PRAGUE – Tech Day
Monday, June 25, 2012 – 11:00 to 17:00
ICANN - Prague, Czech Republic

Eberhard Lisse: Can we settle down please? Good morning. So I'm going to hide behind this pillar that both sides can see me nicely. For the ones who don't know me, I'm Eberhard Lisse; I'm the Chair of the Technical Working Group and I have omitted with forethought the word ccNSO in this regard. Welcome again to our usual Tech Day.

Today I think we have a number of really cool presentations. The first one will be Dmitry Kohmanyuk from the Ukraine who will show us how you set up DNSSEC in six months in a reasonably-sized registry.

Then we'll have a really cool presentation – a short one; one that is a little bit shorter than expected because it was put together on very short notice. Richard Lamb will talk about TPM which I call The Poor Man's HMS but as far as a trusted something platform management chip which can be twisted into amusing hardware signing things.

Then Thorsten Kraft from eco or DCIX will talk a little bit about deep packet inspections – what you can do with this, what you can do or what you shouldn't do and what you should be afraid of.

And then we'll have lunch. Lunch has been sponsored by the Dutch Registry. We hope that we'll get somebody to say a few words or at least we usually offer them a few words. It will be a box lunch but only

people who are here at 12:50 will get something, so don't run away too early.

In the afternoon we will have the usual host presentation but even better this time because Ondrej is going to do it. And then we'll have a round table for DNS servers as we all know also what the different name servers are. There's been nine and coming to be 10 and NSD and then there is the new two kids on the block differ from EURid and Knot DNS from the Czech NIC. So I think this is going to be quite interesting to compare notes and as Olaf Kolkman used to say, we're all among friends here. This is not going to be a slug fest that we try to pick each other apart, but it's good to compare; it's good to see what things can do and why one should try and take a different one.

And then as usual this time it's Norm Richie who has been twisted into making the closing speech. Alright, without further ado – Dmitry. Just one more thing. We have a remote presentation which means if a question comes from remote, I think they will get precedence and we have got two room microphones for discussion afterwards. And you can remain seated if it's easier for you.

Dmitry Kohmanyuk: Thanks everyone. This is probably more interesting to ccTLD operators, especially small ones. I don't have to say why DNSSEC is important but if I can move this forward… Okay, so the regional plan was about a year ago and I was deliberating it for a very long time.

The kickoff point was November 8, 2011 during one of our IPv6 Technology Workshops in Ukraine. We hold those about quarterly.

While it's kind of a revolutionary day in Soviet history – if somebody knows what I mean – we started to deploy in a sub-zone called UA.UA just a test zone. Some interesting experiences were gained.

The key was generated on November 8 as I said and this zone was signed. The UA zone had a secret delegation. I also set up a very small website which had just a page saying it works and an anchor, so I think it's very important to start doing this with something that you don't mind breaking and experiment all the way because you would not get a chance doing this later on the production zone.

So as I said, it's six months – it was actually a little less than that. DNSSEC-trigger was instrumental to verify how things work. All the various options were tried including N63 and 63 was opt out. Also had a wild card record in that zone – something that you don't normally do in production. But it was very interesting to see various interactions especially with unbound and not handling N63 with wild card in some corner cases.

So about a month later we had the official UA Key Generation ceremony. It was scripted and rehearsed. The event was dedicated to the 19th Anniversary of UA domain operation. It was recorded and published on the website. We had about 10 people participating as observers on stage; about 60 people attended the event. Having a small webpage – a mini-blog in our case – was very helpful. I used the parameters we have used – nothing so special here. Also want to add that we ended up using NSEC 3 with opt-out.

So in parallel I was setting up the test environment, so essentially a full replica of the current reproduction but using separate hardware. I list

some details here – BIND 9.8 – the then current release. Now of course we would use 9.9 – we'll probably migrate to this at some point; loss of internal scripting – we had an hourly chron job generating the zone and signing it so it was completely parallel set up, and as it says several sleepless nights – lots of fun. Again it's very good to test things which you don't mind breaking because they were breaking a few times.

So we also had a separate Anycast server with complete replica of the zone – all the records the same; separate anchor and you are able to use it if you just use [Type 4 end] zone and maybe just local inbound. Again DNSSEC trigger is probably the way to use it. Just set it up on your own laptop and try. Lots of people use this; we use this in our company internally.

So the next step was to offer the public resolver which was using that test zone as a liaison. The project was code named Lighthouse. It was announced in our workshop on February 7, so it's what – three months in. It was using that authoritative server I have mentioned so essentially if you use this as your resolver that was easier than doing any magic. It just works and you are able to install Firefox or Chrome plug-in developed in Prague and it was another tool I would never be as successful while there was another testing point.

Next I… I should have mentioned this. I was in a meeting in Costa Rica and Steve Crock and I had a little chat during the dinner. He was asking me how things are working and I was saying, "I'm about to do this but I'm not sure exactly when." And then he just asked me to draw a little plan on a napkin and that was actually the point of no return.

So I just went ahead with everything and we were live on Friday, April 13. Since then we had six secure delegations, two of registrars and two geographical domains, three domains. Every one of those I'll walk us through. Again, it helps to have those test beds set up because half of those already present in the test zone. So we knew, then we go live, it will be working already. So there was no this, what about you turn this on and somebody else breaks.

So that's the traffic. It's been slowly going up. I think it's about – I put this number there – it's about 0.2% of the overall DNS traffic, so the query rate now is about I would say 4 average six max for second. And again, this is only one of our servers, so we have six servers in UA, so that's the one we run. Not so much but with the number of delegations we have, it's not surprising. I guess I'll stop now and that's the status page we had and any questions from the audience please. Thank you.

Eberhard Lisse: I didn't get – how are you signing - hardware?

Dmitry Kohmanyuk: No, it's all software – just BIND 9.8; script only. Nothing fancy really. I would not have the budget for that nor the time to debug. I should also add the zone is quite small. We had about 15,000 delegations in it. So the performance another problem. But the whole signing process takes like five seconds.

Alex Blowers: Hi, Dmitry; Alex from Nominet. I have a different question independent of the DNSSEC. You mentioned you started with a test zone – UA.UA. Now in Holland we did something similar about 10, 12 years ago where we had an NL.NL. What we saw almost immediately or basically what we noticed, what people told us almost immediately is that some resolvers, some OSs have configured NL as a search box and that broke things horribly.

Dmitry Kohmanyuk: I know what you mean. I chose setup UA.UA about 10 years ago exactly because of your experiences. And we have been running this and I was monitoring query rate. This I call it the Double TLD Animale. I should write probably write a small report. We do have query logs. Thankfully the amount of those inquiry requests is steadily falling down, so when I was trying this I was pretty confident I'm not going to break anything.

My UA.UA zone was not a real zone. It was a zone with a [wild card] record we set. This is UA.UA; this is DNSSEC. So I probably broke it for very small people who had a) those [query]-resolved behavior and b) DNSSEC enabled. I think the intersection of those sets is empty. But I of course can't attest to that. But thank you for this point, yeah, it's very interesting. As I said, this is only just a test zone. And we had these reserved since 2002.

Adam Peake: I've been asked by Dr. [Eberhard] to make this point again. So…

Dmitry Kohmanyuk: That's why we had it reserved – because we decided to never delegate this domain.

Adam Peake: I understand. I'm just trying to elaborate on it a little bit. So there is something called the search pass in your [FC.resolve.gov] and so the search pass could have something like Nominet@org.uk or it could have something like .NL. So when you type in the domain name, it might be suffixed with .NL. So sometimes you have something.nl.nl and of course that wouldn't resolve properly.

Eberhard Lisse: Okay, any other questions? Come on, don't let him off so easy. Alright, thank you very much. Okay next one will be Richard Lamb. Okay, for the remote participants Richard was kind enough to just put this presentation together on extremely short notice. I met him yesterday evening in the pub about it because on that FRED Masterclass we had the day before with the cz.nic where we met [Jaomir], the programmer and his team. And Luis from .cr mentioned that they're using a TPM chip and he mentioned Richard's name about it. So I think it's actually quite a cool thing to look at it.

Richard Lamb: Alright, thank you very much. I'd rather stay standing so I don't fall asleep. So I apologize for the roughness of this presentation here as it was a last minute thing. But it is something I'm very excited about and with Luis in Costa Rica – he's the one that kind of started the germ of the idea going as far as using a TPM chip.

The Poor Man's HSN – and thanks to Eberhard. I didn't realize The Poor Man spelled out TPM so this is perfect. I'm slow. English is not my first language. I went to public school system. [laughs]

So one of the things in deploying DNSSEC is everyone goes, "Do you use soft keys, hardware keys or what have you. Do we use smart cards? Do you spend $20,000 for a high-end cryptographic device or what?" For most of us we don't have $20,000 burning a hole in our pockets, but we still would like security. We still would like a system that actually insures that the private half of keys that we use for DNSSEC is protected in some way.

Well, it turns out that Dell computers, Dell servers, I think Gateway servers, IBM servers – a whole bunch of machines out there – commodity hardware – has this cryptographic chip built in it and no one really uses it. It was originally put in place there for the purposes of being able to attest to the integrity of the hardware on a PC and then also attest to check the digital signature on various pieces of software and stuff.

And as some of you may remember the TPM wars, maybe about five, six years ago. Everyone was really angry; this is an evil plot by Wintel – hopefully there's no one in here from Intel or Microsoft – Microsoft and Intel to somehow lock down DRM – Digital Rights Management – lock down software and stuff like that.

Anyway, for whatever reason, maybe some of that politics and some of the other things, the TPM chip kind of is not used. But in any case, Luis identified that there's this TPM chip there. There's been software out there trying to use the TPM chip for various other things for a while and

at least every time I've tried it I've failed. Software is not quite up to snuff.

But now there's added incentive and enough motivation actually to go that last mile and figure out to make this Open Source software that actually works with the TPM chip work. And so that's what I want to talk about really quickly here.

So PKCS11 is an interface to HSMs. It's a standardized API for HSMs. It doesn't matter whether it's a Smart Card or whether it's a $20,000 HSM or whatever it is. And so this is a good thing. If we can write to this, it's a good thing. I think some of the recent versions of BIND have support for it. They worked through the open SSL engine in order to be able to make this stuff work.

There are various ways to do this. I modified DNS signs a long, long time ago to just directly talk PKCS11. Who needs open SSL – so we went straight to this. And so from my point of view it's kind of easy. But it's good to have the standard. The PKCS11 interface has a million different things – delete object; draft object; ways to export keys/import keys; back up things – all kinds of stuff.

But all we really need is two functions. In DNSSEC all we ever really do is generate a pair of keys – RSA keys, what have you and we sign with those keys. So if you actually get into this stuff, you'll find that the PKCS11 is not that… vendors have slightly different variations of this everywhere but all we really need is these two things to work. That's it, okay?

As I said, there's various Open Source software. I know these are very sparse slides; I was just throwing this stuff together; there's not much here. The Smart Card that we all have – and I actually go around doing a training course for ICANN where I say for some of these smaller TLDs all you really need is a DVD with some software on it and a Smart Card reader, a flash drive and a Smart Card to do this. It only does one signature per second but that's actually if you're just looking at the KSK for example – not the KSK and the ZS – just the KSK, how often do you change that? How often do you really use that? Not that often.

So one signature per second actually works. And these things are certified to all kinds of levels. You can get a pretty highly secure Smart Card actually. Again, the idea here is don't let the private key ever out. That's it – you're really just protecting the private key and these things are great at doing this. In fact, most Smart Cards you can put something on here but you can never take it out. Once you've written the private key in here, there's no way to extract it.

The Open Source software that supports Smart Cards is Open SE. It's been around for a while. It's actually improved vastly in the last year. Before that you just couldn't count on it – sometimes it would work; sometimes it wouldn't work. Now it works on Macs; it works on all the Linux flavors – it's actually pretty good; supports a wide range of cards.

So this is pretty cool. But it's one signature per second and when you try to go buy a Smart Card – if I'm getting too long, let me know – if you try to buy a Smart Card, these things are… they're not sold in quantity one, okay? These are things that are sold to large Defense Departments or agencies in quantity – 20,000 or a million – so it's kind of hard to get

these things. And then sometimes you run into some export regulations which is kind of silly.

So let's look at this Trusted Platform Module chip, this chip that's actually already inside the PCs. It turns out there's an Open Source package that supports the TPM. It's called Open Cryptokey and Trousers and if you have Ubuntu or something like that, you can just… I think you can just go app.get and install this and it actually kind of works. It's still somewhat disappointingly slow, but still, this is something that is in devices and you could use right away and it's built into a lot of servers and actually allows you – because of the way this works – it allows you to actually have any number of keys actually. It's not just one key; you can actually use the TPM key to manage many, many keys because actually the TPM chip – it's about this small – it only holds the master key then encrypts everything else.

So you generate keys on it; it has a random number generator on it which is very, very critical. One of the biggest problems that people have found out there with SSL - and there have been recent studies on this – is that the source of the key – the random number generators – are bad - they're using a bad source of random numbers. And so if you use nothing else about the TPM except for the random number generator, you're ahead because the random number generators on these sort of things are very well thought out and certified and people have spent their lives trying to do this right. So that's pretty good.

Alright, so this is where the talk is going to go a little strange here. Maybe I want to… I'm going to go through this slide deck and I'm going to describe what Luis and I did for .CR.NIC after that.

So the one signature per second for the TPM chip that I'm getting there was a little bit disturbing to me and I said, "This is kind of ridiculous. That's kind of slow. Why can't I get something faster?' In fact, if you look at the data sheet for the TPM chip, this is a standard TPM chip that you will find in many machines and so you can actually see on here that it'll do a 1024 bit RSA signature in 40 milliseconds. I think that's about 25 or something signatures per second. Not one – 25. So this thing actually should be able to do pretty good.

It'll do looks like 200 milliseconds for a 2048 bit. So this made me wonder – and you can pull the hook out at any time if I'm going on too long – so it made me wonder and kind of excited about what we could do here. So I looked at the chip set – 25 second generates TPMs inside has a random number generated – this is everything I need. Private keys are always encrypted; infinite number of keys. And there's actually a mechanism – one of the things that you gotta look at when you decide what you're going to use to generate your keys if it's a hardware device is - how do I back up the keys.

So here we are doing everything we can to protect the private half of the key and if it's inside a device that I can never take out, never get anything out of, how do I back up the key? Something can fail. I mean if it fails, I'm dead. Well, it turns out this well-defined TPM and the standards that go along with it – they have something that allows you to do this so I put this up here just in case you do get the slides at some point and this stuff really does interest you like it interests me. You Google that and you will probably then pull in all the documents you need to be able to make this happen.

So anyway, I said to myself, "Gee, I want to get 25 signatures per second. Why don't I just build it? I don't care. This is not as ICANN; let's just solve the problem." Here is a missing… there are HSMs that work at one signature per second; there are HSMs that work at 1,000 signatures per second at $20,000. Today I can do anything I want.

So I laid out the board; I built a surface mount solder oven – this is just a toaster over with a PID controller which you just stick in series with the solid state relay and this is my kitchen. [laughs] And this works really cool. This is… well, I won't say that word but if you're a geek, you really get excited. And so I put the board in there and surface mount – this is tiny, tiny work, right – surface mount.

So I put the parts on there and lo and behold – you know, that's what comes out. In fact, here it is. So I'll pass it around if I get it back. So I better get this back because otherwise I'll kill you. This is not ICANN; this is my stuff, you know? [laughs] Anyway just pass it around.

So here's the TPM chip - $2.50; here's something that talks – USB – to the rest of the world. USB is the interface everyone wants. This is what I keep hearing from people. "We don't have to want to use Ethernet; we just want a USB thing. We should look like a [fog]." So that's all this is a generic 8-bit processor and here's some flash – 8 megabytes of flash just in case I want to put the keys on here.

One of the problems with the HSMs out there, even the expensive ones, is they keep all that HSM material, key material, on your hard drive. So now you need to make sure that the files are synchronized with the HSM. And even at the root – this is something that's very hard to do when we do this because we need to make sure those files are exactly

the same. Because if they're not in sync, you got nothing; you got garbage. So put it there.

That's a crystal – 8 megahertz crystal – and I threw this on there too because this is a really high precision real-time clock. It's not a CZM clock; it's not rubidium – it's not atomic but it's really high precision. Some of the things I had heard from some of the experts in this DNSSEC community was, "There are other things we'd want. We want a perfect sense of time; one that's not coming from someplace else; not something we'd count on, something that might even limit the signature of a validity period for our sakes."

If you look carefully at DNSSEC, there are some… an HSM will sign anything. So if someone gets a hold of… you know, compromises your system and they get into your system and they say, "Okay, I want you to give me an RS sig valid for this bogus information for the next 20 years," HSM will do it; they don't care.

