So I will try to share some of the concepts and actually try to keep the presentation not specific to our implementation but to what the protocol can actually solve in the hope that it's maybe useful for others too.

So the protocol is I would say yet another XML protocol. The [key thing] that is probably differentiating this protocol is that all the messages that are exchanged over the protocol are individually [signed]. The messages are supposed to be exchanged over an encrypted channel but this is, of course, not a requirement. And every message that is exchanged from, let's say, the inner message and all the encapsulating messages carries all the authentication data that may be necessary for the registry or whoever is processing the message to authorize various [object creation] or manipulation.

Another interesting aspect of the protocol is that with every message the version of the nomenclature, that is the lists of various possible values, is communicated and if there is a difference between the server and the client, the client can request an update of this list so that they have up-to-date information.

This is one feature that is, I would say, directly related to the verification of the data that is [inaudible] at the registry database because the server is supposed to refuse any data that does not conform to the public nomenclature lists. And also the messages can be nested in many

levels that [inaudible] more than two levels of communication that we commonly see in our implementation.

So what is the goals that we achieve with this protocol? One is that we want to have a secure communication channel with the registrars. So with the registrars, we have contracts. We know who they are, and therefore know how they will present themselves – okay, identify themselves – to the registry.

Another goal is to have secure communication with registrants. This is, I would say, the most common problem all registries have. And the problem mostly boils down to the situation that almost no registry knows who the actual registrant is. They all get this information through some kind of proxy and with not strong enough authentication.

So the idea of this protocol is that if we can identify the registrant, we can have communication with that registrant with the confidence that they are the party that requests the [object additions] or modifications even if they go through a proxy like the registrar.

Of course, the protocol achieves end-to-end encryption and the signing of all messages. In the case of our implementation, we keep record of every message that was ever exchanged with the registry so that in the case of any disputes we can produce those messages.

Another very important issue that the protocol solves is the separate authentication and authorization of the registrars and the registrants. In the most common model, the registrant will authorize the registrar to do anything on their behalf. This is often not the case in real life and is, I would say, probably the key reason why we had to develop a protocol

like this. So in essence, the authorization of the protocol messages will follow the contractual relationships between the different parties.

And here I would like to just say that the registry users are ultimately the registrants, and the registry should have a means to fulfill their contractual obligations to the registrants even if the registrar is not functioning properly.

So in short, how the protocol works, this is more or less related to the current implementation. We expect that the registrant will prepare and sign the message they want delivered to the registry. They could do that either through some interface that the registrar will provide, or they may use some other kind of protocol or the same protocol to send a message to the registrar.

Then the registrar will encapsulate that registrant message in their own message and will communicate that message to the registry. This is the role of a proxy that we often expect from the registrar.

Then when the registry will receive the message, it will of course check the integrity of all the layers by verifying the signatures at every level. And according to the signatures that are used for each message, it could authorize the party to do the different, let's say, parts of the request it intends, like I think new domains or some modifications, object removal and so on.

And, of course, if the registrant does not want to deal with any of this, they could authorize the registrar to submit messages on their behalf. So in this case, we fall back to a scheme that we have now. We will trust the registrar to do the object modification on behalf of the registrant.

And then the last [inaudible] is that since we already know who the registrant is and can authenticate them securely, we can let them communicate directly with the registry so that they can, for example, automate the management of their DNS and DNSSEC data or contact data updates and similar with the registry without the need of the registrar to participate in this case.

So the implementation, we had this protocol completed two years ago. As of now, three of the registrars that we have for dot-bg already use the protocol to communicate with the registry. We have also implemented several proxies that use the protocol and provide a subset of the comments for those registrars that have not yet implemented it.

The current server implementation expects to run over TCP/IP with TLS. This is not a strict requirement because, like I said, the protocol is [stateless]. It can use [inaudible] transport. And only I would say the negative I think we have so far in the implementation is that we are not yet ready with full translation of the documentation in English. So that was my presentation more or less.

EBERHARD LISSE: Okay. Thank you very much. I have one question. I personally as a registry don't want to communicate with my registrants. Why would you want to do that?

DANIEL KALCHEV: Well, if you ask the registry as a business, of course we would not want to do that. We will prefer to do less work and get more pay. But the

reality is that we ultimately have a contract with the registrant not only with the registrars. This may be in some cases for historical reasons, but it is a fact. And we get a lot of questions from registrants to implement, let's say, more direct communication.

Then there is another issue is that we often see a situation when the registry wants to implement new technologies like, for example, DNSSEC is very good example in this situation where most of the registrars simply don't care, and there is no other way to make them care than to let everybody else be able to handle their tasks. So this will actually, and as we see from the examples in real life, let's say, motivate the registrars to move faster.

EBERHARD LISSE: Okay. The policy of our registry is totally opposed. We do not communicate with individual applicants unless there is gross misconduct of the registrar. We have a help desk and we have the most favorite reply is, "We don't talk to individual registrants. Please, go to your registrar." I'm just saying, different models. The less you have to talk to thousands of individual clients, the less help desk stuff you have to employ and things like this. For a small registry, we have no employees, so I have to do the help desk. I don't really want to talk to three and a half thousand applicants.

DANIEL KALCHEV: So this is why we implemented the protocol. We don't really want to involve people into this.

EBERHARD LISSE: Sure. But that implies that registrants understand how protocol is written. Our people don't really know how e-mail works. My Internet is not working. Okay. It doesn't matter what. It's different philosophy. We feel, and I think many registries feel, as little communication with the individual applicants so that we can cut down on the help desk. CZ.NIC has got a help desk of a few people. They deal with individual applicants. You have got, we as a very small operation, we just don't want to.

If we wanted to or if we had to, of course that's one way of doing this. Yeah. So I'm not trying to criticize the implementation. I'm just trying to sort of ask from behind why would one want to do that.

Warren was first, and then you will be next. We are good for time, and we have to have a few minutes' break for technical reasons because we must sort of get the translators who have in the meantime arrived so that we can get them onto the system.

WARREN KUMARI: Warren Kumari, Google. Quick question. Wouldn't it have been simpler to have just used an extension to EPP or something like that?

DANIEL KALCHEV: An extension to EPP would mean that we had replace the, let's say, normal EPP with [inaudible]. And it will again create a completely different protocol. The idea we had – and the need actually – we had

this to be able to authenticate the center of the message, not the postman who delivered the message to us.

EBERHARD LISSE: Oh, yes, Jason. I forgot.

JASON SHULTZ: I just want to clarify one thing. So this protocol is meant to be between the registry and the registrars?

DANIEL KALCHEV: Yes.

JASON SHULTZ: Only?

DANIEL KALCHEV: Not only.

JASON SHULTZ: And the individual domain registrants as well?

DANIEL KALCHEV: Yes.

JASON SHULTZ: So recently there have been a string of attacks, for example, by the Syrian Electronic Army where they basically steal someone's credentials and they go in and they alter name servers and things like that.

DANIEL SHULTZ: Exactly.

JASON SHULTZ: How does this – someone could just as easily steal a certificate, yes?

DANIEL KALCHEV: Yes and no. I mean, okay, in this case this is yet another point here. I actually left it out of the presentation on purpose because I was sure that – I actually had a slide that was labeled "Why Not EPP?" But the point here is that if they steal the certificate of the registrar, what has happened in most cases nothing will happen in our protocol. We will get an outer message. The stamp of the, let's say, postal service forged. But we will still either have a valid internal message signed by the registrant or we will have something that we don't recognize and we simply don't care.

JASON SHULTZ: Discard it. Right.

DANIEL KALCHEV: And then the best part of it is we will have good evidence of what we actually received because it will all have digital signatures on it. So that's [inaudible].

JASON SHULTZ: Thank you.

DANIEL KALCHEV: And just one more point. Of course, if they wanted to, for example, forge the [size] of Google, like what happened in many of those cases, they should have gotten the certificate that Google used to send the messages. And if they can do that, they can hack their sites too so that's [inaudible].

EBERHARD LISSE: We've got Google here. Let's talk to them.

WARREN KUMARI: So I deal with a number of the DNS hijacks that we get, and in the large majority of the cases it's not actually that the credentials have been stolen or guessed or anything like that. In what seems to be the majority of the cases – often it's hard to get actual details – but it seems to be that it's actually the registry itself that has been owned through either SQL injection attacks, which are unfortunately still really common, or just Web parameter changes.

So, you know, the [ERL] will be whatever either the registrar or the registry/user=7 domain=something. And then somebody will just change the domain parameter to be Google.whatever. That is a really old attack, but it still works on a lot of places. Oh, God, yes. Probably upwards of 75% are a combination of SQL and the other.

EBERHARD LISSE: Anything else? Alright. Thank you very much.

DANIEL KALCHEV: Thank you.

EBERHARD LISSE: Next presentation will be dot-co, and then I would think we would take a five to ten minute break after the presentation. We are a little bit ahead, not much though, so that we can connect the translators to the outside [thing].

GONZALO ROMERO: Hi, everybody. I am Gonzalo Romero. I am the chief security officer of the .CO Internet. We are currently the registry of the Colombian ccTLD. I'm very glad and proud to be here for showing you our strategy in the ccTLD in regard to the security issues.

I have these items in the agenda. It's very, very short, so I would like to show you the motivation we have, the security policies we are currently handling, our knowledge transfer and cooperation action efforts, the

malicious activities monitoring process we have in place for three years, our current challenges in cyber security, and at the end you can ask me whatever you need.

I think you know that we are a private company, and we have a concession contract with the government of Colombia. And we work on the promotion and engagement of the community in terms of a global perspective, not only the in-country perspective. So we are working more on the spread of the community worldwide, and I think we have a very nice responsibility with our registrants worldwide.