So anyway, so I built this thing and I wrote all the code. It's very hard to probably see this from back there but this is a code that sits embedded on the chip and here – gosh, this is going to be hard to see – but I'll try to describe some of this. This is just running DNSSEC sign zone. So I run DNSSEC sign zone with the hyphen V flag to make it just be really voluminous.

So you'll see it's signing with DNS key up here. Then here it's signing I think an NS record; an MS record; a DNS key – it's actually retained because I used another KSK to do that. Anyway, this is just proof that this little thingy that any one of you could build – nothing should stop

you cause there's no barriers here – actually works with DNSSEC sign zone and creates a valid sign zone.

So anyway, that's the fun part of this talk. If I may, I will now try to describe a little bit of either verbally or with something on here what we did for cr.nic.

I am re-purposing slides here; I apologize. Apologize for that right off the bat. These are some of the pictures I show all the time. This is how some people store stuff which is perfectly valid. You don't need a fancy setup to do secure deployment.

Okay, so once you have this TPM system in place, here's one approach. This was not the cr.nic slide that I was actually looking for. Let me see if I can look for that. I apologize. Things never work when they're supposed to work, do they? Okay, well I think I'm just going to talk through this then.

Once you get the TPM system working, there are a couple things you could do. You could use the TPM chip to both generate your KSKs and your ZSKs and for .cr we decided to… Luis decided in the sense of coming up with the most secure system to actually secure both the ZSKs and the KSKs this way and you can do that. So even if the signatures per second is not the fastest thing in the world, you can actually use this to secure both your ZSKs and KSKs.

And the way that you would then do the key management in that case is that you would generate – so that you're never transferring private keys in the clear, you would actually generate the ZSKs on your signer, your regular signer that you had signed your zone a couple times a day and

then all you'd need to transfer across is just the public half of those ZSKs on a flash drive, sneaker net to your secure facility and at your secure facility, then you could actually just transfer the public part of the ZSKs, then you could use the ZSK which is generated and maintained with the TPM on a off-line laptop and sign everything there, generate your signed DNS key RR sets – this is exactly what we do at the root and this is what a lot of people do – you pre-generate these DNS key RR sets and that way you never have to put the KSK online.

And you get these set of DNS keys that you've now created – signed with the KSK; put them on a flash drive; take them back to your signer and just append them to your zone file that you want to sign with the DSK and away you go. That was the approach used there; it's been working very well as far as I could tell and it's a perfectly valid approach. I've been approached by other people that have said, "We wanted to use the TPM stuff as well," and I think at some point we would like to be able to publish all the details of this and I think I'm going to look with Luis but at some point I think we will publish all that so that everyone can learn from that experience and try to improve upon it because I built this board and all that, but if the TPM chip is in fact on most of these motherboards, we should be able to use this more effectively.

We should be able to get that full performance out of it – that's the goal. I'd love to be able to get 25 signatures per second and also be able to take advantage of some of the import/export migration tools used to be able to back up the keys between the TPMs. And in fact – just as an aside – the way that works, there's nothing magical. So every TPM chip generates a master key inside of it - it's called the SRK.

And so if you have two TPM chips or two machines with TPMs on it, they both have their own unique public/private keys in them that no one sees; never leaves the thing, okay? But you can read the public key. So what you do is you actually take the public key of the TPM chip that you want to transfer the data to and you actually submit that *via* a special TPM function call to the TPM that has the keys and you say to that, "Please allow me to use this key to migrate."

And so you do that; it spits back a specially encoded little blob and now you can use that blog to transfer anything, so this is something we might want to do in the future, right, in just the future – transfer any encrypted blob from that set into the other. So you're encrypting with the public key of which the other TPM only knows the private key – is the only thing that knows the private key so you're safe.

So once you've encrypted that with the public key of this other TPM, only that other TPM can read that data and only internally, only inside the TPM, so security is maintained. So pretty common sense tools that are in place; it all kind of makes sense cause it's hardware – it just kind of works.

So that's all I have to say. Sorry for being a little disjointed but it is possible to use some hardware that's actually out there already. And if you're really aggressive about it, build your own. With that, I think I'm done. Any questions?

Eberhard Lisse: Any questions? Thank you very much; that was very cool.

Male: So, Rick, do you have any extra boards left over, just a PC board?

Eberhard Lisse: No, the answer is do you have a microwave oven.

Richard Lamb: No, that wasn't a microwave oven; that would blow everything up. They come in box of three and I solder them myself. I'm looking at trying to figure out a way to do this in a more mass production way so that I could at least generate something to hand out to some good friends.

The parts are pretty inexpensive. The TPM is a couple dollars; the 8-bit MPU is like $10, so the whole board. And then the board is 3 for $100 or something like that. But that's really small prototype quantity. So I'd be interested in making a few more of these things because I think it's fertile ground to people experiment with. This is all just GCC AVR. This is all Open Source tools; it's C programming and unlike the pure software world, things actually do what they say they're supposed to do.

Dmitry: It's very nice talk; it's very good experiment, but let us say really your approach is poor man, but I want to say also for a rich man where it's also special clusters produced by Intel specifically, specifically for computer clusters and for cloud that has specific environment to do this kind of remote secure environment, that you can do certified and encrypt and store all information securely with a remote environment.

So if we want to run it really very powerful, you have already solution, but truly not for poor man. So just embedded on there.

Richard Lamb: Yeah, it's all embedded in that; that's the whole point.

Jay Daley: Hi, Rick. Thank you for that. Now I'm no hardware expert by any means, but it's clear that we have a gap somewhere between your $100 board and somewhere between the ridiculous amount of money we pay for things like [SEA] 6000s. So how much work does it take for you to take your board and put a Nitrox 3 processor from Cavium onto it so that we can then get 200,000 sigs per second off it?

Richard Lamb: Not that much. I probably don't even need that processor to get it to be up to that speed. A lot of it is a question of certification so I don't know if you guys are familiar with the certifications that follow HSMs, but pretty much people follow the U.S. set of certifications – the FIPS 140 standards just because there's a lack of anything else, not because the U.S. has any particular say here.

So the reason I went this way, Jay, was because the TPM chip itself already has a certain level of certification associated with it and in my heart of hearts I'm dreaming that I might actually go for one of those certifications for this thing. Well, okay, Warren laughs, but anything is possible, okay? There really are no barriers here in my mind.

So I think it's about $15,000 to get the basic certification so the whole point was that I could maybe do this with this. I think the problem with going with something else much faster – it's not a software problem; it's not a hardware problem – it's now I'm going to have to really certify the whole… I mean go through a much longer process in getting the thing certified and I think more of a six-digit cost to get something certified.

That is not impossible either. I mean I think there are actually plenty of organizations, in the U.S. that would be happy to fund this sort of thing, this sort of effort just to improve DNSSEC and security and general good. So I'd be interested in doing that. But that's the reason; that's why I went this route; that's why I went the slow route. I didn't go for that because I thought what's the point of that? I'm never going to be able to get that out of my pocket.

Eberhard Lisse: Rick, did you notice HanSeng from Korea is in the room?

Richard Lamb: Oh he is! Okay.

Eberhard Lisse: They can make it smaller.

Jay Daley: I would encourage you to go with that because I think that's the one thing that is missing is somebody who is willing to take a lead on doing this and I'd be very interested in contributing to that and I imagine

many others would be. I don't know how much any of us actually desperately need the certification if we're part of the process by which it's created.

What we need is reliable cheap hardware that doesn't have anything to do with Oracle or any other idiot manufacturer involved in it but we can trust it does the things we need it to do. And I think many of us are willing to put money in to do that. We just need somebody who's got the talent to lead that please. [laughs]

Eberhard Lisse: I'm wondering is this not something that you can put on a little laptop, netbook, inexpensive; run Ubuntu to get it all sorted and use that thing – chain it to the wall, put it in a safe – why does one really have to build a…

Richard Lamb: You don't. See, that's the… it's a fuzzy wall. At some point it is just a cheap laptop that you leave offline and have the proper procedures and practices to protect. That's just as good and that's going to do well. But I think what Jay was saying – he needs something online that's fast. As soon as you take that laptop and put it online, forget it. It's no longer something that… It may be secure but it's harder to prove…

DNSSEC has a certain amount of theater associated with it, okay, admittedly and so there needs to be some way to provide comfort for not the engineers in this room. I know I'm preaching to the choir with you – you know this, right? It's not just comfort to the people in this room but comfort to people that don't really know this stuff – the

financial community. And this is why those weird little standards kind of matter to me. I think in order to get the credibility from that wider audience that we need to…

Maybe someone has a better answer. I'd love the answer to be no to that, but I've got that sense that well, for ICANN we need to get things certified and we need to get things audited through a SysTrust audit and by PWC, alright? And it's very clear – it's got to be an HSM, FIPS Level 4, Level 3 – something like... they are very clear rules where I have no option.

That being said, of course, I agree with you. I mean from a security perspective a laptop that's been carefully brought into the system and checked out, generating the keys is just as good, if not better than some stupid little 8-bit processor.

Eberhard Lisse: I'm looking at it from an entry perspective. Dmitry just mentioned that's how they did it. We do it exactly the same way. We also do the soft signing. But for a small little zone like that, to start this one of the big barriers to using hardware is the price. If you take a little netbook that costs about $200 and you make a nice little welded cage in your office that has a very big lock on it so that… that thing has two locks on it – two locks that only two people at the same time can access it. Things like this that should really be for an entry level operation 15,000 names – that should do it for the beginning.

Richard Lamb: Yeah, well it's all about satisfying your relying party, people that are using your system. If you can get their trust, you're done. So two locks that you don't have the key to both for I think is perfectly reasonable. Yes?

Male: I'd like to congratulate you. That was one of the most interesting and entertaining technical presentations I've been to at Tech Day. More impressive since you seem to have done it between pub and breakfast. What I was going to say was just a little comment in passing when we were picking up what you said about theater. What you're actually saying is it works better if you put it in a steel box that's an inch thick.

Richard Lamb: I mean, yeah.

Male: It's part of life really, isn't it?

Richard Lamb: That's partially true. If you actually look… I mean, I think you know this as well. If you look at the standard, the standards and certification, there is some logic behind them. It is not just theater. And it does provide… I mean, we live in a legal environment; it does provide credibility that engineers would not otherwise build easily garner. I don't know if I'm getting that message out correctly.

We could have the smartest guys in the room and this is a different community. [Financially] this is a different set of people. And in order for that trust to jump across that barrier we need unfortunately steel boxes.

Male: Just a quick comment. Thanks, Rick, for an excellent presentation. I mean, I'll buy one of those if it will be available. Can you just… Forget the certification – can you just kick-starter the thing and make a set of solvent? People would pay the kick-starter. I think in California you know what to do. You'll be reaching next year, you know? Seriously, and don't worry about certification. Somebody else would do the certification for you.

Richard Lamb: Let me first see how I can get a larger number of these things produced and be able to get some experimentation going because you're right. At some point you go to all this work and you say, "What have I just built here? I've just built another PC and why don't I just use a PC?"

My purpose in doing this was to try to grasp at that higher performance that those chips could actually provide. And for some reason or not, because of the way the operating system is laid out or the way the driver is laid out – I don't know what it is, but there's something stopping you from getting that full performance.

And those TPM chips continue to improve. Every year they improve; there's another one that's faster and faster. At some point we might hit the 100 signature per second kind of barrier and then it becomes useful.

Eberhard Lisse: My dad used to be a radio amateur which means I'm quite sure there are a lot of people in the Ukraine who will solder this together for you if you get the plans. I think once you figure this out, that's… but you will publish the plan. I'm sure. This is the…

Richard Lamb: I'll publish the circuit diagram.

Eberhard Lisse: This is stuff basically Open Source altogether so you can probably…

Richard Lamb: There's nothing special here.

Eberhard Lisse: So then you get one of your guys, "Here is the stuff. Put it in a microwave oven and it will work."

Richard Lamb: Not the microwave. Toaster oven.

Eberhard Lisse: I personally…toaster oven. I actually prefer the idea of the little netbook because it's one thing; you can put a name on it; you can put… you can lock it up; you can even paint it in a way that you know that's the one; it must be there; nobody must touch it.

Richard Lamb: I agree with you.

Eberhard Lisse: This [donor] gets lost and things. The big thing about the laptop – it has got software on it. I just need to put a batch [light]; I just need to sync and run the program and it's done. As long as you make sure that this thing remains offline, we could even disable Ethernet; we could even physically disable the Ethernet then after we could even physically disable the radio and in the [blue] so that it is actually not possible to connect this thing.

Richard Lamb: We do that at the root. We have a laptop that is a special laptop that we pay a lot extra for that has no wi-fi; has no hard drive and we yank the battery out and throw it away because batteries can go bad if they're sitting in a closed environment for a long time. So yeah, no argument.

Eberhard Lisse: A EEE PC costs $200. Simple. Okay, alright. Any other questions?

Richard Lamb: And I would like that board back at some point.

Eberhard Lisse: Rick, so would everybody in this room. I think this was very cool as Nigel said. One of the most fascinating presentations I've heard in a long while and particularly since he was so kind of flipping it together on short notice. In the meantime, Thorsten Kraft will set up. Thorsten Kraft works at D6. I happened to fly together with his boss in the same plane to Costa Rica and he mentioned that Michael Rotert is the Chair of ECO, the association of the internet industry in Germany and he mentioned on the plane that they also do an internet exchange and they do some packet inspection to prevent botnets and so on.

And in my country they wanted to try and pass a law that not only forces the telecoms to listen to us; they want them to presentations, our telephones, our communications with the doctors and the lawyers, but they also want them to save this data preemptively and retrospectively make it available and things like this and I got worried.

So I asked whether if they couldn't come around and tell us what they're doing, what one can do, what one should do and what's possible and what we should be afraid of. I'm not sure whether this is exactly the gist of your presentation, but I hope you can delve into these issues, at least at the end of it.

Thorsten Kraft: Hello everybody. My name is Thorsten Kraft. I am from eco, the association of the German internet industry and one of our subsidiaries is DCIX, the German exchange point in Frankfurt and what you see here at the walls is the daily traffic we manage in Frankfurt so in a peak time we have I think around two terabit of data packages in a second – quite

a lot to do a lot of bad stuff like deep packet inspection or something like that to prevent botnet.

But what we are doing in Germany is a little bit different. I want to hold a presentation on the German Anti-Botnet Advisory Center to give you a short overview what we have done and set up here in this place and I want to give you the possibility to listen very close to it. Afterwards I want to have a question into the room – what you have missed in there and then I want to start a discussion with you.

Okay, The German Anti-Botnet Advisory Center – why has it been set up here in Germany – because millions of computers worldwide are infected with a piece of malware unnoticed by their computer owners and Germany was in the top 10 rank and the German government wanted to change this because they are not happy about this position in a worldwide ranking.

So the goal of the Anti-Botnet Advisory Center was to support the end customers in internet security, to reduce the botnet behaviors, to free the infected machines from malicious software and to get Germany out of the top 10 ranking.

So what we are doing is helping the customers by removing the malicious software and we are working together with the internet service providers and anti-virus vendors. The target group we have to find where the Microsoft computers – because Microsoft is the major botnet's router and so we are focused on that.

ISPs and banks detect by themselves. The infected machines are notify their customers and point them to a centralized call center, including a

ticket system without personal data so we only have a ticket ID to not have to possible IT to refer back to the end customer.

We got initial funding from the German government of 2 million Euro – that's a lot of money – and we had to deal with the BSI, the German Ministry of Internet Security and all the ISPs have to take the necessary technical and organizational steps to implement the initiative to inform customers and to detect infected machines.

The website has been set up in three steps. We have an inform section; a clean section and a prevent section to inform about what we are doing exactly because we have to be very, very open in how have we detected a malicious or the infected machine.

[break in audio]

Thorsten Kraft: This is a tool to disinfect the machines and prevents how a customer is able to work or to not get infected five minutes later after we have cleaned it up.

So in the prevent section we say check your computers on a regular basis, install service packs and security updates and run anti-virus scanners on a regular basis. These are the ISPs that have participated – Deutsche Telekom; Vodafone; 1 & 1; Web.de and some others and we have some banks involved there.

What you now see is the number of downloads we have received in the time of the Anti-Botnet Advisory has run more than 1.4 million

downloads and activation of the tool and the deletion has been started. The web access – more than 2½ million visitors; 4½ visitors on a block with additional information on step-by-step advisories and we have a forum with more than 30,000 support requests.

What you now see are two different stuff, two different anti-virus problems. On the one hand side bka-trojaner – that's ransom ware that is known as a police Trojan and a DNS changer. You have heard in the news and therefore we have installed different landing pages – two different landing pages to point the customers on that.

There you see the systems we have scanned for instances in March. What you see is that a lot of people have [moderate] infections – more than 8.5 infected files on the system on an average.