We are working on this because we feel that we are responsible in the short term of these kinds of things. And as administrator or the ccTLD, we are very committed to protect the integrity, the stability, and the reliability of the IT and the services of the TLD and as well as the image and reputation worldwide.

We think that security is an added value for the registrants worldwide, and we participate and contribute with efforts, activities, and initiatives to maintain the SSR of the Internet global ecosystem.

This slide shows you the security policies we have in place. One of the first – one thing is a very successful point is our price. We are not cheap, and it's a deterrent of malicious or abusive or illegal domain name registrations. And that's very good for us in terms of we are not a five dollar per year domain.

We have very good practices IT, security, and business continuity. We have a very, very strong relationship with NeuStar, our partner in these kind of things. We have a very interesting and innovative Malicious

Activities Monitoring (MAM) process that has a security position and a reputation of the ccTLD globally.

We have a very, very strong relationship with our registrars in terms of conditions. We are working on the process of suspend domains or notifying registrants and registrars in terms of malicious activity. And we have a lot of work in terms of global, regional, and in-country initiatives for working together in terms of security.

All the time we are working on knowledge transfer and security awareness. And we are working on joint efforts for awareness campaigns and initiatives to work with private and public entities not only inside our country but regionally and globally as well.

In terms of knowledge transfer and cooperation action, I think that's the key of our success in terms of security. We are [all the time] building community in terms of security as well. We have a CO-DNS community in which you can participate as well. We have a DNS Technology and Security Day every year in August, and we have a very interesting active participation and commitment with the national, regional, and global community.

We support every initiative, for example, to certify national CERTs in the FIRST Organization. We currently have now five certified CERTs in our country, so that's very good initiative and process. We are very close to have our national CERT certified by the first, so that's good news.

And we have a lot of agreements internally with law enforcement agencies, with [inaudible] and the IXP. And we have a very strong

relationship as well with the academic community among other agreements with them.

And globally, we participate. We are members of the Anti-Phishing Working Group, and we have very interesting agreements with Microsoft and NCMEC. We participate on DNS-OARC. We have a very strong relationship with RSA and other very high-level companies in security. So that's a very good point for strengthening our security in the TLD.

This slide shows you the Malicious Activities Monitoring we have. It's a nonintrusive process of just monitoring the zone looking for malicious URLs or domain names. The feeds are provided for different several trusted third parties. We validate phishing, pharming, malware distribution, malicious hacking, defacement, and child porn cases.

The NeuStar CERT is our investigation body. They take the potential alerts and notify to the registrars and eventually to their registrants as well and work with registrars and registrants given a timeline from six to 24 hours to resolve or at least to put the domain in [separate] hold or to obstruct from the DNS zone.

And all of our work in terms of these kinds of security issues is based on cooperative action and terms and conditions. We have a very interesting and strong agreement with the registrars, and they have a [transitive] agreement with their registrants so that we ask for them to notify the registrant at the end to alert or to solve the problem we are notifying.

In terms of special cases like rogue pharma, content, piracy, or spam, we forward that cases to the Colombian law enforcement agencies – we have a cooperation agreement – for research and actions.

Our challenges, I think the most important and relevant we have now is the cybersquatting. It's a very difficult and challenging process, and obviously any infringement or violation to our terms and conditions. We have a very proactive monitoring of hourly dot-co domain name registrations looking for trademark and intellectual property violations, and that's very important.

And this is the letter we sent to the suspicious registrants doing some kind of cybersquatting or [inaudible] squatting in our domain. We proactively are doing this kind of efforts just to not allowing these kind of issues in the domain. And that's it. Thank you very much.

EBERHARD LISSE: Thank you. Any questions? How many domain names did you have?

GONZALO ROMERO: We have 1,600,000 domains currently.

EBERHARD LISSE: And what did you say your price was?

GONZALO ROMERO: $30 per domain per year.

EBERHARD LISSE: How much?

GONZALO ROMERO: $30.

EBERHARD LISSE: $30?

GONZALO ROMERO: Yes, that's domain.

EBERHARD LISSE: That's very cheap. Maybe we have a premium domain model. We have even less foreign registrations than you have. They just can't afford it. It's just not worth the effort. Any other questions? I think we'll break for five to ten minutes. We don't have to go. It's just they need to do some technical stuff so that the translators come on. Can you indicate when you're ready for us? Okay, thank you.

Okay. Can everybody please settle down and sit down? Everybody who doesn't speak fluid Spanish should get themselves a headset please. There's headsets lying around here and in the front. I'd like to do this relatively quickly so we can carry on, please. There's sets laying here. There's also a number of sets laying here. We're having the English translation streamed. I don't really know how this works if there is questions from the floor. As long as they're in English, it's not a

problem. If they're in Spanish, we must just listen. So that should all work.

Gabriel Brenta is the technical director from dot-ar and he is going to give us what we have come to refer to as the usual Host Presentation. We have traditionally invited our hosts to sort of give us an introduction on how they run things and if they have any special interests on research and something to talk about that as well.

GABRIEL BRENTA: Thank you very much. Sorry for not delivering this presentation in English, but I think that you would appreciate that. The organizers asked me to let you know how we're managing our ccTLDs. And to tell you the truth, this is a very interesting case to share because some time ago we have been in charge of the operation and it has certain funny characteristics.

My name is Gabriel Brenta. I am the CEO of Computer Systems, also technical and legal secretariat of the [Argentine Presidency]. This is an agency managing the ccTLD NIC.AR.

Going back in history, many people – many of you know this – in 1987, a governmental organization Chancellery in charge of all international relationships started to play with the Internet and started to move forward to take Internet seriously in Argentina. In 1998, the first structure of ccTLD managers was established. At that time, an application structure was created allowing people to transact between nic.ar and to register domain names such as dot-com, dot-ar, dot-edu, dot-net and we started to turn into a ccTLD as it is today.

As it usually happens in the computer [inaudible], there were some ups and downs. There were many personal efforts [applied] and technology from the [inaudible] and we had an application that allowed us to work until August 19th this year. In 2001, [audio cuts out].

Let me tell you about the structure. Nowadays, we have 2,300,000 domain names and this is quite a lot for the region. Then I will let you know how we compare ourselves with the rest of the countries. The structure and openness of domain names is, as you can see in this picture – I don't know if you can see it – we only have 444,000 domain names and then the rest is divided.

When it comes to dot-tur, that has to do with tourist agencies and we are trying to identify the specific purposes that [inaudible] the rest of the world and we are part of that.

This is the relationship between domain names and population. Argentina and Columbia, we're approximately on the same basis in terms of population. We have a difference of 6.7 [inaudible] at 2.8 and Brazil goes below. And they have a greater amount of domain names. But in terms of population, we are high above.

Part of the complexity about this process has to do with the migration process from a ccTLD working and that was migrated to a different operating structure without the capacity of having [inaudible] migration, not only migrating functionally [inaudible]. This was carried out by part of a working group. It was an application reengineering. We had to adapt the infrastructure and we have to work with the latest

[inaudible] in software development and we started to understand what we could add in the terms provided.

So we performed the technological [startup] and we started to create the networking layers. We have the IPv4 and IPv6 and the DNS, and all you can see as a [startup] of a ccTLD taking into account the structure that is now operating should be kept and we have to build on a structure according to the new infrastructure that should support [inaudible] being implemented. New security dynamics were implemented inside and outside the application and we applied the migration of the data center.

Migration was conceived or implied [inaudible] the old applications, data migration taking into account that one of the things or the paradigms being changed in the application had to do with leaving or putting aside a certain type of registration and have new registration form, which is by using a username and password. So we had that data migration trying to have a friendly [inaudible] of information. We had to have the IP masking for nic.ar to start working with this new infrastructure. We had to begin with the infrastructure. We have to have the approval of the user. We generated the first zone and we have the master DNS, all that during the same weekend. This was something very interesting.

We have 20,000 daily users accessing – and that might be escalated, according to certain contingencies – taking into account the new infrastructure of the application. That is not a problem. Access was tripled regarding the previous – or in comparison with the previous platform. We are now publishing two [inaudible] per day, and in the

short-term, we want to advance with the publication of eight automatic posts.

The DNS zones are uploaded to the masters in IANA or defined in IANA, and within Argentina we have three and two outside the country, and one of them is Anycast network. We performed a migration and update of BIND to incorporate new features that will be implemented very soon in order to improve safety – [inaudible] safety and query safety. And we are now working to diversify the DNS operation to reduce the impact on the DNS.

As future projects, we would like to see our internal network, national and regional network, be mounted on Anycast so as to have a different level of collaboration so that the availability of the service is the one that we have at an international level.

I would like to tell you that we are about to implement IPv6 – [native] IPv6 – but my teamwork works very quickly, so that is not a project anymore. It is being operating, so you can have a full IPv6 access.

And the other expectation – strong expectation – that we have, taking into account our recent access to the ccTLD world is to empower networking among all the group having or sharing the same problems that we have.

Any person understanding or knowing that we can help or anything we can share or can be shared with us, please contact us because we are willing to help you. We want to improve the things we are doing and we want to implement things that perhaps you might have already implemented. So we don't want to have any problems or obstacles to

solve. So thank you very much for your time. If you have any questions, please let me know.

EBERHARD LISSE: Okay, thank you very much. One question in the back from Robert who speaks Spanish.

[ROBERT]: Hello. This is Robert from [PCH]. I would like to know which kind of problems you had getting IPv6 from your providers and how many of your providers are able to deliver or claim to be able to deliver.