What we would like now to set up is ACDC – that's the Advanced Cyber Defense Center with 28 partners in 14 member states and how it works is the detection has been in this solution as well.

So on the one hand side we are trying to deal with stolen credentials, with [swipe by] exploits and malware on hosted machines with spam campaigns and with DDoS traffic detection. We point all those data to a centralized clearinghouse trying to analyze the data and pulling out the data through the relevant stakeholders.

The hosting companies the access provided a mobile network and to let the customer do cleaning up his PC over the national landing pages. So what you see here is the collecting part. We are using honeypot systems and drone reports, various sensors and user reports pointing to the centralized clearinghouse, working with different entities on the

same data feed, pulling out the data for the stakeholders like police ISPs, universities and security vendors and we would like to have value added services like on the protection side on the detection side on the advisory side and the mitigation stuff.

And this is what we have done in Germany. Here again the ISP is notifying its customer and pointing him to the national platform. So this is how we have done in Germany in the Anti-Botnet Advisory Center. What you might have mentioned is that there is no deep packet inspection in place. Why is that? Do you have any ideas why we are not using deep packet inspection?

We have [Wisdek], a very big data source. We can find out infected machines – that's possible. The ISPs can do it by themselves but why don't we do that?

Male: Privacy?

Thorsten Kraft: Privacy. It's forbidden. In Germany it's forbidden to take a look into the data packages because I'm able to look in each package; I can look into each communication between one and the other. For instance, if credit card credentials are transferred from my PC to a shop and it is not SSS encrypted and I see private data, the credit card credential's to pass phrases. Any other idea why we are not dealing with deep packet inspection to fight botnet?

[background conversation]

Male: Are the botnets using encryption to counter deep packet inspection?

Thorsten Kraft: Any other ideas? It costs a lot of resources to look into each packet. How many CPUs do I need to look into each package and to decide what is evil and what is not evil? How many traffic do I have to throw away?

Male: Maybe because you don't want to from a moral perspective?

Thorsten Kraft: From a law perspective?

Male: No, from a moral perspective.

Thorsten Kraft: Moral.

Male: You just don't want to.

Thorsten Kraft: Additionally yes, that's possible or that's a very, very hard point. The community is not wanting it. If we are as a German ISP talking about

the packet inspection, there's a very, very big cry in the community. What? You are looking into my packages? I cancelled the contracts. So I need to have what I have mentioned again, contract. I need to have the permission of my customer to do so and this is such a hard stuff I will place there that I have to put it very in red letter, very, very big written on there – "Be careful. I am looking for fighting botnet into each package in blah-blah-blah."

I need to be transparent in the contract. Any other ideas why I don't use deep packet inspection? Or make it vice versa – any reason to use deep packet inspection in a public network? We are not talking about hidden [dark net] or something like that, a traffic that is going into a dark net I'm looking into the traffic. Normally there is traffic that shouldn't be there, that can be handled – that may be done. Is there any reason why I should do deep packet inspection in a public network?

Male: Maybe because you also need to analyze the traffic to understand the communication protocol that the botnet is using. So if you really want to fight it, you also need to understand how the nodes in the botnet communicate and you need to create some sort of sinkhole or something.

Thorsten Kraft: If I sinkhole it – that's the main difference. I'm not doing it on DCIX, on a very big network and every traffic [spreading] out. I'm doing it locally on a box [I] controller. There's traffic coming to this specific box and I'm looking into traffic that I receive there. That's a little bit different to a

DCIX or traffic on a peering point. No reason for deep packet inspection?

So why are we talking about deep packet inspection? The German government I have shown it on the slide here, the DNS changer who asks to use similar technique to redirect malicious DNS traffic to controlled DNS servers to not have the internet black out on a given date. Does this make sense to you?

On the one hand side the German government is wanting to rise attention to infected machines, to infected customers and on the other side they took care or the ISPs took care with the redirection of DNS traffic. Wouldn't it be wiser to let them go down, to not see the internet anymore and to rise the attention with a broken internet?

Peter: I'm Peter [Koftinek]. I'm a bit confused as usual. Could you elaborate on what the German government asked you to do and on what legal grounds they did?

Thorsten Kraft: The German government is possibly [deciding] and it was something like, "Can we do something technically, something more technically to not have an internet blackout on a given date?" It was asked by the BSI if there are technical possibilities to not lose the internet.

Eberhard Lisse: Does that answer your question, Peter?

Peter: Well that leaves me to conclude that I'm not the only confused person.

Thorsten Kraft: We can do it. If you like to, we can do it bilaterally because I think English is not my mother language.

Peter: No, I guess that's not a problem of language or communication here. The point is we're now going back to the whole DNS changer topic and how to prolong this counter-measure which, by the way, has interesting policy implications all over the place that I would like to dive into more rapidly than into the technical details, but that's probably not the right venue here. So then I understand they asked you in terms of they really asked a question instead of ordering you to do something.

Thorsten Kraft: Yes, it's just a question.

Peter: So a real question – not in a question like, "Would you please do that."

Thorsten Kraft: No, no.

Peter: Okay, thanks for that clarification and that eliminates the needs for the legal grounds because of course the government can ask questions all the time. How did you respond? So what is happening at DCIX or in corporation with BSI on the topic of DNS changer? Could you elaborate on that a bit?

Thorsten Kraft: What we have done then was receiving the data from the DNS Changer Working Group and pushing all the data to the affected ISPs, proceeding notice and take down and not to look into traffic or to re-route something to a box that is controlled by the German government or by the ISP itself.

Jay Daley: If I could ask a question about privacy again, at what depth into a packet does looking into a packet become illegal?

Thorsten Kraft: The IP address for instance is in Germany is private data. The IP address in combination with time stamp is private data. If there is a domain name or for instance it's in a communication, domain name is private data.

Jay Daley: The protocol field?

Thorsten Kraft: I'm not sure.

Jay Daley: Okay. There is a researcher in New Zealand who has an alternative to deep packet inspection that looks at I think the first four bytes – that's all – and can tell you a great deal about what's going inside by using statistical analysis techniques rather than looking any further and doesn't need to look at the IP addresses either and the first four bytes of the pay load. So if you're interested in that, that may provide you a way around to those things, but you may not want a way around.

Thorsten Kraft: I think in Germany we need to rename the stuff because deep packet inspection and the term is… you clearly say we took a deep look into a package. Maybe if we are doing something like that, then we need to use another term.

Male: Just as we're on the subject, are those existing tools for the DPI IPv6 compatible? Cause I'm thinking that many of the older ones aren't.

Thorsten Kraft: To the audience.

Eberhard Lisse: I'm more interested in what's possible. What is theoretically possible if the legal framework was different, nobody would care about the legal framework.

Thorsten Kraft: But I have mentioned if you have something like a private data involved, then it doesn't make sense in Germany for instance to do deep packet inspection. A domain name by itself is a private data.

Eberhard Lisse: I mean the other way around. What could you do if you were allowed to or if the legal framework was different? What's technically possible if that?

Thorsten Kraft: You see everything. You have the possibility to see everything – the complete communication if you have the possibility to put the packages afterwards.

Eberhard Lisse: Two terabytes a second – can you analyze this?

Thorsten Kraft: No, that's impossible.

Eberhard Lisse: So you would sample.

Thorsten Kraft: With a million, yes. Maybe each one-millionth data package or something like that. Maybe that's possible. But in this… with volume we have it. DCIX, it's impossible to proceed it.

Eberhard Lisse: Alright, any other questions?

Thorsten Kraft: Not in lifetime.

Eberhard Lisse: Can you, for example, copy the data onto a server? It's not possible.

Thorsten Kraft: If I lose, then I lose a lot of time if I write it down to a hard drive and I analyze it, then afterwards I have written down one hour and I need to analyze, let's say, four hours, five hours.

Eberhard Lisse: Alright, any other questions? There's a remote question.

Female: Hello. Yes, Mark Lampert wants to know, "Instead of deep packet inspection, why not flow based analysis? Do you have agreements, plan to obtain flow data from ISPs?"

Thorsten Kraft: This is something the ISPs are doing as far as I know – flow data.

Eberhard Lisse: Can you go a little bit more detail?

Thorsten Kraft: The flow data must exclude the IP addresses as well so we only have the HTTP headers in the very, very short way and you only see some… yes, everything with a private data involved is problematic in Germany. And if you are doing it in, let's say, in a [dark] area or something like that, then you can do that because there is no traffic that's valid or that should be valid.

Eberhard Lisse: Alright, anymore questions?

Male: Sorry for insisting on that. I think there are two issues mixed here. One is the ability or not to do deep packet inspection at an internet exchange which I wouldn't really see in Europe somewhere. The other one is dealing with the peculiarities of DNS Changer which actually doesn't really need deep packet inspection but there are mitigation strategies available even through the ISP. They had been available and wouldn't have needed all this policies [sent forth] and everything so far.

So what's probably going to happen is that at I don't know what date in early July – 7th – thanks, when the centrally run name servers or resolvers, I should say, are to be turned off and maybe the announcement withdrawn, then this traffic directed to the well-known formally malicious IP addresses… or sorry, I shouldn't say malicious IP addresses – but IP addresses of these rogue servers that had been taken down so long ago.

So when these disappear the traffic will probably still flow but just sync. Did you talk to ISPs – what they are going to do with this traffic? Would that fall into your mission here? And could you elaborate on that? Like setting up their own version of the server or doing something completely different?

Thorsten Kraft: No, they will not do that. They will let the customers call into their support lines in Germany and guide them to fix the DNS server on their own by using tools or whatever. So nothing will be set up from the ISPs I have talked to. I know at least Vodafone and Deutsche Telekom that they will not set up something like that.

Male: I just want to give somebody a notice that was information that before USA security discovered a lot of botnet and dark net of internet. They actually run this kind of alternate DNS services for around nine months. So this is a [good] reason if you would direct traffic to different and create something where you can observe the activity. That's what I read so this can also explain how it's used for what purpose.

Eberhard Lisse: Anything else? Alright, thank you very much. It was a little bit shorter than expected, but interactive and it's much coming from the public. There's not much we can do about it but in any case, thank you very much. In particular I appreciate that you came out on relatively short notice and on your company's own expense. And I wanted to mention

that next time I see Mr. [Walter] I will tell him again how much I appreciate it.

Alright, the next thing will be a rather pleasant thing, at least for me, not for Antoine also because the Dutch Registry is sponsoring lunch so we usually give a few minutes time for a word from our sponsors, in particular when the sponsor is non-profit.

Antoine [Kantiza]: I'll fill in the gap then. First of all I'm very pleased to be here again. SIDN has been one of the first ccNSO members and I remember I also visited the first ccNSO Tech Day which I think was back in Puerto Rico or was it before that? Anyway, a few words about SIDN of course. I won't be in the way between you and lunch.

SIDN runs the .nl registry and also some small [enom] zone. As of this morning we have 5,900,082, 256 delegations which means we will be running over 5 million very soon. Since the 15th of May we have DNSSEC in production and we now have 2,176 DNSSEC signed delegations. We just announced a two-year incentive for the DNSSEC deployment which will probably mean that we're hoping to have 200 .nk DNSSEC delegations by the end of the year.

It's targeted at registrars and we have a lot of large registrars that already have signed into the incentive. So I think it will give a boost to DNSSEC deployment, at least in the Netherlands but we hope it will be outside the Netherlands as well.

I have something new to tell you. I know for DNSSEC deployment there has been some discussions about secure transfers of domain names.

We just performed some tests on… not really tests… in a real production system we have just performed secured DNSSEC transfers, so we know it can work. I will have a blog on this later this week and probably for all of you who want to talk about it later for lunch. That's a great opportunity, of course. I'll give you some subjects to discuss for lunch of course.

SIDN has recently increased their DNS resource and development. We have formed an SIDN lab like most registries are now going into some more research and developments. We have just done the same. We've done increased cooperation with NLnetLabs but that's not the only outreach that we're doing. We're also sponsoring the BIND 10 project. We also have sponsored Power DNSSEC and in that context the discussion of this afternoon where we will be sort of discussing multiple vendors or multiple DNS implementation is one of our interesting topics.

We are setting up at this moment a DNS work bench where we will run and test all this DNS software and we will report on it and even give others the opportunity to make use of this DNS RND program. If they want to do tests on name server software, they don't run themselves. SIDN has been an independent contributor to the internet's development and we want to continue to do so and that's why we want to sponsor the ccNSO Tech Day because we believe in sharing knowledge and I think the ccNSO Tech Day is a wonderful opportunity to do that. So without further ado, I want you to enjoy your lunch and have nice topics to discuss.

Eberhard Lisse: Where is the applause? I think the biggest applause came from cz.nic because I had them online to sponsor lunch if we haven't found a sponsor. So we had, as we say, we had a cunning plan. Anyway, thank you very much everybody who presented so far. We'll meet exactly at 2:00; we'll not even wait a second because we are going a little bit early now.

We'll start with the host presentation. Ondrej will tell us about what they're doing and especially what they're doing on DNSSEC. I like the word incentive about DNSSEC. Does that mean you give your registrars who do DNSSEC a discount, Ondrej? Are you giving your registrars who do DNSSEC a discount as an incentive, or how do you do that?

Ondrej: Yes. We're giving them the wholesale price. They will get a discount and it's also… it's a lower incentive so it's going to be two years and I don't know if there are thresholds. Do you know? There's no thresholds.

Eberhard Lisse: I'm thinking about the same thing. As you may know, I may be ask 2,500 names of which seven are DNSSEC signed and we are also thinking of pushing this by telling our registrars that they get the discount, so they have to start not only signing it, but they also have to install resolvers that can resolve this because there is not a single resolver that can resolve. When we do this we all have to use the OARC ones. That's all.

Alright, thank you very much for your attendance. We'll see each other exactly at 2:00. Lunch will arrive at exactly 1:00 - at least that's what the plan is. So we must be patient a little bit, hang around, talk, communicate, whatever.

[break – next session's audio begins in progress]

Eberhard Lisse: …usually the ones that are there that get insulted for the ones that are not being there. But I hope that even after in this postprandial bliss we can listen to what Katrina has to say about measurements of update propagation speeds. Take it away.

Katrina Sataki: Good afternoon. So my name is Katrina Sataki and I do represent NIC.LV. That's my first presentation at a Tech Day so just please be tolerant. Definitely I'm not the geekiest person in our registry but I'm the one who does the talking. So I am from .LV which is the country code top level domain of Latvia, small country in Europe and the registry is rather small too. So we are approaching 100,000 domain name threshold – really slowly approaching. It's like two steps forward; one step back; I have no idea why. We have 98,000 at the moment. I hope that we're going to have 100 this year.

Speaking about our registry, we introduced our registry/registrar model in December 2009 and at the moment approximately 26% of our

domain names are registered *via* registrars. We still have a pretty impressive share of direct registrations. So our customer support center gets a lot of calls from our own director, registrants and from our registrants who register their domain names by registrars.

And basically quite often they say, "My domain is not working. I've changed the IP address using your online system," or just on their own DNS infrastructure and, "When I try to access my website, there's still the old one. What's wrong with the registry?" Well, there's nothing wrong with the registry but the thing is that the moment you change for example IP address for your website, it won't start working immediately.

Why? Why not? It really takes a lot of time, a lot of explaining just to try to make people understand how this domain propagation works and they do not and that's quite okay because they don't have to… the experts in this field. But at some point we realize that we need something to make life of our customer support easier. We need some tools which would lead them to find out where the problem is.

Therefore we decided to develop some application; we call it digSense and it basically consists of two parts: its server side which is web engine and engine register collects statistical information from clients. We call them sensors; I think it sounds more scientific and way cooler than clients.

So we have sensors and the initial idea was to plant sensors somewhere and to do the measurements to let our customer support just to enter a domain name and see how it looks from different resolvers and cache servers. So that was the initial idea and so we implemented it.

But of course the main target here is to have as many sensors as possible and try to plant them in different ISPs. We deliberately decided not to contact ISPs because we did not want them to influence the results of our measurements. Therefore when we had these initial sensors ready, so we started to think about our strategy – how to build this network of sensors.

So the first and most obvious solution was just bring your work home. So we just asked our employees to bring and install these clients which are simple software and written and patent, so it's really very easy to install it. So they just bring it home, install on their computers and just leave them there.

Unfortunately the number of our employees is quite limited and of course, when we hire new employees we cannot discriminate by their internet service providers. Therefore, we just try to think about other ways how to plant these sensors and so we decided let's let others to do it.

So we decided to look at our Anycast servers – okay, I will call it Anycast cloud – and so we gathered data from our Anycast servers and tried to filter out those resolvers querying our .LV servers and filter out those which are located inside the country, at least on IP addresses of the country. And it turned out that we have about 6,000 resolvers inside the country.