GABRIEL BRENTA: In Argentina, all the providers are full compliant with IPv6 in the PPT. So having the operational startup or achieving the operational startup status was a little bit harder. In fact, we had certain problems with non-compatible ASNs. The ASNs, they were not compatible with the core. And thirdly, two of them were incompatible with the core. In the third case, we didn't even try it. We insisted and insisted and insisted and we paved the way for IOS credits and today we are happily using four or five IPv6 applications in a demo version.

This is a big problem in Argentina. I understand that everybody's doing or making a great effort in order to solve this issue. However, the low traffic and the low amount of requests from customers means that providers do not focus on the technical update that is required. However, this is operational to date. Thank you, Roberto, for your question.

[UNIDENTIFIED SPEAKER]: Hello, [inaudible]. I would like to know what is the trigger for dot-ar to have more than two million domains. Is it related to marketing or to technology development or anything else?

GABRIEL BRENTA: I would like to say that people really love us, but in fact, what helps quite a lot is that they are free of charge.

PATRICIO POBLETE: Patricio Poblete from NIC Chile. The change that you implemented was quite a big one. You changed from one paradigm to another one. You had a paradigm in which registrants identified themselves in a certain way to a new paradigm in which registrants use a username and password. Can you share any learning experiences? Is there anything you would have done differently? Because we are about to start a similar process in NIC Chile.

GABRIEL BRENTA: Initially, the shift or the paradigm shift, implies making an effort and going along the right way. In as much as the base that you want to migrate is as clean as possible, your day after will be happier or as happy as possible.

In 2.3 million domains, well, as you can imagine, we have several users that have registered things that were not sustainable through our time. So while they were not able to prove or demonstrate that the identity they used when they registered was still valid and updated, so we had

to study this on a case-by-case basis in order to give legal certainty to the registrant, and in order to start fighting cybersquatting.

So you have to implement certain massive actions, but on the other hand, analyze things on a case-by-case basis. Let us get together later on to analyze this with more granularity and focus on things that we're 100% positive and others that were positive, but painful at the same time. Is there any further question regarding to this point? So thank you very much for your attention then.

EBERHARD LISSE: Thank you very much. I'm quite thankful that we managed to get translation services organized so quickly and I want to take this opportunity to thank the translators for doing a very good job so far.

Next is Warren Kumari from Google. We know him. Most of us do. And he will talk about Automating Delegation Maintenance.

WARREN KUMARI: Hi, everyone. I am going to be talking about two drafts that we have in the IETF. Please let me know if I'm going to quickly or you can't hear me or anything like that.

So this is about automating the maintenance of delegation information. It actually covers both DNSSEC and other delegation information. I am going to be focusing on the DNSSEC use case because it's the one I'm most familiar with, and also I find it the most interesting.

First off, what's the actual problem here? Well, rolling DNSSEC keys is hard. Actually, it's not really the rolling of DNSSEC keys that's hard. it's the publishing of new DNSSEC keys that's hard. Currently what happens is I roll my keys and then I need to find my registrar's log-in information. I've usually forgotten that, so I need to go through password recovery. They send me another password, then I log in.

Then I need to find the domain I want to manage, then I need to go click on "manage DNS information", then I need to click on "DNSSEC information", then I need to click on "bulk upload of DNS records." And then I cut-and-paste my records, and for some reason, they want it in a different format, so I swear a little bit. I add some spaces, I remove some spaces, I click "submit". Eventually it works.

Then they want to know do I want to add a new DNS record or do I want to replace the current resource record [set]? I never know which one to pick, so I sort of choose at random, and then half the time my domains go bogus. This obviously does not make for a good user experience.

Publishing your new records is also somewhat dangerous. As I said, it's fairly easy to screw this up. There's also a problem if you outsource your DNS operations to a third-party. [inaudible] the ones who do the DNS key roll, but they don't really have a way to upload the new key.

There are two solutions for this. One of them involves the DNS operator e-mailing you the new record or DNS key and saying, "Can you please upload this at 10pm on Thursday?" If you don't see this e-mail or you forget, your [zone] goes bogus.

A solution I've heard proposed by some DNS operators is, well, you just told me your registrar's log-in information. Basically, you give me your [inaudible] credentials and I'll log in as you and then I can just manage this. This obviously is fairly scary from a security standpoint.

So, how do we deal with this? Well, the solution that we're proposing in these drafts are actually in CDS draft – is simply publish your new DNSSEC key information in your [zone]. Basically, take your new DNS record or new DNS key record, you stick a "C" in front of it and you publish it in your zone. Technically, it's a new resource record type, but it's easiest to just think of it as stick a "C" in front of the type.

So if your new DNS record looks like that, it becomes this. As you can see, it's identical other than the word [DS] has changed to CDS. The "C" stands for Child.

If your parent prefers DNS keys instead of DS records, exactly the same thing. You publish the new DNS key record and you stick a "C" in front of it. Now it's in your zone.

So, how does this actually end up being published? Well, what will happen is your parental agent will come along every now and then and will poll the zone and will look for any CDS or CDNSKEY records. If it finds them, it validates them. And if they validate correctly, then it could only have been you who published them. Then it actually sticks them in the parent.

So, hang on a second. Who's this parental agent person? So we have some new terminology here. In the RRR model – the registry, registrar, registrant – the child is the registrant or the DNS operator. The registry

is the parent, just like you would expect. And then the registrar is who we're calling the parental agent. We specifically chose these terms and not registrant, registry, registrar because we want this to be a generic solution that can be used in multiple environments.

So what ends up happening is the parental agent will poll the child. It does a standard DNS lookup for CDS, or CDNSKEY if it prefers that, and the child is running a standard name server software – BIND, NSD, whatever it likes. And it simply replies. It's a standard DNS lookup. It's a standard DNS record, standard DNS server. The parent gets the CDS record. It strips of the letter "C", make sure it validates, and then it pushes it up into the parent or the registry using whatever it currently does. This is likely to be EPP. Basically, whatever protocol it currently uses. And now the parent has the updated DNSSEC information and you can start using the new key.

In a number of environments, the parent is the same as the parental agent. This is very common in things like enterprises or educational institutions, universities, things like that. Basically, exactly the same thing happens. All we've done is we've moved the labels around a bit. The parental agent will query the child, [inaudible] example.com. The dot-com name servers or a tool on the example.com name servers, querysales.example.com. The child replies. The parent publishes it [inaudible] using whatever protocol it would currently use.

In the RRR model, a fairly common setup is that the DNS operator is the same as the parental agent. Basically, [inaudible] places, the registrar will run DNS services for registrants. In these cases, the organization can continue to use what it uses now or it might actually choose to use CDS

because we're expecting that this will become a standard and it provides a keen separation of [sort of] duties. And then it gets published in the parent using whatever protocol the registrar currently uses – EPP probably, whatever.

So a few additional details. Polling –really? Yeah. Currently, what we have specified in the draft is a simple polling mechanism. The parental agent will come along every interval – whatever it chooses – and will query all of its children for CDS. Or, if it prefers, CDNSKEY and will publish them.

We fully expect that people will extend this and write additional triggers, because the system self-authenticates itself. You could have a restful interface or an interfaced based on updates or notifiers or the phase of the moon. Whatever you figure you'd like, you can use that as a trigger to sort of initiate this process.

CSYNC. There's another draft which complements the CDS and CDNSKEY called CSYNC and this stands for Child Sync and was written largely by Wes Hardaker of Parsons or SPARTA or whatever Russ's company is called this week. And it does similar things, but it is mainly for name servers [inaudible]. Basically, the CSYNC record is a type bit map of which resource record should be copied and the parental agent will poll for this resource record type. It will look to see which messages should be copied into the parent, and it will simply copy those verbatim into the parent.

So this sort of sounds like two very similar drafts for accomplishing very similar tasks. Yes, they are very similar. CDS and CDNSKEY is designed

and optimized specifically for DNSSEC stuff. It allows for some things prepublication of keys. So, keys that you want to publish the DS record for, but not actually use yet. And similar things.

CSYNC is designed specifically for name servers, glue, potentially also as a trigger for CDS, possibly other future records. Because it simply says what records to copy, it doesn't really work hugely well for prepublishing of keys. And I think that that's it. Questions? Roy?

ROY ADAMS: Just not so quickly. I can't run that fast. I am properly handicapped. Thank you very much, by the way.

WARREN KUMARI: No worries. Hopefully, that was at least vaguely understandable to people.

ROY ADAMS: Thank you for the presentation. My name is Roy Adams. I work for Nominet. This is basically an overview of what's currently going on in the IETF, and the last part I've seen was basically presented in Vancouver two weeks ago. However, I've also seen an alternative – a [cert] alternative, if you will. [inaudible] and they complement each other, basically. But a third one that I've seen is the Mark Andrews update. Could you elaborate a little bit on the difference between CSYNC, CDS, CDNSKEY and updates?

WARREN KUMARI: So the CDS draft was actually adopted by the [DSSEC] Working Group like three days ago I think. We actually misspelled the draft name, unfortunately. We couldn't remember how to spell maintenance. But I actually have Mark Andrews' thing here as additional slides.

The Mark Andrews solution uses standard DNS updates to do this, to achieve similar things. I personally thing that it's a fair bit more complex. It involves, first, the registrar and registrant or DNS operator have to exchange TSIG keys. If you've ever done stuff with TSIG keys, they're a little tricky to get right, unfortunately.

Then the child has to query the TLD. We'll pretend this is a registry/registrar/registrant model. It has to query the TLD for child.update.tcp.parent. So, for example.com, it would be example.com.update.tcp.com. And this means that the TLD has to have a sort of parallel DNS tree. Either that [inaudible] delegated TCP or actually stick the entries directly in itself so that the zone file gets way bigger.