I have to admit that was really surprised by this number because there are like 300 internet service providers in the country. Of course, most of them – like 95% of them – they are very, very small, regional. Some of them probably serve some residential building so of course, the

number looks really, really big and frankly here I think about our inspectorate which has to supervise gambling in the country because they wanted to have legislation … they are going to have legislation which will allow them to request the holders of recursive servers to block access to domain names. So I'm just wondering how they're going to speak to 6,000 entities to do that.

But okay, that's their problem. Anyway, this number is quite impressive and out of this number we have 584 resolvers which are open actually. They are open for anyone in the world. Later when we found out these numbers… we found out there's a tool – I forgot the name, sorry – but there's a tool which claims that there are around 470 open resolvers in Latvia. We found out that there are at least 584. So what we did, we installed these sensors in some instance on our server and just used these open – not all of them but at the moment like 70 of these open resolvers – to query for domain names.

So how does it work? So now we have a tool for our customer service. The real data from one of our customers who complained about their changes that they made for their website. Well, to our great surprise, for example, some resolvers, they just do not have any information about particular domain name and we have no idea why and that's actually the DNS server, of course the server of one of our largest service providers.

What do we have here? First of all we have those sensors, we have separated them by different internet service providers. It's okay to have several sensors from one internet service provider - it just allows us to do some comparison. Then we can see the status of the request

from that particular server. For example, in the last line there was one fail. In the last two columns you can compare information from IP address, information from authoritative server and the second one result from dig.

Well, particularly we're very interested in ttl on authoritative service and ttls they have on their particular cache service. This too is really, really helpful for our customer support, but that was not all that we did. Now going back to the previous slide, we had quite many open resolvers and we found out that more than 120 of them are really honorable and I think that the record here is one server which still had BIND 4.9 on it, so a really bad thing.

Then we contacted our national cert team - cert.lv - and together with them we started calling the owners of those servers and trying to explain to them that that's not a good idea to have. First of all it's not good to have it open for everyone… well, in many cases when those are internet service providers, that's their policy. They just want to make life easier for their customers so they just leave their resolvers, cache servers open. Well, that's okay.

But for enterprises or even for home users who have their own servers, that's not a good idea to have them open and we tried to explain it to customers, but unfortunately they do not understand the problem. They say, "Come on. Why not? Google has it." Yes, but your research says and research says of Google that they're different a little bit. Still do not understand it. They try to explain that it's not a good idea to have BIND 4.9 for example on their server.

That's another – not help but from this project and I think that we have convinced at least some of them that they have to do something, upgrade their infrastructure.

So what are the future plans? First of all we want to make digSense public, well at least semi-public meaning that offer it to our customers. Those who have accounts in our online system – they can for example use it to query for their domain names and see how it looks on different internet service providers. And of course we want to add more and more sensors so we would like to invite people to add those sensors, to cover as many internet service providers as possible.

And for that of course, we will make the software Open Source and just invite people to install it. It's really very simple. You can install it on Windows, Linux or anything. Well, once we have an extended network of sensors of course we would like to add more tests. One of the ideas is to use it to just measure responses from our Anycast notes to see how it works there.

So basically we want to have some public tests which any who has joined the project could measure and also some internal tests for our own needs just to make sure that our infrastructure works properly. So that was a quick introduction to the project. Are there any questions?

Eberhard Lisse: Any questions? There you go. I was getting worried, but not from you.

Robert Martin-Legène: Hi, Robert from Packet Clearinghouse. You have at least 6,000 sensors. Is that true?

Katrina Sataki: First of all, as I told… at the moment we have 70 sensors – 6,000 servers which I used for resolvers and we can't use all of them because only 584 are open. So theoretically we would be able to plant 584 sensors, but we haven't done that yet.

Robert Martin-Legène: So the 6,000 name servers that you have diagnosed as being a recursive server, what you want to do is to go out and see if it has something in its cache and the detail of that?

Katrina Sataki: Okay, we analyze data from our Anycast servers to find out how many servers within the country query .LV. Then from these 6,000 we try to find out open resolvers. So the 6,000 – forget about 6,000 – doesn't matter.

RobertMartin-Legène: But 6,000 is a good number.

Katrina Sataki: Oh that's true, it's a surprisingly good number. But it's irrelevant actually. It's just 6,000 and out of those 6,000 584 are open, so theoretically open for us planting a sensor.

Robert Martin-Legène: And how did you get so many – what 70 – probes, sensors? How did you get so many? I mean people are stupid enough to actually download something just because you say it?

Katrina Sataki: First of all, some of them are made by our own staff at home and others are… I'm telling… okay, we have this server, digSense server, and on the server we have, for example, find out one open resolver you want to use, you just make a folder for this particular server, you install [pliant] sensor in that folder and configure it the way that it sends requests to that particular open resolver.

And so you add it to the network of sensors and the moment when you request the query for domain name, it just goes through… sends a request to all the sensors and then on one of the ports and gets information from all of them then stores the information in the database for future analysis. You can see how it used to work in previous requests, for example. So you can see history of requests previously for this particular domain name and well, that's it.

Robert Martin-Legène: Okay, thanks.

Tomas [Mackus]: Tomas from .lt registry. I want to ask do we have good relations with your biggest ISP who has DNS resolvers with most of clients in Latvia. If your client doesn't work for some part of and you see it maybe doesn't

work for largest ISP because its cache is all done. So do you have relations to ask them to flush its cache resolvers?

Katrina Sataki: First of all, yes, of course we are very interested to have those sensors from our largest ISPs. For example, if you look at – went back to this slide – the first two. Actually the largest internet service operator – later I will explain to you why we haven't contacted it. I'm not that nice. But yes, that's one of the ways of course to…

You see, large operators, normally they have many customers and they do change their settings quite frequently. Okay, not quite frequently, but they do change. If you have many customers you face more changes. Of course internet service providers would be interested to fix their cache server not to flush every time when somebody has problems.

So yes, that's the next step of course to help internet services providers to insure that their servers work properly. But anyway, at least now we can identify the problem and if there's a customer who has problems we can tell where this problem exactly is.

Eberhard Lisse: To abuse the Chair, at the prerogative of the Chair, how are your customers taking to this? Do they like being told that [they can now] get it sorted? If I did that, I would be promoted from page 16 of the African speaking dailies to page 3 probably again. They don't understand things but I mean I could go into some detail which is really great fun. It's always the registry's fault as you mentioned. Why is it

not working? My internet is not working. When you tell them you've got this and this, how are they taking to this?

Katrina Sataki: It depends on the person I think. Since we're trying to be nice and help our customers, they are okay. But honestly I'd love to be promoted from page 16 to page 3 or even higher, but no, that's not the case. They want solution and if this tool helps us to give them solution or at least point out where the problem is, it works fine.

Male: You probably don't know this but in my country page 3 is reserved just for special people. [laughter] And I must say, we are English-speaking too in my country.

Eberhard Lisse: Or as they say, you would be quite becoming. Alright, anything else? Jacques.

Jacques: Quick one. You said you have about 300 real ISPs in Latvia. Can you ask them to white list your IP addresses so you can query directly instead of putting sensors and stuff?

Katrina Sataki: Well theoretically, yes but again, as I mentioned at the beginning of my presentation, we did not want to contact ISPs directly because we did not want them to tamper with the results. So first of all we want to

understand the problem then we can address the problem. Understand the problem then to help our customer support and then we're moving forward to just offer it to the public and help ISPs to fix the problem.

Eberhard Lisse: Anymore? Anything from remote? Nothing. Alright, thank you very much. Okay the next presentation will be Ondrej Filip. As it has become sort of a custom, we always have the host give a host presentation, tell us a little bit about what they're doing, what their special interests are – some do research; some do other things. And Ondrej's ccTLD is quite famous for having a large number of signed DNSSEC signed zones so it's quite interesting to hear how they're doing it.

Ondrej Filip: Good afternoon everybody. It is hard to anticipate what I should present to you because we are trying to be very active and we are presenting almost every Tech Day. So I covered many things just to explain how our company works and what issues we are solving, so I don't expect any really deep details. And I hope maybe it will inspire you in something.

And of course I will spend some words about the DNSSEC but mainly from the technical point of view because I don't think there is some time to talk about marketing stuff we are doing then, but honestly the number of domains is more marketing than technology of course.

So at the beginning I have a few let's say boring words about the company but just for you to know who they are. You are probably very surprised but we are the administrator of top-level domain .cz and also

the ENUM. And we have crazy, crazy legal forum - Special Interest Association of Legal Entities – which basically means that just companies can be members.

We have open membership and we have about 100 members. You just pay a very tiny fee to be a member; the membership is growing. As many others, we do not deal with the industries directly so we have registry which is Tom O'Dell and we have about 46 registrars.

Majority – at least majority of the largest are based in Czech Republic but we have also 10 abroad, so that is the structure for registrars. What's more interesting we have roughly 50 people – I didn't count it properly but that's about it – and I think more than half of the company is not really focusing on the domain itself so we have very strong R&D Team and we really spend a lot of energy in the community projects and Open Source software and things like that.

Those guys are based in two cities – Prague and Brno. Brno is the second largest city and the capital of the Moravian part of Czech Republic. So that's our other locations. And some of you could see yesterday our Prague offices – and we are not-for-profit, so Eberhard is smiling because it's his favorite phrase. That basically means that we don't do profits. We need to spend the money on some community projects, so that's probably why we are so strongly focused on that.

So what do the incomes come from? We have domain .cz and that's not a very sexy suffix, so just Czechs usually are used to that. We have like 94% of registration from the Czech Republic and the rest is from Slovak and then countries like U.S. and other countries, so we are nothing important. It's a meaning just for Czechs of course.

But we don't have any restriction on numbers presence so everybody can register .cz domain if he or she is crazy enough to do so. And we are strictly first come/first served – nothing surprising. One of the quite nice features is the arbitration… I mean the dispute resolution use the Czech application court which is a company providing not just dispute resolution for us, but also for .eu and also as a UDRP provider.

Although we don't deal with the end users directly, we do things for the end users. For example, if an end user wants to protect or to be protected against transfer or against any change in his or her domain, we offer end user domain locking so such guy can contact us and we will lock the domain for him and his domain cannot be changed.

We have 24/7 end user support so not just registrar, but everybody can call us 24/7 and one feature that is quite we are proud of is the comfort notification. If the domain is going to expire, we are trying to abuse the end user with call, email and snail mail so we are really trying to notify the domain name holder that something is going to happen that he should pay the fee of course.

Another step which now we have used a lot is database cleaning. We are really moving forward to identify the end users. It's another big area – I don't think I will discuss that farther but we are really trying to be sure that the WHOIS database is correct. We have certain ways how to validate those data and we have several levels how an individual or company can be identified.

And the one thing that might surprise you if you walk through Prague and you see a lot of tiny little funny things above the letters but we do not supervise IDN. Every two years we ask the people, we make a

survey and every two years we get a response that the Czechs don't want to use IDN in their domain so we don't support that.

The numbers are happily growing. That's the same graph as you would probably be able to show me, so it's growing; and we will reach a million of domains approximately – I don't know, by October of this year, so we plan to have a nice celebration at the end of the year.

Now some more important stuff for you as Czechs probably. This is a very complicated picture how the registration system looks like. Here in Prague we have three main locations which two are used for production. Those are two tele-houses. And also we have some small but (inaudible) in the offices but they are not used for the [production] things.

So those three locations are circled so in case of any fiber cut and anything we don't lose connectivity and of course, every location can work independently so in case of some blackout or something we should be able to work. And one more thing which is not on this picture – we have also a third location which is outside Prague so just in case something really horrible happens here, we will not stop operation and we will be able to register domain and provide services from the other people outside the city.

So we have three independent locations that are replicated and running simultaneously so that's how the registration system looks. More about the DNS system. We operate our own four Anycast clouds; they are called A, B, C, D and you can see we have very good coverage in Europe because that's the key for us but we have our servers on the east and west coast of North America; we have several in Chile; we have a server

in Japan and if you try to take all domain, you will see that we have actually five name servers so the fifth one is a Unicast one and they are afraid to remove it from the zone so it's still kept there. So we have four Anycast clouds and one Unicast node for DNS. And as I said, everything is operated by us so we don't outsource anything.

Just very briefly, the other activities - we have a lot of stuff for the local community. We are printing books, making conferences, provide training. We have our own academy in training. Actually the Masterclass was held yesterday. We try to somehow play with openID technology or open ID-related technology so we are working as an identity provider because as I said, we try to validate the data and WHOIS database so we use those data to offer the local communities and new services and things like that.

We host some of the root servers here – first was F from ISC and second is L which is the first node of L outside U.S. and I think it used to be – I'm not sure it still is – the main distribution node for Europe. And also we are helping some other domains. We are hosting some secondary service of some other so we have Tanzania here and we exchanged a name server with Chile so it's a mutual and very good cooperation with nic.cl.

Now to the software. The whole day I'm talking about software so excuse me. One which is related to software, we provide Open (inaudible) or DNSSEC (inaudible) resolvers that's somehow integrated with another software that Java introduced in another slide. So because we couldn't see any really (inaudible) resolvers before we built that now I think (Inaudible) and other people. But that's a service that we offer.

Briefly the keys for a package is because about FRED you heard a lot so it doesn't make sense to introduce you again. Knot DNS which is a new thing will be introduced by Andre later so I will not touch that. Another package which is worth mentioning is called BIRD. It is a very successful software, not really used by the normal user, but it's the most popular out server in the ISP world, so the biggest internet exchange points use this software as root servers because the demand is really high and this software is very efficient and either open source or commercial implementations we have problems to work out, so that's why I brought this various expertise in this area.

Now I'm going a little bit to the DNSSEC so first project that was related to DNSSEC and that was something that helped us create the visibility of DNSSEC was the DNSSEC hardware tester. It's one thing what side and also on the website you can download some client software which will help you to test your network. So what you do – if you go to the website you can download the client for Windows, Linux and Mac OS.

And this client tests your network and sends a report back whether you will be able to validate at the end users' side, your network, your, I don't know, ADSL router and everything which is between you and the server is okay and smoothly support DNSSEC. So it gives you an overview. And honestly those results from the download is very nice looking, so validation at the end to decide would be a real problem in the future.

So we keep the database of the results and we provide the (inaudible) which is English, Czech and Hungarian as the Hungarian registry was really interested in the project and helped us to translate it into Hungarian.

Popular stuff – I think one of the most popular stuff we ever developed is the DNS Validator. We started as a Firefox add-on so it's an icon that is similar to HTTPS and it shows you whether the domain on which the website is is signed and whether the validation process is correct. That's the software that uses [our open] DNS validating resolvers as I mentioned earlier.

I think last time when I spoke about this we supported just Mozilla Firefox, but now we have a Google Chrome and also some version of Internet Explorer so whatever browser you use except Opera – I'm sorry, we don't support it yet –

Male: And Safari?

Ondrej Filip: And Safari. So you can use it.

Eberhard Lisse: I use GoDaddy.

Ondrej Filip: Okay, so even for Eberhard it works, so that's perfect. Some new stuff I don't think we have ever published it to you – we use it for some national projects – is this kind of HTML widget. If you go to www.nic.cz, you can see it on our main page. It's a widget that informs you about whether your connect supports DNSSEC and IPv6 in very simple graphical way and if you wish you can also measure the speeds. The

speed measuring is taken very informally; it's something for information; it's not very precise, but it just shows that there's no huge difference between IPv4 and IPv6 if it's configured properly.

So that's another tiny project that is going to be a little bit customized so you don't have to support the speed measuring if you do not wish so. Here's the URL – you can go for that. Another thing it is Open Source and if you like it, just use it. We will be happy to help you with implementation on it also if you wish.

So those are tools we use to increase the visibility of DNSSEC to have something to talk about with the local community. Before I show you the results, let me give you a brief history of DNSSEC in Czech Republic. We started quite early – I think it was April and September 2008. We started with ENUM zone because we hoped that nothing can be broken if the process did not go okay. So we started with ENUM and if nothing happened, then we continued with .cz which was signed on September 2 and I think it was the fifth TLD signed by DNSSEC.

Then we opened registration for the end users so it was September 30. And we started to use NSEC at that time because NSEC3 was not deployed. Then the root was signed on July 15, 2010, so two years later and after that we decided to make a very exciting exercise and it was the change of the algorithm and it was a trick one. I think I had a presentation about it on Tech Day so if you are interested there is a presentation online about it. That took almost a month of work and fine-tuning and at the end we succeeded, so now we use NSEC3.

We don't use [Opt Out] because we believe it wouldn't have much sense for us because we really wanted to grow the number of signed

domains and so otherwise we plan to have a lot of DS records in the zone. And here's the current situation. About 36.5% of domains are signed so it's roughly 350,000 of domains of the less than 1 million.

We publish all the numbers directly on the web page so if you go to www.nic.cz you will see the numbers online. They change every second so you can look there. What happened or how we could achieve those high numbers? We identified some stakeholder groups that we need to attract, to talk with them and we used very different ways how to communicate with those guys.

At the beginning of course we started with the registrars; we did a lot of seminars. As we have the accommodations, the room for training, we also offered them training and we of course discussed with them directly and we had a lot of talks about local conferences and offered technical support to anybody who was planning to implement DNSSEC. So it's gone very well.

And it wasn't bad. Since the beginning we have a reasonable number of registration, but of course we need to do a little bit more, so we also offered a financial incentive. We do not give discounts to registrars – that's not that way. But all registrars can make a marketing campaign and if they use .cz in the campaign, we are able to pay our half of the campaign back. There is a gap depending on the market share of the registrars.

And in case they support DNSSEC and they have a really large number of signed domains, we can increase the gap, so that's the financial incentive we offer them. So they still need to pay 50% of that but the gap is much higher so they can plan a larger campaign.

We also started a public campaign, a campaign in the media. I think it is also discussed not directly here but I think on ccNSO. We created twins; we found the twin of some celebrities and we make those twins to make something really crazy like something which is strange that those original guys would never do. And we make an amateur looking video of those and then we sent it to the media and it is of course spread around and everybody is watching what happened and they were really shocked. Those videos have millions of visitors and things like that.

And then we revealed that it was a fake; that it is just a joke. The message was you can be faked in real life and worse things may happen in the real world, so use DNSSEC. We have about a million visitors and they said you cannot a video where DNSSEC is. So it was very popular and we of course continue a direct communication with the website and ISPs.

And that time after this campaign we are very successful because many large ISPs in the country started to validate so that was great. And now the validation at the ISP side is very well covered. And I mentioned the tools we developed. So now we have the number of domains I mentioned.

Another good thing is that all major registrars with more than 90% of market share support DNSSEC and the majority of them signed by default so that's why we have so much a high number. But it's not the only number, you know, small end users that signed their domains.

Many major ISPs validate; two or three cell phone operators; Telefonica which is the largest incumbent in this country, so the largest company providing a DSL and everything, Vodafone as well. A lot of B2B ISPs are

validating, the major ones as well. And with the direct communication we are able to finalize it on some important websites so many important websites sign, like news, magazines and e-shops, so that's very good. We still need to work a little bit on the government institutions but that's a little bit hard.

And I guess almost 20 minutes so this is my last slide about DNSSEC and it's the forecast for the end of the year. Of course, the second quarter is almost over so the red graph is roughly good. We plan to have about 40% of signed domains at the end of the year so it should be around 400,000 domains. Now it doesn't grow so quickly. It's hard because we picked the easiest parts of domains that are hosted by registrars and now we are trying to really communicate with every single website, especially those from the top in ranking and it takes some time. But those sites are of course more important than the high numbers we could generate easily. So that's all for my side. Thank you very much.

Eberhard Lisse: Thank you very much. Any questions? Steven.

Steven: Hi there. Thanks for the talk. Are you guys satisfied with your uptake of DNSSEC of 35%?

Ondrej Filip: No. No! But we didn't stop the campaigning; we just set up the communication. We really want to get more and especially we need to

work more on the (inaudible) sides here in the country – that's more important than the numbers.

Steven: Do you know how that percentage compares with other registries that have implemented DNSSEC? I'm thinking of .se.

Ondrej Filip: I think it's… I don't know. Is there anybody having more than 10%? I would doubt…

Male: .gov.

Ondrej Filip: .gov. I don't think any… but I'm sure we have higher numbers than them even if they… there's a default policy but still they didn't sign everything.

Male: They're at 15%.

Ondrej Filip: 15, 15 yeah. So in .gov…

Male: .gov have about 15-20%.

Ondrej Filip: And that's one that radar so it should be 100.

Male: No, well, not quite. There are entities in there that are exempt from the mandate.

Ondrej Filip: Oh, I see. Thank you for that.

Eberhard Lisse: Yeah but how many of the mandated ones have signed it? That would be interesting. Okay, Russ Mundy, can you answer that question?

Russ Mundy: I can't answer it authoritatively but with a reasonable guess around 60-80% of the mandated ones are signed in .gov.

Male: (Inaudible) from Brazil, .br but I've heard that in [Conviar] there are about 300,000 domains that are signed. [Conviar]? I guess that's like 30%.

Eberhard Lisse: You mean .pr for Puerto Rico? .br – Brazil because I see Oscar sitting there.

Male: Some of TLDs have their DNSSEC horror stories. Have you ever had any incidents yours – maybe customers?

Ondrej Filip: You mean what…?

Male: Like the bad stories, you know, like the UK had its own horror. Have you ever had…?

Ondrej Filip: We reported the story about the NSEC to NSEC3 transition. We tested everything with BIND and it worked perfectly and then some unbound has problems with validation so it was for a short period and I think Olaf is here somewhere – he informed us of somebody from the internet discovered that and informed us and we reversed it back and redesigned the process. So yeah, that was a bad story but for a very short time and because we were working on things that never happened before, a change of the signing (inaudible). So it's good to know.

Eberhard Lisse: Alright, thank you very much as usual. So the next is a round table where we will have the four eminent authoritative name servers being presented one after the other. And then Jay will ask some pointed questions to stimulate some discussion. And afterwards we will have a short feedback from the Masterclass that we had with FRED yesterday

and then I have found that Luis was not too vehemently opposed to close the session.

Jay Daley: This year's been an exciting year for name servers. We've had two new name server implementations released and the two major implementations are both working on new versions. And as we're all good friends in this community, we're all here to talk all about our marvelous work that's going on so we're going to go through a presentation from each of the four name server operations or name server developers and then the chance for you to ask questions afterwards as well for about an hour about their future plans and elements of technology they may not cover in their presentations. So first we've got João about BIND10.

João Damas: Can everyone hear me? I hope for the best. Okay, my name is João Damas. I work at ISC and today I'll be talking about BIND9 and BIND10 just very superficially and then we can leave the details for the round table discussion later.

Generally speaking, what is BIND? BIND has had four main versions so far – BIND4 which some people are still using amazingly; BIND8 which was basically a renaming of BIND4 with some added work, but it was a time when everyone thought that 8 was the cool version to have like in [send mail] and all those programs.

And then by 9 which was finally a complete implementation and now the BIND10 version that we are working on which again is a complete

rewrite. But what is BIND? BIND is mainly two things – 1) it is the DNS reference implementation out there and the second thing is it's Open Source software. Those are the two main factors that define what BIND is.

When I say reference implementation, what I mean is basically what's written down there. It's an implementation; a piece of software that follows and implements the protocol standards. We also make it available to use and test new protocol ideas through some side releases – snapshot releases, special branches. And this was for instance, very useful I would say even crucial in the development of DNSSEC because BIND basically already supported all the previous items it was easy to extend to test out new ideas and enable to protocols to evolve.

The third one means basically that when behaviors deviate from the protocol, from the grade standards, they'll either be off by default – users will have to turn them on if they want them – or they will be available separately – not in the main trunk. So what you get from us when you download the thing is a pure implementation of the protocol.

Open Source is another crucial aspect for us when producing BIND. Basically it means three things – 1) is that the ISC license – you can look up in Wikipedia – all you want to know about ISC license. Basically it's an edited version of the BSD license or a simplified BSD license. It basically says… it poses very few restrictions on what you can do with the software.

This enables anyone to modify it for their own purposes without actually needing to give it anything, any modifications back to the community. Of course they are welcome. Anyone who produces

modifications or extensions or any sort of alteration of the code is more than welcome to contribute it back and we'll do our best to put it in the mainstream if it makes sense.

And third, another good benefit of it being Open Source with the kind of license that it has, is that if ISC goes astray and anyone else can pick it up and run with it, continue development and do something with it. So that provides for a kind of stability and security for your infrastructure. That applies to all versions of BIND that ISC has cared about.

The current ones we have in (inaudible): the first one, BIND9 is the one that we would recommend for you to use in production right now. This will change in a couple months. It's been available since 1999 or perhaps even a bit earlier if you had access to the better versions. That's like 13 years, going on 14, so it's had a fairly long evolution. It started 9.0 back in that year; we are now at version 9.9. Each successive major version has added a substantial number of new features. I'm just going to mention a few new ones that may be because of them being really recent people may not be aware of, but I'm not going to go through all of them.

So I'll mention inline signing where basically you can drop BIND into an established workflow so that it takes unsigned zones – like normal DNS zones – signs them and spits them out signed the other side without having to change your systems too much.

RPZ – Response Policy Zones – which is not a DNS standard by itself; it's more kind of a filtering mechanism, a local policy aspect that allows a recursive server to modify other what it does when it receives a query.

In every single release we have had ever since 9.0, and particularly for all those since 9.4, we have increased performance in all different areas. The reason why I'm not going to go through new features is because that for instance, is the listing of new features in 9.9, so I don't think it makes much sense to go through there.

ISC has this thing called the Knowledge Base which is open to everyone who wants to go and read it. It's just basically a bunch of answers and questions. You can find all of these things there, otherwise you put everyone to sleep.

So that's BIND9; what's up with BIND10? BIND10 is the new version of BIND. It's currently in development, in very active development. We have a group of eight people working on it. Some are directly working as ISC employees; some others work with different companies, for instance, CZ.NIC, our host today, a country with one engineer who is active in the development and then so do CN.NIC and GPRS.

It's a sponsored effort. We had a number of mainly registry operators who pulled together to add their financial contributions to enable to project. And then we started the work. It's on the third year of development. Well the third year finished in April, so actually the fourth. It is based on a completely new architecture. It uses new programming languages; it uses C++ rather than C as before. Basically the things that need to go faster are using C++ as a compiled language.

The things that you can… where it's more interesting to have system administrators to change things quickly to meet their needs rather than to have top performance are done in Python because that's easier for other people to mess with. And it has a completely new [open] style.

Even though both BIND9 and BIND10 are Open Source, BIND9 comes from a time where it was decided, due to many circumstances at the time, that we would not be able to rely on a lot of the infrastructure that was provided by the operating systems. It was a time where you had things like OS4 and Solaris was emerging and IX and HPUX and ULTREX and all these other flavors of Unix that weren't quite the same.

So what ended up happening in that environment is that a lot of the memory management, basically for structure [of our software] was put into BIND9 itself. So it yielded a piece of software that's rather complicated, complex that even though it's completely open, makes it very hard for anyone that doesn't work with it on a daily basis to basically alter it or modify it.

BIND10 relies a lot more on off the shelf solutions. It levers its Open Source libraries, it also is coming into existing in a world where the variance of Unix are greatly reduced and distributions of Linux, but in the end, the kernel is always the same basically. The VSDs are also a lot more unified than they used to be. So it levers that to its advantage to try to make it more understandable and more accessible for anyone who wants to do anything on it.

We also learned quite a lot on these 13 years of continuous work in BIND9 and that's been put to use in BIND10. So what this means is basically that BIND10, instead of being one big chunk of software that does everything you could hope for, is actually a lot of smaller programs that work together. That allows you to choose where and how you start your (inaudible) [list] because the needs of someone with zones that are in the thousands of records may be quite different than the needs of

those who have zones with millions of records or those who have small but many zones like hundreds of thousands of zones.

And there is no way we can come up with a solution that addresses every possible combination of different needs. So what we are doing is basically making it modular and as of now it supports a BIND9 memory like thing for the most common cases, but it also supports a sequel back end so that you can load the zones from a single database immediately.

You can also choose which functionality is available. With BIND9 the recursive and the authoritative server are combined in one and it's quite common for people to make the mistake of enabling both at the same time or do it intentionally and then have an intentional consequences. In BIND10 those two things are actually separate so you can have a pure authority for a pure recursive server.

And then we are adding these hooks for special processing. As the name server goes through the processing of the query and produces the answers, there are many, many steps where different people want to do different things with the queries – modify them, look for this database instead of the other one and so there are special places in the code in BIND10 where you can introduce your own script – little scripts in Perl or Python or whatever your language of choice is to do whatever you need to do.

We are also separating the libraries so that they can be reused by others that want to do DNS software. In the BIND9 case its libraries sometimes forced you to have the memory management mold that BIND9 imposed on people and that made it less than useful for most people, so we're making that better now.

So as I mentioned before in September, so just after the summer, we will be releasing the first beta with the authoritative server for complete testing. There are many more things that we can discuss today. I'm sure you have a few questions but we can address them all later I guess.

Jay Daley: Okay, thank you to João Damas from ISC talking about BIND10. Next we have Olaf Kolkman of NLnetLabs who will be talking to us about NSDv4.

Olaf Kolkman: And there we go. I'll be talking about NSD. NSD is Name Server Demon. It's a product from NLnetLabs and for those who don't know NLnetLabs, we are a not-for-profit foundation chartered to develop Open Source software and open standards for the benefit of the internet. It's paraphrased from Article 1 of our foundation's charter. We're funded through donations and gifts, mainly from the NLnet Foundation but we are also seeking further methods of funding, maybe more on that during the Q&A.

NLnetLabs is known for several software products and that is essentially a segue in what we're hook on to what João was just saying as BIND9 was this big machinery of functionality combined. At the moment we started developing NSD in our line of software we made a decision one to one job. So we've got NSD as the authoritative name server and it does authoritative name serving only; Unbound as a recursive name server so for your ISPs.

And we are involved in the Open DNSSEC product which is a signing and key maintenance functionality that you can hook up in your provision

chain. And then there is a bunch of libraries for which we are known. Another thing which I will not talk about is something we have baptized Credence which is a checking tool for if you want to deploy DNSSEC and more on that during the DNSSEC panel on Wednesday.

Anyway, NSD. What are the goals behind NSD, and there's a little bit of history around this. NSD came into being – and I'll show the timeline in a bit – but NSD came into being around 2000. It was first thought of around 2000 when Ripe NCC operating one of the root servers K, Daniel Karrenberg at the Ripe NCC talked to (Inaudible), the Director of NLnet Labs at that time and they both came to the conclusion that there was essentially one single code base on which all the root servers were running. That code base at that time was BIND8; 9 was not yet put into production at that time at many root servers., so most of the code base was either a clone of some variety of BIND8.

And they figured that is essentially something that is bad for something like a global infrastructure. Biodiversity is a good thing and that was one of the primary reasons to set out and develop NSD. And that's also been the primary reason for us to go into Unbound. Another thing is that at that moment, DNSSEC also came on the horizon. People were talking about it and it seemed ready at the time; that was of course a mistake at the time but we didn't know at that time.

And we were afraid that having only one implementational DNSSEC would be blocking, so biodiversity again providing alternatives to an implementation is important.

Simplicity – keep it simple. That was one of the leading principles during our development of that. Simple and therefore secure. The less

functionality; the less code, the less chances for errors. So want to design to do the job and try to keep as minimalistic as possible so that you could actually do that job well. And we tried to shoot for the highest possible performance at that time.

So authoritative only; a reference implementation – I just modified this slide. It said "reference implementation," and since BIND said "the" reference implementation, I decided to put "a" reference in front of that. And in fact, I could have copied the slides that João produced. We want to implement the standards as developed within the open standards community.

Secure, independently written. We don't take code from other implementations so while we're not claiming that we're bug-free, we do guarantee that our bugs are different than the others. And again that has to do with that biodiversity.

History – we started off with NSD1. NSD1 didn't have yet DNSSEC implemented. We first wanted to make sure that we could make something working. Was developed over the 2002-2003 period and its basic feature was packet pre-compile. So changing CPU for memory; having a pre-compiler that writes your zone file, compiled everything into memory, and basically had the packet completely ready to ship at the moment the query was asked.

Of course this is a bit of a memory consumption thing, but that was the design choice that we made and we sort of stuck to that throughout the history of NSD. NSD was put into real production I would say when it was deployed on the root and that happened about 9½ years ago – 19 February 2003 - one of the K root nodes started deploying NSD.

NSD2 was developed to support DNSSEC. We change the compilation inside of the memory a bit. Instead of complete packets we started to pre-compile RR sets – those are the sort of atomic units of data and at the moment of answering you have to look for those atomic units, pile them up and then ship them as a packet.

That piece of software stayed around until end 2006 at which point we released NSD3. NSD3 is still around and it sort of answered to one of the questions that we got from the community – can you please support incremental zone transfers, incoming incremental zone transfers? Because we are a big ISP – we've got a few hundred megabytes on zone content and we cannot afford waiting to see that transfer over a very slow light to a remote location in the Himalayas. And so we facilitated for that.

I made a very small note that we've been supporting IPv6 from day 1 from the first time the network code was written. That almost speaks for itself and I know this speaks for itself for everybody on the table here.