It then replies with the serve record telling the child where it should send DNS update messages to. Then the child has to do a DNS update message to the registrar, authenticated using TSIG. This of course means that the registrar needs a TSIG capable name server. The child needs to be able to speak TSIG. They need to remember which particular TSIG key or credentials they're using. And then the registrar sends this to the registry with the EPP or whatever, which then publishes it in the name server.

So this all technically works fine and relies upon existing DNS stuff. It just requires I think a lot more work on the registrant side than they're comfortable doing. Registrants have a hard time with this. It also requires the TLD either doubling its size or having a parallel tree, and the incentives seem odd to me.

Somewhere we actually have a very brief proof of concept, sort of semi pseudo-code implementation of the CDS thing and this is largely it. As you can see, it's 10 or 15 lines of [python] which we think is way simpler than this. Anyone else?

EBERHARD LISSE: No further questions. Thank you very much, Warren, as usual. Next one is David Peall from [inaudible] DNS Africa or whatever it's called this week.

DAVID PEALL: Domain Name Services.

EBERHARD LISSE: Okay, that's what it's called this week. Where's your thing? There it is. He's going to speak about – I gave a short [inaudible] on the agenda. He's going to speak about domain name services, right protection mechanism that [it] probably implemented in the new registration system, if I'm not mistaken. Go ahead then.

DAVID PEALL: Thank you. Domain Name Services is a technical provider for the ZA Central Registry, which is a gTLD applicant for a number of domains – dot-africa, dot-capetown, dot-durban, and [inaudible]. We also operate the technical backing for the ZA Central Registry, which is the ccTLD ZA.

When we were looking at our proposal to ICANN for the gTLDs, we conceived a rights protection mechanism and we wanted to group the protection across all the domains, including our ccTLD. So this is basically what we're going to cover.

We looked at the TMCH and it's a great tool for big corporates who have trademarks registered with trademark authorizes and are looking for global coverage. But it does have some limitations and it covers pretty much the tip of the iceberg when it comes to rights protection mechanisms.

Firstly, it's a direct word match on your trademark. So [inaudible] registered with your trademark authority, that's exactly what you're going to get obviously with compatible string conversions in terms of what's a legal DNS.

And what we've noticed is that there was a lack of awareness and outreach from the TMCH in Africa. Obviously these things changed and we've had a discussion with the TMCH operators and they are looking at collaborating with us in terms of marketing in Africa, so that is changing which is great.

But what we saw was that Africa is full of small micro-enterprises that aren't big corporates that don't technically deal with things like trademarks and international law and that kind of stuff. So we saw a big

gap between what the TMCH was offering and what our customers were going to want. So we were proposing that during our sunrise we wanted to enable local businesses, companies, trusts, common law trademarks to participate in our sunrise process. And we did this through conceiving this Mark verification system.

It also does give protection to these guys in our ccTLD. We are launching three extra second level domains. [The org] is being relaunched and net and [inaudible] are also being launched. [inaudible] has been running and is currently sitting on about 900,000 domains.

We didn't want to reinvent the wheel when it came to the Mark verification system, so we modeled this on TMCH's technical spec that was in the IETF. So basically a Mark holder will register there [inaudible] system. We will validate it and then issue them an SMD, much like you would with the TMCH and they can then approach any registrar with the SMD and do a domain registration. This obviously avoids the registrars doing any major development, which is great.

One of the additional things we wanted to add to the sunrise was priorities so the SMDs will have, depending on what kind of Mark you register with our Mark verification system, you'll either fit into one of the priorities.

The first priority is obviously a full trademark registered with a trademark authority, and those will get priority one if they are registered in an African country or district. And then registered trademarks outside of an African region will be rank two. Then the common law business names and registered trademarks will be priority

three if they are in an African region. And priority four would be the unregistered trademarks. Anything greater than four will be [inaudible] premium names and generic registrations. If there are two registrations on the same priority, we will resolve it through an auction.

Being a relatively new company – this is our first foray into the IETF process and we had two very different experiences. The first is to make our envious work in the TMCH model. We wanted a small change done to the EPP launch phase [inaudible] was proposed. The [inaudible] EPP launch phase is an extension mapping for provisioning of management of domains during a launch of a registry or a name space.

So what we wanted to do was add a validator ID so that the registry could tell which Mark verification entity the SMD came from. We submitted a change. We discussed it on the mailing list. We engaged with authors and we submitted a change and the source for the IETF draft is in GitHub. We were able to create a fork, do an update and a poll request and our changes were published in version 12 of the draft. That was quite a smooth process and we were quite encouraged by that.

The second change that we were trying to get done is on our TMCHs in [inaudible]. Sorry. I've made a mistake here, sorry. The validator IDs is for the application – the EPP domain application. What we wanted to do with the reciprocal was the SMD to send to the registrant was to include a validator ID as well. And the [inaudible] draft that's in IETF.