A typical use case of NSD is one as either a hidden Master, although it doesn't do well with outgoing IXFRs and does very well with incoming IXFRs. So usually it's being provisioned as a secondary authoritative server or a slave server, as people say. And of course, in order to provide the biological diversity, there's always another implementation in that same cloud of authoritative service, at least if you follow our philosophy. What you also see in that case is that sometimes NSD will be one of the instances in a cluster that externally represents one instance of a server.

So that is NSD3. There's actually very little to say about NSD3 – it's simple; it works and it's been working for quite a while and we've tried to make it fast. That said, we are developing NSD4 because there are still a lot of people who are looking towards NSD and saying, "This doesn't really fit our needs," and then I'm thinking in particular to the people who have very many zones and a high flux of zones that they want to publish or stop publishing.

NSD3 has the feature. It's there by design so I cannot call it a bug, but it has the feature that every time you want to add a zone to your configuration, you have to shut the server down and do a full recompilation of everything you have. And of course, if you are a registrar and you serve a lot of zones for a particular customer and you have some turnaround time, then you cannot afford the down time.

So one of the key design things that we wanted to add to NSD4 was the ability to – without reconfiguration, or without restart – add zones to the configuration, transfer the data in and keep the server going, so make NSD more suitable for the hosting environment. That is sort of the key goal for NSD4, while maintaining the other features. Keep it relatively simple for that job, as simple as possible, and keep it high performance.

So we started developing NSD4 and this is a list of the features and one of the things that I want to stress is that there are no changes that we envision except for sort of maintenance in the query logic. The query logic of NSD3 is the same as the query logic of NSD4. First thing that we did – what you see here is a timeline of milestones. These milestones are all published in the SVN Directory which you can follow if you go to

our website and we've completed five milestones until now and have four more to go.

One of the first things that we tried to do was do an increased performance by implementing a radix tree instead of red/black tree and I've got some images to show you what the effect of that was. We changed the way that we store the database. NSD3 uses the patching system and writes differences to files – incoming zone transfers; and we now do that in core with specific database that we call the micro database. NMAP matched so using NMAP features to directly map that into memory of the machine.

Addition of new zones now works – that has finished – and in order to be able to do that, we have introduced the concept of patterns, sort of configuration sniplets that you can load using your remote command to align to control the server and say, "I want to add a zone based on this pattern," and the pattern would say that you have a specific master server from which you pull the XFRs or it would give you a shared key that you would use for the zone transfers and so on and so forth.

Also the remote control facilities have been implemented and what we wanted to do now is make sure that all the zone transfer logic scales and doesn't become a memory hog. We are going to assess whether we want to improve or actually start supporting outgoing incremental interfaces.

There's a milestone in which we want to improve the logging and other usability features and then review those. And then as final milestone, there is full production grade release.

Performance increase – these are comparative pictures where the blue line is NSD3; the green line is NSD4 at its first milestone, so things might become a little bit less in performance once we start adding features. But you can see here that we actually gained a little bit of performance – about 10% or so – in fact for all the use cases. The red line is the so-called echo and the echo is reading a packet and echoing it back without actual processing, so that's the theoretical limit of the machine.

During the development of NSD4 we actually tried to do the full packet compile – that was sort of the intention. We decided not to do that because the memory hog for that was drastic – 35 gigs of memory for a relatively standard TLD zone was one of the outcomes and we just decided to not follow that path.

Features – NSD and at NLnetLabs we used the BSD license that allows for proprietary extensions; it also allows to pick up the code, run with it and the same considerations as João just mentioned. If we were to go belly up or disappear, then at least that code is still available.

We have support to the community – Bugzilla interfaces. We provide free support to the community as good as we can and we have committed ourselves to announcing if we ever were to stop providing support of net software, to announce it two years in advance so that you can either pick up the code and place it in another entity; find yourself a support engineer or decide to move to another product of which there are now multiple at this table at least. There's also the possibility of paid support and you can talk to me about that if you would like to. And that's made me talk at least one and a half minutes more than I was supposed to.

Jay Daley: Okay, thank you to Olaf Kolkman from NLnetLabs. We have plenty of time, gentlemen, so don't need to hurry up. So next we have Peter Janssen of EURid talking about Yadifa.

Peter Janssen: Thank you, Jay. What can I say? More or less after what João has said, what Olaf has said, I will sort of repeat the same thing and I can only imagine the horrors that Ondrej will feel when he will repeat a fourth time more or less the same thing.

For those that don't know us, my name is Peter Janssen; I am Technical Manager with EURid where EURid is the registry that runs .eu and has a contract with the European Commission. We built something that's called Yadifa which is yet another name server.

To give you an idea why we started this, I have listed some of the use cases which more or less you have seen also in the previous slides of Olaf. It's almost like I stole his slides and then dressed them up a bit differently, which I didn't, by the way.

The first use case which is for us the most important one is we're TLD; we have name servers out there that's relative. We're running that on BIND NSD as public slaves and we wanted to have a alternative to these two that exist. And if BIND is *the* reference and NSD is *a* reference, well, I came out with a alternative and let's stick to that.

We built Yadifa as a clean implementation, so more or less the same reasoning as NLnetLabs so our mistakes will be even more different

than those from BIND NSD if that is still possible. The major goals were to have a high query rate, being portable, running on different sorts of platforms like all the different Linuxes and BSDs out there; obviously being RC compliant where the most important aspects are; it should be authoritative; it should support DNSSEC and it should support AZFR and IXFR both as incoming and outgoing DNS servers, so it's master and also slave. That's for us the most important use case.

A second one is what registries in general have as a setup. Some of them do dynamic updates so if domain name gets registered, it gets fed through the TLD zone immediately; some do its own generations. But more or less you have a hidden Master that needs to do the DNSSEC thing. We at EURid, we actually are doing the dynamic updates process where when a domain name gets registered, it gets immediately fed into the .eu zone and signs on the fly and pushed out to our public slaves.

Guarantee we're running that on BIND with a nice script in front of it that looks at each changes being done on the registration database and feeds the changes in the form of dynamic updates messages to the hidden Master which is BIND which signs everything and pushes it out. The whole idea is that it differs with the [open] replacement for that and it can take the role of BIND as an alternative to doing that.

The flip point is that we actually want to go a bit further than that and eliminate the separate process that watches the registration database, whatever that may be, and generates dynamic updates and pushes that through the hidden Master, but actually make that part of Yadifa in itself. So the idea is to come up with some sort of a generic extensible

storage back end that will be sitting directly on your registration database, would pick up the changes that are necessary, format them correctly and push them out to the public slaves.

A third use case mentioned before as well is what we call zone management where you actually want to be able to actually add and remove zones. So we're not talking about the content of the zones that we can change with dynamic updates and AXFR and IXFR, but actually adding a zone on a name server to start it to become authoritative for that zone.

And the fourth use case João again mentioned before is that you have a lot of ISPs out there that actually have name servers in place for the end users that actually type in something in their browser, email or whatever, with the domain name in there and there is a need for a recursive name server that's always able to go out and finds the answer and potentially also validate it should the answer be signed hopefully.

So where are we? As of today we have the official 1.0 release. I guess you can see the basic building blocks of key indicators or its authoritative. It supports its own files with all the types that you can see there. It does zone transfers both the master and the slave. It does notify; it does TSIGs; it does ACLs; it does dynamic updates and it supports DNSSEC as far as its protocol 5 or algorithm 5 and 7, so NSEC, NSEC3 and it does online resigning. So if a signature comes up to expiration, it will automatically regenerate a signature and push it out to the slaves.

We have packages available on our website for the popular Linux distribution that you can see there – Centos. Debian and Ubuntu – both

32 and 64 bits. We do FREEBSD. Incidentally, we develop on Linux and FREEBSD so that all the basic building blocks I would say where we are basing ourselves on.

What I did not say is that we have sort of gone through the same ID that NLnetLabs had so it's a basic C implementation, basic C to get the most performance possible. We also have Mac OS available binary package and as of this morning – 10:00 about – there's also available for downloads. It's a BSD 3-clause license, again very much in line with what Olaf has said before. But you can see there you can download it from Yadifa.eu which incidentally is the name servers that is by Yadifa.

So we're eating our own dog foods and even more if you look at the authoritative name servers for EU itself. We have cz.dns.eu – one of the name servers where you can imagine where the name server sits, in which country. Yes, that's the Czech Republic and that's being served by Yadifa. So we have a steady name server. One of our authoritative names for .eu that's actually running Yadifa as of a few days ago.

Where are we going? We still have some work to do on the functionality side. We need to implement all the algorithms that exist in DNSSEC and I just listed a few of the ones that are more obvious there and the new ones being proposed and coming out all of the time – [Elliptic Curve] just to name one will make it into Yadifa eventually as well obviously.

We still need some work to do on the key rollover mechanism, so currently Yadifa will be signed, but it will keep on using the keys forever unless you tell it to do it differently and the whole idea is to come up with something where that will be automated as well with some sort of

information where key holders [are] handled automatically by a default in its master hole at that moment in time.

Dynamic provisioning and configuration – remember one of the use cases – we have something in our labs running in alpha mode where we actually are able to tell a bunch of Yadifa s running to start provisioning a zone on the fly without stopping, without dropping any queries, without anything bad happening.

We have made the – at least what we think – interesting choice to use the DNS protocol itself to do the provisioning of all the name servers. I'm not going to go into the holy war I would say about the split of data plane and (inaudible) plane – that's not really the point. The point why we have the DNS protocol there is – it is there; it's protected; it's secure; it has DSEC; it has ACLs – why not just use it? And obviously we're looking at all the developments in [net.com] from DNS CCN and like that to see where that is going.

On the conformance and performance level, we have worked recently quite a bit on the performance level by implementing work threads. Up to some time ago we had a single thread machine that would actually answer [queries] in a sequential order and we have implemented parallelisms by doing [market threats], and I'll go into detail a bit more on the next slides.

We need to do some more work on stability prefixes and mostly the conformance aspect. We have spent an enormous amount of time coming up with different scenarios of different malformed queries and things like that to make sure that Yadifa first of all doesn't crash when it gets something weird, but secondly does the expected thing. And that

is a lot of, of course crystal ball looking that you oversee from time to time. We're not that crystal clear about what it actually should be doing.

So what we're doing there is we have built a lot of scenarios that we're firing off to the prime candidates being BIND and NSD and see how the three of us, and actually the four of us these days – how are we reacting to the different types of queries out there.

Documentation – as you go through our website you will see that there is a reference manual available. It's not complete in a sense that Yadifa is not complete, but the documentation is not complete either. We feel that more than the basics are there – that should get you started quite quickly in getting Yadifa up and running.

And two of the older, more important aspects, at least how we feel: platforms we still need to see if it runs. It should run but we haven't done any field testing on that which is Solaris and OpenBSD. There are some issues with certain system calls to do P selects and things like that but we'll see where that brings us.

Performance – as we are a new kid on the block, one of the ways that people will look at us is pure performance level. Obviously on the performance level there is obviously some benchmarking to be done but more on that later. What I've shown here on this slide is before the round table here in a test that we have been doing since the beginning. We're running on a 2.6 Linux kernel so one of the older ones and configuring the name servers with one working thread.

So the idea here is to show with one thread one machine that does all the heavy lifting – what can you get out of the hardware that is. And you can see there the purple, the green, the red and the blue. And on the left you will see the different versions that we used to benchmark this.

What you see on the X axis is the number of queries per second that were shooting out of the server and on the Y axis you see the number of responses that we got back. So you should have sort of more or less a 45° linear meaning that all incoming – that's not exactly 45° because the scale got skewed a bit but anyway you get the feeling.

The general idea is as long as it's going straight up every query being sent will get a response and in a short amount in time you will see that you drop off a bit and queries get dropped for reasons that the software can't cope with it anymore. You see that we all go linear up quite a bit. On the one thread on a 2.6 kernel you see that we were able to do a bit more work than our friends around the table here.

Over to a 3.2 kernel, so a bit more recent, and here we have configured… actually we did some tests by running all the name servers at once. We had 2, 3, 5, 7, 8 and so on and so on. And what we've shown here is the optimal work addressed as far as we could see that. And what you see there is that the differences are not actually quite on the same line all the way up there to the – I don't know – 300,000 queries per second mark and then things start diverging a little bit. But more or less what I would say is that on the performance level there are no real issues in the respect that we need to go higher.

What's interesting here – we have spent quite some time on interpreting these results because our friends from Czech Republic have been publishing slides as well and they got some different results. One of the big differences is, as far as we have been able to determine obviously is the work [threats] that they configure, but the second thing is the actual hardware. That is quite some importance in the network cart is it's able to handle so many packets per second just physically getting them in and out is a big bottleneck; and secondly – now I forgot what I was going to say; not that important apparently because I forgot.

So what you can see is we have been doing quite some tests here. The idea is that actually the four of us around the table here will gather together - hopefully somewhere in July if we can make it stick – to actually come up with some sort of a universal benchmarking tool that will allow us to – in a generic I would say optimal way – test name servers on two different aspects – pure performance obviously, but also on the conformance level – is it actually doing what it's supposed to do.

Résumé – Yadifa.eu - website with all the information. You have an email address; you have mailing lists where you can subscribe to and obviously you can download to buy and lease as well as source codes to play around with. [SSAT AB] is de-licensed so you can do basically whatever you like with it; use it in your own produces; break it – make it better. Please talk to us. Thank you.

Jay Daley: Thank you. That was Peter Janssen of EURid. Finally before we go into the major round table, we have Ondřej Surý of cz.nic who will be talking about Knot DNS and showing us some different graphs I imagine.

Ondřej Surý: Hi, my name is Ondřej Surý from cz.nic and I'm a leader of cz.nic Labs which is our department where we do funny stuff. What is Knot DNS? Well, I think almost everything has been said and so we are also Open Source authoritative only server, but there's a difference. We don't have BSD license; we use GPL so if you like GPL more, pick us. [chuckles] I found that one difference.

It's developed in a very open way, including our mistakes – we try to publish everything which can be painful sometimes which we're facing, but… We try to strive to accommodate for everybody but obviously as we are TLD – well, our first focus was on TLD – it's fast, feature-rich and I'm sure there is not enough articles in English to say (inaudible) [implementation].

It's portable, modular. We are around several BSDs, MacOSX, Linux as well. The portable mainly depends on the library we use which is called userspace-rcu library and if you want to know more nifty details, just go to the presentation Knot DNS for the first time.

We are standards compliant – maybe I should say "le" reference implementation. We do both normal transfers and incremental transfers, both master and slave. We support all known RR Type, including all those TYPE#nnnn. We do DNSSEC with NSEC3 from v1.0. We also support TSIG and root zone and NSID. We try to fast track new standards and it's not only because I am chairing that working group. We have a standard for DANE protocol from 1.4.

The configuration – if you like C language and curly braces and semi-colons, you will like our configuration. You can configure some stuff there like interfaces, remotes, zones, keys, logging. We do support Runtime configuration. You just change the config file and say, "Hey, no, it reloads," and it will do that, including all those zones which it should add or remove.

We were inspired by NSD and we also precompile zones so we offload the parsing of the files from [banked] servers to other processes. If you have parallel loaded server, you can even compile the zones onto other machine and just copy them to the production and then say, "Hey, reload."

A little bit about Knot DNS design. We tried to minimize the amount of lookups in the memory for one query so we have acquired a lot of zone structures where we keep all those references to related data which also means we consume memory. Well, not a huge amount but you need to calculate and try before you deploy Knot DNS.

To minimize lookup time we found a quite nice hashing function called Cuckoo hashing scheme which is a hash table in the worst case of a one lookup time and it's mostly lock-free architecture. I also mentioned the RCU library – we use that for this. It's designed for non-stop, run-time updates that's also that RCU that was frequently updated and we also do shallow copies of the zone when there's incoming transfer.

Where we are now and I think this is what we plan for the summer. We already sped up the incoming transfers but there's more improvements coming into 1.1 release. We also want to focus on stability and bug fixes. The documentation, I think it was the most asked thing. When I

mentioned no DNS, the people asked, “Does it have documentation?” So, yes, it almost has documentation right now, but it still needs to fill the blanks.

I think that more speed up of the answering and the performance will be just comparing the graphs. We want to focus on zone parsing and loading because that’s the [worry] and they did a great work.

Even more future plans – we already have some support for dynamic updates somewhere in the repository but it’s not [margin] domain branch yet and we want to include that during this year. We want to support NetConf for the DNSCCM support. Also we want to improve the DNS hosting support for a huge number of zones and enhance the Common Light Interface. Also if you want something from the Knot DNS, just talk to us and we are very open to any idea you might have.

Just a little bit about testing. I have a few slides about – we added file 1.0 which was raised today but it was on a very short timeframe so it’s just two slides. So we have Yadifa, Knot DNS, NSD BIND and Trafgen is the echo, it’s the theoretical performance of the hardware because we also found out that the network interface card can make the difference.

You can download the test zone from the location there – it’s two mio of random mix of unsigned records. The zones artificial generated. The test queries is one million of queries and half of them is in the zone and half of them is out of the zone. And we use just commodity server I think they are HBs and right now the tests were down with the Broadcom network interface card but we were recommended by several people who saw the testing that we should try the Intel cards, so we are going to do that in the near future.

First performance testing is a little bit different from what Olaf and Peter shown. It's based on dnsperf and independent variable here is thread processes and it's note, only applies for Yadifa for release candidates, the new Yadifa can be tuned for that. And the dependent variable is queries per second. We run that on Linux and FreeBSD and I only have new pictures for Linux. And I would say that the Yadifa people did tremendous work but we haven't slept and we also improved our performance.

So the red bar is Knot 1.0.6 which is our latest release; the yellow is the new release from Yadifa, then it's Knot .03, then NSD BIND and… I'm sorry, it was Yadifa and the old version, the release candidate and then BIND. The brown bar at the top is the Trafgen; it's the theoretical limit of the hardware. Here we don't have a new picture so I will just skip that.

The performance testing No. 2 is the same as this picture has shown and also we haven't slept here so I think we should definitely sit at LNnetLabs because they were too nice to invite us to sit together and work on the common testing framework.

Okay, so we also have packages. The Debian, Ubuntu from wheezy.sid, they will be in the repository directly. If you run Squeeze, you can use our repository even to same thing Quantal will have Knot DNS by default. Otherwise you can use PPA. Federal packages will be available shortly. I was promised that it will be in days or you can use packages we provide. For Free BSD there's Fresh Port of Knot DNS so you can install Knot DNS from Ports.

Other resources – the home page resides at know-DNS.cz where you can find more information and source code. We also started using Google+ page where we put small news about the development and interesting things about all the stuff we find interesting. There's also issue tracking in Redmine. I have recently changed the Redmine interface so it allows the [NMNLs] to report back so I hope it will not be overflown by spammers. And you can also find the development source code in the repository. We also have a mailing list. Obviously it's not DNS users. And I think that's all.

Jay Daley: Okay, thank you to Ondřej Surý from cz.nic. So for the round table session now I'm going to focus on a set of topics, introduce the topic and ask some questions about it and then give chances from the floor about that same topic.

So the first topic is going to be about advanced features. Now I think many of us are going to be in a position soon where we're running two, three or maybe four different implements of the name server from the four very capable offerings here and that's going to potentially present us with a problem of managing these four different implementations, all of which may do it differently.

Now we've already heard just from Knot DNS of their interest in supporting DNS, ECM and NSCP underneath it as a potential standard for managing name servers, so I'm interested to find out from the others what plans they may or may not have for this and how much they really care about us users who will struggle with it otherwise. Can I start with João up front?

João Damas: Okay, we participated in the initial efforts to define things like NSCP. I think the design is quite okay; we would be more than willing to support something like that - that's diagnostic of whatever the software that's running underneath it.

However when you go down to details, the particular choice of language that is used in NSCP to express the whole thing actually makes life harder rather than easier. It's based on this thing called [Yang] for which the simple specification without an explanation is about 250 pages long. From my point of view, it's unrealistic to call that an aid to assist an administrator. If that problem could be solved, we would be right there. Until that problem is solved, what we'll do is basically make a lightweight configuration mechanism that's hopefully reusable elsewhere.

In BIND9 and in BIND10 really, probably we will achieve these by putting some [shims] in front of the software in any case. So whatever the solution or solutions is up to the IDN and we can cope with them. But right now in the case of the NSCP, we have these major hurdles.

Jay Daley: Okay, would anyone else like to talk about that at all? No? Any questions about that?

Olaf Kolkman: In the sense that there is a standard solution to this problem – and I should say that we're looking at the developments in the

standardization we know, so to speak, but once something is available and standardized, then we are likely to follow the same approach of a shim between the existing name server control protocol that is proprietary in our case, proprietary in an open source way of course but is of our own design.

And we would facilitate for whatever comes there. It is the design goal to be one of the implementations in a cloud. And having open interfaces at the incoming configuration as well as the data plain is of high importance.

Male: Just a comment to this. I agree completely with what Olaf was saying. The problem I have with opened sometimes is that opened requires a little bit more than publishing. It requires things to be accessible. One of the things that we have seen with, for instance, BIND9 is that though everything is Open Source, the barrier to understand the codes is so high that very few people have actually been able to contribute. And that's the mistake I don't want to have done with the configuration language here the way we are having.

It's very nice to have this true and complete language that can describe every problem in the universe – that's not what the system is looking for. So it's kind of a plea to the idea – please come down to the earth and make something that's usable and not…

Olaf Kolkman: And it's a plea to operate us to join in that discussion.

Male: Yeah, maybe to add just… although it will be very tiresome I think, as I said our current line of thought is to do it on the DNS protocol itself. And again the whole Netcom for DNSCCM thing, we have exactly the same ideas as these guys. Ultimately we'll do some sort of a front end proxy, whatever you want to call it that would listen to whatever is the standard and then translate into whatever that we ultimately come up with.

Although on the other hand we have a hopeful feeling that when we talk amongst each other that we might come up with something that is a – if not *de facto* standards, the real standards about how name servers can be remotely configured and I feel… well, I have the same feeling as João that the current Netcom thing is a bit too complex to actually be workable I think for the normal operations people.

Jay Daley: We have a question on that.

Petr Kout: Yeah, Petr Kout. I happen to be the Co-Chair of that ITF Working Group that came up with the requirements document that you may or may not have been referring to here, so just to clarify and then a question. The proposal using Netcom's yang and being named DNSCCM – that's just a proposal; it isn't even worked on in the ITF. It is worked on by an individual or a small group with some external funding which is appreciated, but not yet under consideration.

The question I have is… I guess it was Olaf or João – both of you – inviting operators to join the discussion and I heard that there is some tendency against Yang system – a) have you taken into consideration RFC6168 and that is the collection of requirements for such a protocol when you designed your own interface; and 2) your hesitance to go the Yang way – or Yeng I should say probably, and I think I heard Peter lean towards a do it within the DNS, use it as a control plain as well. Does that reflect operator feedback or is that based on some implement or some preference? So, RFC6168 relevance and…

Male: Maybe me first? 6168 states rather clearly that that is a requirement for at least one protocol but that it doesn't say that it is only one protocol that should cover for everything. So the current line of thought that we had was whatever we can do within the DNS protocol, we would like to do that.

For instance, restarting the name server is still feasible in DNS, but starting it if it's not running will not be possible. So there are a few things that will never be possible unless you have a DNS proxy front end, but where are you going with that? So yes, 6168 requirements document what should be possible to be done is sure an input to our process.

Secondly, as we all, we are running name servers, so we have an operations department that we talk to that actually tells us, "Not yet another protocol; not yet is a thing that we need to open up on the firewalls need to configure if we can do without this." And the beauty of doing it within DNS here is exactly that. We have the secure channel

in place; it is only one; you don't need to open up the SSH port on a different port if I am to believe the Netcom thing where you have to configure the keys and things like that. So operations input, surely from our side, is there, yes.

Jay Daley: Okay, we'll move on then.

Olaf Kolkman: Petr is looking at me intensively I've got the feeling. I'll take the bite. And the bite essentially is, as I just said, we are following at somewhat arms length. We've got a control protocol that's been around for way before that RC in the 6000s series was published. And in fact, I am as a vendor, interested in hearing from operators in this area. So following it and committed to implement something whenever it's standardized.

Petr Kout: Well I should probably defend Yang a little bit because we also have one guy, (Inaudible) who is very active in Netcom Working Group and it's not that [bad] that the operators will not have to speak in the Yang language; it's just a matter of using the right libraries and there are some libraries. But as everybody – peers – sit down, we will closely follow up the development in the field and will implement what will end up as a starter solution for this.

Jay Daley: Okay, thank you. So let's move on. João mentioned in his presentation BIND10 adding in hooks for special processing which I haven't heard

from any of the others. I'd be curious to understand your views on that, whether ISC are doing something crazy or if there's a proper need for it. Olaf, would you like to go first on that?

Olaf Kolkman: With NSD we don't have that feature. It's also not on the roadmap. I should say that with Unbound, the recursive name server, we do have such a feature. So the simple answer is actually no, we don't have that in our roadmap. I could imagine that there might be features that people would like to do special processing, but again, we want to do simple things in a very simple and fast way and keep to that sort of core.

Other plug-in mechanisms… I could think of a number of things that people would need, for instance, the throttling under potential denial of service attacks or being used as a potential denial of service attack is something that is under hot debate over the last few weeks and I could see people wanting to implement that with a plug-in. But this is something that could also be done as a patch and those are the kinds of things that I can see of real use, of authoritative only name server. Plug-ins with respect to signing or maintenance of DNSSEC records are for us outside of the scope of the DNS functionality.

Jay Daley: Anyone like to add anything?

Male: Yes, I think three things. First of all, as I mentioned briefly, the general idea with our storage back end is that you generically can configure

whatever you want, whatever ideas that you configure, your sequel commands or your scripts that actually connect to the database and fetch data and forma to present it to the Master name server, that that is somehow…it fits into what João was saying.

Secondly, general IDs as a side effect that we have – or side products – that we have a set of libraries that you can easily use on your own software so you would have the database layer, you would have the DNS layer, they would have all sort of modules that you can read, purpose and reuse in your own software.

And thirdly, which is going a bit into what Olaf said again as well, is that we have been looking at how we would implement some sort of firewalling, rate limiting on a [per-query] basis depending on the number of parameters like source, IP address, query type and things like that.

But that as well we see mostly as a standard feature of the product that you can activate and use or not, but it's not necessarily something like a plug-in that you can leave out. It is actually the other way around where you use our building blocks to build a different product or another product.

Ondřej Surý: Well, I think with Knot DNS we are on the same page as Olaf and Peter that we will probably implement everything what is needed, just drop it into code so there will be no need for such things as plug-ins and it also is not compatible with our philosophy how to do that project.

Jay Daley: I have a question from the floor.

Peter Van Dyke: Hi, my name is Peter Van Dyke; I work at Netherlabs; we develop power DNS. We have hooks in various places and we find that they're mostly great for prototyping. What you don't want is people having lots of additional calls sitting inside of ETC being part of the functionality.

But for prototypes, these hooks are right and as Olaf suggested indeed, we are currently using our (inaudible) scripting host to prototype write limiting and it makes it really easy for people to experiment without having to learn CRC++. So even if you don't want your hooks or your extensibility in your core functionality, there are reasons to have it.

Ondřej Surý: Just maybe adding that one of the dreams of one of my guys is… well, more than a dream – idea – is to have a Perl library where we have actually the modules that we have built which are C but actually are available as Perl modules so you don't have to write C, although you would have to write Perl which some of them is better for some.

Jay Daley: Okay, one of the big trends at the moment is big data which – for those of you who don't know – is collecting every piece of data that you can, for example, a full packet capture off the wire and analyzing that to find what interesting things you can find in it, possibly keeping it for some years until you actually get round to looking at it.

I have a question for you which is whether you feel that your service should have any specific hooks in them – functionality or otherwise – for instrumentation and data collection in depth, or whether you think this is something that is best done by putting a network tack in place and have something listening to the traffic and doing it separately.

Male: Short answer – in-depth analysis – no, it's something that you do in within [that]; you actually don't do in your software.

Jay Daley: Anyone disagree? No? João?

João Damas: Bulk data analysis, no, I think that's best done somewhere else. What BIND10 is meant to have – and I don't know exactly when we are going to put it in as a configure – is introspection, being able to ask the server how particular queries process internally to find out how you arrived at the given result because sometimes it's not quite clear when you start combining zones and different behaviors in the server, it all is a bit of a mystery what goes on inside servers. But bulk access to data – no.

Jay Daley: Alright, any questions? Move on? Two of the implementations have mentioned ACLs in one form or another – we have Views and BIND – and Yadifa has ACLs as well. My question to the other two then – is this something you're interested in doing? Is this because you're aiming for a different target market?

Male: Clarification question. Do you mean with ACLs that you get different DNS answers based on the ACL or do you mean you get servers or not based on the ACL?

Jay Daley: I think we mean different answers.

Male: In that case, no.

Peter Janssen: And in that case no for Yadifa either at this moment. ACLs for us is typically a security measure where you… well, there is a bit of a difference because IXFR, AXFR you do [T6] and things like that, we are just part of the ACL so it's IP checking and things like that. But the answer is the answer for us for the moment it's not that the answer is different, depending on where you are coming from for the moment.

Olaf Kolkman: Same here. We use ACLs just for (inaudible) and transfers. We don't even have ACLs for refusing the query.

Jay Daley: Right. So I think we can be clear then that none of you will be attempting to compete in the market for monetizing NX domains or anything like that then? Is that true?

Olaf Kolkman: Yes.

Ondrej Filip: First of all we are a not-for-profit organization, so we're not competing at all. But no, the thing that we are putting out there is as I said, mostly a name server for TLDs, high performance, high grade name server.

Jay Daley: Okay, let's move on then. Yadifa, I think is noticeable compared to many of the other implementations, that it appears to have a marketing strategy behind it. It appears to have had a designer set up the pages and a website that is marketed, shall we say? Can you explain to us whether that's right, Peter, and what's the thinking behind that?

Peter Janssen: Mostly as I said, we had a .eu registry and one aspect which I will not deny is that if Yadifa becomes known in the world, it will be known as Yadifa, a .eu product and as such it's a sort of backhand way of attracting attention to .eu without any doubt the most sexy domain name that is out there in the world, so yes. We have a marketer and we have a designer that did the logo which we trademarked as well.

Jay Daley: Right, okay. Do any of the other of you have plans for this or plans for a particular target market at all to talk about? João, perhaps, what will be the plans for BIND10?

João Damas: Serve the internet – I think that's quite enough for us. That's why BIND has actually all these other features I told you about not choosing to implement. One thing you find out – when we were putting together BIND10 for instance, we went around asking people what were their main features that they wanted to see in the new name servers. For instance, the top one was views.

Of course, none of the ccTLDs care about views; they have one zone. That's what they want to show the world, but there are soon to be 1,000 of them – TLDs. But there are infinite amounts of enterprises that have this problem every day. So what we are trying to do is basically address all these many ways in which people want to use DNS and minimize the amount of breakage that comes from it because people can get very creative sometimes.

But as a specific market – no. We just want basically this to be a reference implementation so that there is a stability in the internet and that's quite enough.

Jay Daley: I think that it's fair enough that there was a time when there was one implementation that writing it and they will come would work. Do you think that in the current competitive environment where we have, say ANS from Nominum being marketed as a product and sold for extremely high price and having very little difference from anything any of the four of you produce – don't you think that's something that can continue or

do you think that you'll need to potentially copy Peter's more market-focused approach?

Male: Not necessarily. ASC has never been good at marketing; that's a well-known fact. What I have to say is once my intent is done, I will be selling my Nominum shares.

Jay Daley: Right, okay. So this brings us on to the future of your products and I think there are two parts to this question, one of which is the sustainability of it in terms of we trust that it will go on further and the other one then is your… so that's about your commitment to it; and the second is about your funding model for this and of how those will go forward. Olaf?

Olaf Kolkman: Yeah, I just wanted to respond to something that… I want to take a bite on something that João just said. We made the different choices and we did that with our authoritative name server, but we also made the conscious decision to have different products for different purposes. And yes, we don't do views, but we also decided to have a recursive name server or participate in the development of a recursive name server because I should notice that started as a collaborative product and have that develop as a reference implementation. And again also for the good of the great internet because that's sort of our vision.

Same goes for DNSSEC. We think that the maintenance and operational procedures of DNS key management are not necessarily part of the main server implementation, but is something that hooks more closely towards the back end systems, and still is a sort of generic feature of DNS provisioning and hence we incorporate it in open DNSSEC.

We do not have a marketing strategy as such but for us that openness of the internet is a driving factor and having high quality cheap – the cheap as possible – and free is pretty cheap – in at least one access – having high quality and cheap products that lower the price for deployment of technologies that will skill the internet and make it more secure – that for us is one of the driving factors. And we don't have a fancy webpage to demonstrate that but word of mouth counts for a lot.

Ondřej Surý: In that case I need to ask some words about it as well. We see the stuff we do at .cz labs as something we give back to the community and to the people of .cz and we give that back to all the internet community and that's our marketing strategy – to give back the trust and the resources we have in cz.nic.

Jay Daley: Thank you. So let's move onto this question of sustainability into the future. How do we as customers, shall we say, know that if we use one of your products it will be there for the next five years or 10 years? And to be blunt, I think all of us who are customers – given that all of your software is Open Source – are getting a great deal of value without necessarily paying you anything for it.