We weren't able to solicit a response on the mailing lists or from the author in terms of getting an update done to this. I think that's what the

power of these meetings is all about. I've been able to talk to people who are actually actively helping us find and communicate with authors of these drafts. So, a very positive thing is ICANN and IETF meetings that enable us to communicate with the participants.

Additional features that we added to the Mark verification system – we call it [inaudible], which is much like the LORDN process where TMCH will notify Mark holders of registrations of [their] Mark. But what we saw was a gap between the registration and the update of your registrant details. So if a Mark's registrant is altered, we thought it was valuable for the Mark holder to know that their registrant details had changed. Effectively, changing the registrant on a domain is changing the owner, so we thought that that was important for the Mark holder.

Something else we've worked on for the MVS out of bands claims notices. A lot of registrars are going to battle, especially in the African market with additional development of things like claims notice, which is a process where the registrars have to spare additional information during a sunrise and [inaudible] period to a registrant trying to register a domain if there is a claim in one of the clearing houses. And they have to accept this before they can continue with the registration of the domain.

So what we've devised is a process where the registrar, if they send through the domain creation with doing the claims check, we will hold the domain registration until we've done [inaudible] out of band through e-mail to the registrant to confirm that they are aware of the claims before proceeding with the creation of the domain.

And any moderated demands that need a central moderation location, but would like any registrars to participate in the registration of domains, the MVS is the perfect fit for that. Thank you.

EBERHARD LISSE: Thank you very much. Any questions? How many people are using this already?

DAVID PEALL: We've involved the top IP firms in South Africa and currently they are entering Marks of their clients into the system to participate in the launches of [our] domains.

EBERHARD LISSE: Okay. Any other questions? Alright, thank you very much. So that gives Jay his usual spot to tell us something very interesting.

JAY DALEY: Hello. I'm Jay Daley from the dot-nz registry and this isn't the very interesting slot. This is the windup and finish off slot today. First of all, for those of you who don't now, this is organized by the ccNSO Technical Working Group. Can I ask all members of that to stand up? There are three of us here. There are one or two others who aren't here. So we would like to thank all the speaker for coming along and giving their time.

Today certainly the presentations have been very professional, increasingly professional, and very much sharing best practice. And I think that this reflects the maturing of our industry as people work through many things and learn and share with each other.

We are already planning the Singapore meeting. So if you have presentations, please let us know. We've already [agreed] two of them today. So please, if you've got these ideas, come forward to us. Today, for one of the first times, we've actually turned down a number of presentations because, for one reason or another, either they weren't technical enough or they didn't meet with our particular goals today. So please make sure you get them in nice and early because that almost guarantees you the slot.

We've discussed in the working group some plans for the future we want to share with you. One which we've mentioned a number of times before is that we want to bring in a wider audience here. Could I ask, if you are not from a ccTLD, could you please raise your hand so we can see how many non-ccTLD people there are? So that's great. So we have seven in the room. We would hope to have about 50 at some point.

We know that if you're a technical person coming to an ICANN meeting, the number of opportunities that you have to interact on a technical basis are relatively limited, and this is an important one so we want to broaden the scope of this to other people.

We would like to see, if possible, all of the technical meetings in ICANN tagged in some way just with the word "technical" and some effort put in to make sure they don't overlap because there's the security sessions

one has [inaudible] gone on this afternoon as well, which is a poor conflict of time – so that we can maximize things.

And I think that many of us do dual roles. I'm the chief executive of a registry, so I do the technical stuff at times. And other times, I'm doing the finances for weeks on end without ever turning the light off or going to sleep. And I think there are a number of people around here who do a very similar job. So we want to try to ensure that we make this more useful in that way.

We're also talking about bringing in one or two external speakers if we can – potentially some high profile speakers who will be able to share things with us.

The other thing is the things that we talk about generally have a policy implication and we are looking at how we can ask people to provide presentations on the things they have learned here or the things they have talked about here to the ccNSO, removing the technology details and just talking about the policy implications. I think it is important that we who often are the first to encounter a new issue because we encounter technology can be the people who then relate that to the rest of the ccNSO by explaining the policy implications of it. So we'll be asking more and more speakers possibly to do two presentations or consider two presentations – one technology and one shorter policy implications one that they can deliver in a different way to a different audience.

So that's it from us. Did you want to add everything, Eberhard?

EBERHARD LISSE: Yes, I wanted to mention two things. The peak number of attendance in the morning was 90. In the afternoon, it was 85. Norm Ritchie didn't raise his hand when he was supposed to be counted. He is a member of the group and another member just entered the room, so we are five and not three.

In any case, I appreciate all of you who remained until the bitter end and I'm leaving you now to enjoy whatever party you're going to tonight.

[END OF TRANSCRIPT]