What type of funding model should we move to – should it be different for each of you? How would you like to see the funding of these things go and how does that tie into the sustainability? In other words, if we are paying you for it – cause that may make a greater guarantee that it will be here in five years time. Ondřej, would you like to start on that?

Ondřej Surý: That's a different question for us because first we have not decided yet what we will do about separate contracts. We probably will do separate contracts. And for the funding and sustainability, well, you know, anything can change in the internet but well, from right now, standing here, I can say that we want to support Knot DNS for the foreseeable future.

But of course, anything might change. We might get incorporated by government or something might happen. And in this case, it's still Open Source and I'm sure if not DNSL will be used among internet users and somebody will pick it up. But that's really about painting the nightmares about the government. Right now we will support .dns as long as cz.nic will exist. That's our plan.

Peter Janssen: Well, we built Yadifa to scratch a few itches that we had and by doing so, we actually are making a product that will be useful for a whole lot of people that are out there. By making it Open Source, people have the cheapest of possibilities to get a product that actually scratches their itch hopefully.

Yes, we are committed to the product for the very simple reason that we're actually running it for our own name server and infrastructure and we'll handle it internally so again, the same sort of answer that Ondřej gave us – as long as EURid is the .eu registry and again everything can change in the future, but as long as we are .eu registry, we will keep on supporting the product.

And as for the question for paying support, we haven't even thought about that. It might happen; it might not happen. It is something that is still some way in the future and we'll see where the markets will lead that question essentially.

Jay Daley: Olaf?

Olaf Kolkman: Yeah, this is a nice subject for discussion. The history of NLnetLabs is… I think I should explain it. The history of NLnetLabs is slightly integrated with the history of the first big commercial ISP in the Netherlands – that was NLnet. NLnet was an initiative at the start of essentially very enthusiastic computer scientists that thought that internet packet-based technology was a cool thing and they build out this incorporation which at some point they sold off.

And then the not-for-profit that they had set up had tens of millions of guilder at the time – Dutch guilder – which they had to use for their chartered purposes and chartered purposes were for the good of the internet. This is the LNnet Foundation and the LNnet Foundation is our mother – was our organization. They established us as one of the

projects for long-term funding in 2000. Now it is the case that by the end of 2015 they will not be able to commit to us the funding needed because the end of that pot is in sight, the end of that pot of money.

So we are actually a bit worried about our financial future. Fortunately, SIDN, the Dutch registry for .nl, had committed to supply some of the funding towards us and we are looking towards other ways of funding by holding up our hands for people who support the cost of LNnetLabs which has not only developed DNSSEC software but is for an open internet in general.

So this is something that we're looking into to see what type of parties are interested. Reaching out to specific parties that run our software and operations is difficult because we actually do not know who picks up our software and runs it. It's freely available on the internet; it is usually distributed *via* all the packages that the operating system vendors make like Ubuntu and Debian and so on and so forth, so we don't have a direct tie with our customers.

We do offer for TLDs and so on and so forth, we do offer support contracts; we do not advise them that [well]; we don't market them as such, but they are definitely available. Of course, the community support is paid by the generosity of the community. So I guess that paints the picture, Jay.

Jay Daley: Okay, João?

João Damas: We for sure don't have any sugar daddy or sustained income that's going to be there no matter what, so we have to resort to different funding modes for different activities and we're combining it. For instance, BIND10 is basically being sponsored by a set of people, most of them in fact, in this room who contribute actively to make the project actually be able to happen.

The normal functioning of ISC is then… the income that we get is divided amongst support activities, which is kind of a typical activity that Open Source companies devote themselves to. It provides us with a very stable and dependable source of income. We do for instance in the case of BIND9 sponsor development. If someone is looking for a given feature for whatever needs they have and particularly that feature was already on our road map, we talk to them and come to an agreement and therefore, cover the development expenses that way.

We are toying with other ideas to increase the funding we have to make it more stable, long-term, because ISC's idea is to be around forever. The DNS is quite critical to the whole internet and it has to work so we need to continue to provide this software for everyone to use and to prevent fragmentation.

Jay Daley: Okay, do we have any questions on the issues of funding? So if I can just make sure that we are all aware of this, every one of us in the room relies I'm sure on one of these four products and without them our business would not exist I think. So anybody who would like to disagree with that? No, I didn't think so.

There are two of the four here who don't need our money because they're registries and there are two that do. Now I'm certainly one that pays some money to one of them and we only have 480,000 domains, so we're not as large as many of the others here that could do that.

But it's clear that we have another one that needs money as well at some point in the future and I think it's beholden on all of us to work out how we can support these products because we need them if we are to survive. Lecture over. I shall move on to the next topic then. Oh yes, Antoin, please do.

Antoin Verschuren: Your question about funding raised this to me. Not all DNS vendors are on the table here, but only the ones that supply software that we use in our operational environment. I also want to stress that there are a lot of resolvers out there that query our name servers and they need to send good queries or otherwise our servers are going way down - which means that we sponsor these vendors that are on the table here; but we also sponsor vendors that are much more in use by our resolver and mostly the registrar community that usually use different software. That's something to think about.

Jay Daley: Okay thank you. But let's move on now to the importance of DNS servers. I think we're all aware that DNS is an important piece of infrastructure, but with DNSSEC coming into play, we now have new technology such as Dane being developed which is moving X.509

certificates into the DNS and this makes all your products potentially far more important than they were previously.

So how are you responding to this? For example, are any of you considering external audits, using tools like Coverity for code correctness checking or external certification or anything like that? Any of you like to pick up on this? João I think first? Oh, Ondřej first.

Ondřej Surý: We use such tools as start-up procedures so we control our (inaudible) by using such tools. There are even some free tools – the Cloud Compiler has something called [scan build] which also can scan for, well, do the static analysis to code and we consider this as standard procedures for releasing the code. Well, yes, but we will not do anything special.

João Damas: While it is again a multiple pronged, we have a huge amount of internal tests that we do to make us comfortable with the code itself. In the past you provided us with access to Coverity and we run BIND9 through Coverity tests regularly. Unfortunately things changed and that's no longer available to us. I would like to have that back if possible at some time. I think it's providing [some cases that we haven't talked about].

We added recently some first testing equipment by some commercial vendors to the lab to basically throw all kinds of unthinkable stuff at the server, see what happens – we're going to continue that. In the case of BIND10, we are talking to a software company that we think we can

reach an agreement with to have a full review of the code done. So it's important, yeah.

Jay Daley: Alright, the other two of you?

Olaf Kolkman: We don't have access to Coverity. I believe that the NSD code has gone through one or two of those tests but I'm not 100% sure, I should say. I do know that external parties have been looking very carefully with third-party reviews at NSD. Unfortunately, the results of those were shared in a non-disclosure so I can't… It's one of these things that you have to believe my blue eyes and I don't even have blue eyes, so take away what you have to.

At this moment we do lots of internal testing, obviously always internal code reviews. We try to maintain a high professional standard if it comes to testing and releasing our code. It is Open Source and I would appreciate that if people who have the resources to do that kind of test review or certification would share those results with us so that we can make those available and share them. I'm not afraid and I'm actually willing to cooperate with whoever wants to put us through that ordeal.

Peter Janssen: Repeating again but still, lots of internal testing; it is Open Source; the more eyeballs look at this, the happier these eyeballs potentially will be. But personally, as already mentioned, the temporary meeting… well, the meeting that we have temporarily set for the beginning of July with

the four people here or the four organizations here across the table – I personally believe that that might be something where we can work from to actually come up with what is necessary to actually have a name server implementation and what is necessary in terms of tests, both performance as well as conformance to actually come up with a creative set of tests that more or less proves the correctness of the implementation.

I don't know how. I have some ideas. I'm sure all of these people will have general ideas. If you put those all together, that we come up with something and we as you are more or less willing to sponsor some sort of a, I don't know, hardware setup that would be the reference setup or software implementation or I don't know what where we would get together, I don't know, two times a year where all of our name server implementations would be run through the mill and some sort of report potentially authored by a third-party – one of the big five – or whatever you want to call it, comes out, actually puts some stamp on it – this is what we have come to the conclusion of.

So yes, I'm very willing to move forward with that. But I think we first need to make a few steps on what exactly on the technical level is necessary to go forward there.

Jay Daley: Great. That was actually going to be my next question – whether we needed some form of cooperation to set up… no, it's very good… common tests in a common test environment and things. Is there something you would be looking for volunteers from the audience to help with and be involved with as well as the four vendors?

Peter Janssen: Well, sure, there's DNS benchmarking, (inaudible) at DNS OARC or the [SOP] so you can come join the discussion or read the archives or order some stuff from there and certainly, we would be happy to have more people onboard.

Male: And what I can say – the three others will say the same thing – testing costs an enormous amount of time to actually get it done. It's incredible the time that you waste on this to actually do some simple testing. So if indeed we could maximize this by having a conformance test that is actually one thing where you just run it through and that will give you your standard results, that would be great. And yes, if any person or any company or organization in the community is interested in this and has experience in this, or even if they don't have experience, and is willing to chip in, that would be much appreciated, yes.

Jay Daley: Wonderful. Anybody have any questions on this? Shall we move onto the next topic?

João Damas: One last word. Conformance testing is necessary but it's not sufficient. Testing for the things that you know should be happening is a lot easier than testing for the shit that's out there in the internet. So in that instance, anyone who has suffered from weird attacks, strange scenarios – all that information is the most valuable to be able to harden software like this cause going through the standards and replicating what should be happening is not that hard.

The hardest thing – I'm sure that we all have exposed our codes for the first time to the internet and found out that it's one thing to follow the RFCs and another one to cope with the internet.

Olaf Kolkman: And then I should add that with authoritative servers you're at the conservative part of what you send end of the spectrum. Recursive name servers is that to be literal with what you expect and there is a lot you can expect on the internet. It's scary. Yes, I agree. And having one set of tests for all of them as a sort of final check-up or some external audit is probably good, but that doesn't replace your individual tests because you don't want to be caught by the same omission in the test setup and then find that you still have the same bug over for implementation.

João Damas: Just maybe to add one more thing – as a community that actually is in the room here, we actually have access or should have access to a lot of shit out there on root servers, on TLD servers, on ISP name servers. I think that taken altogether will actually be a nice benchmark – just throw it at all the implementations and at least automatically see – are there different answers; does it fall over, which already will give you an interesting aspect of how are they coping with all the problematic packets coming in.

Olaf Kolkman: And the obvious place to point at now is DNS OARC.

Jay Daley: Okay, so that's a very useful point for us all then that we need to discover all of the bad things that we see happening to our name servers and tell DNS OARC about them so we can build reference tests.

Okay so the last topic then is about community engagement. This I think is potentially the first time all four of you have been up to talk about your name servers all at the same time and we're doing it with a good community here. And we've heard Olaf talk about not necessarily knowing who his customers are. If only he had a Facebook page everybody could like, that would do it I'm sure.

So let's talk then about how we…well, I have one question to start off with which is all of you who said that you take contributions to your code from other people and I'd like to know how many contributions you actually get. That is, something that is on offer but people don't very readily take up or if it's just the same people.

But then more generally about how we insure that there is ongoing community engagement that we know can deliver these kind of tests to you, can deliver the funding we've talked about, can keep us going as a community here – how to insure that carries on short of Peter taking everyone to dinner every night. Who would like to start then? Ondřej, you want to start?

Ondřej Surý: Well, to answer the question how many contributions we had – we already had some and what I did I think a few months ago, I said, "Well, if you want to work at .cz Labs, then go find a bug in our code," so it was a part of Edmonton's test for new candidates. So yeah, I had two

contributions from those people as well and I even hired one of them. And we also had some people from the community, some people [found bugs] in our code.

As we tell to the community, I'm not sure. I already said that we set up the Google+ page where we will try to keep posting interesting stuff, not only related to Knot DNS, but that's only one way. Well, we'll see when the community builds up around our DNS server.

Jay Daley: João?

João Damas: We realized that BIND is used by a lot of people out there, but most of these people are in what we now call the long tail – people who use it for their own purposes and very small installations and who cannot be expected to contribute, particularly financially. Some of them contribute with their time which is actually quite good, but not financially. So we're trying to see which ways we can find to get the big guys to contribute so that we can actually continue to address the needs of the rest of the population as well and it's not entirely easy.

Traditionally ISC has had a lot of sponsors in this community here – the ccTLDs – because mostly you tend to be the sort of not-for-profit organizations who, by nature of the business, tend to accumulate money even if you don't want to and also happen to have on other occasions a charter of giving back to the community.

But we need to diversity; we need to find out more ways of reaching out to other people who use. So we are talking, for instance, to appliance manufacturers which there are a lot of them – more and more every day – to the ISPs. Seems that the ISPs have not warmed up to these… well, actually I said the telco ISPs haven't warmed up to this idea of internet being something that's their business which is very strange. And you enter these bizarre situations where they are willing to pay commercial companies a lot of money and willing to sponsor Open Source for a small percentage of the commercial understanding.

And that's because of a problem of different language that's being spoken by the two types of organizations when we address these people. And we are trying to adjust to make people understand it's not easy but we'll keep trying.

Jay Daley: Olaf?

Olaf Kolkman: I'm thinking what I can add to that. As far as a community around NSD and Open DNSSEC, that definitely exists. We've got people contributing batches; we've got people asking questions on the mailing list and so on and so forth. How big that community is – hard to tell. I guess it's a tip of an iceberg or the high part of the long tail that we interact with. Again, it's hard to know who your customers are.

And with respect to the funding and getting back financially from some of our users, it's very hard. We do know that some of our products have ended up in [iron], have ended up with different vendors under

the hood. Some of those vendors have given contributions back in monetary as well as a scout and knowledge and know-how but that has not proven to be sustainable.

Hard to say how you can engage a community of… when you deliver something that is essentially plumbing. A DNS is not sexy. It compares to electricity wires, to water plumbing, to gas tubes and so on and so forth and those are not very sexy, they just work and tend to be forgotten at some point.

Jay Daley: Okay, before I ask the audience how you'd like to be engaged with, let's have Peter's words.

Peter Janssen: Well first of all, considering that the source code of Yadifa has been available since 10:00 this morning, have had exactly zero contributions from the outside world unless… let me finish – unless one of you in the room here has been very busy and actually sent something that nobody at EURid has cared to send me an SMS about. So we'll see about that.

As for the efforts that we have been doing – we're trying to warm up people in this community as well as the larger internet community in general I would say to at least look at Yadifa to see what it's worth, where it can take them, how it can help them and how it can make their jobs or their lives easier.

It's very preliminary but we're talking to universities to see how we can work with them in terms of research and code coverage tools and things

like that. Nothing for the moment that is specific to what will be done, but this is actually venues of attack I would say that we are looking at and we are very committed to this product. We want it to be as sane alternative to name servers as BIND has been through all these years, as NSD has been through these years and make sure that Yadifa will be through the years to come.

Choice is good. I think four is enough. I think – no? [chuckles] Unless the Chinese Open Source – their implement – we have five. That would be very interesting to look at their source codes. But apart from that, choice is good and we are very committed to this project.

Olaf Kolkman: You're forgetting power DNS. You guys should have been at this table.

Jay Daley: Is there anybody out there who would like to give us their views on how they'd like to engage with anybody who, say, runs an organization that very much relies on this? Robert.

Robert: I was busy logging into a webpage so I didn't pay attention to the last part.

Jay Daley: How would you like the implementers here of DNS servers to engage with people like you who use these on a daily basis – your infrastructure?

Robert: That's not very easy to answer. I guess it all sort of depends a little bit on what they expect, but I guess to follow up on what João said earlier, some of the things that they could get would be probably packet dumps, streams of traffic. In our case it would be a little bit problematic because of the NDAs we've been signing and promises we have given people. But I think that might be a good way to actually test what goes in and what goes out if that correlates to what is expected.

Jay Daley: Okay, thank you.

Olaf Kolkman: Robert, would you be – you as an organization – would you be willing to actually use your packet dumps that you can't hand out, but actually put out four or five or six and let's not get hung over a certain number – but put us through the mill and see how we react? Would you be willing to put effort in there, and I'm not saying the effort should all be on your side. We can work something out in terms of hardware, software, engineering, resources and things like that.

Robert: I think… well, apart from being a resource problem, we are definitely interested in helping how we can because it's… I think we all pretty much share the same idea of being of benefit and stability to the internet. So I think we should talk about the possibilities because we do have researchers accessing some data and I'm sure we can figure out some way of generating something. I don't know how. I mean we do not have a lot of extra manpower available but let's talk.

Jay Daley: Okay, thank you. So let's finish off this session then. We've heard from four implementers of DNS servers that all of us rely on, all of whom are doing this for the greater good of the internet. - each one of them has made that very clear; each of whom are very dedicated to this part of a team that is very dedicated and that takes this very seriously and I think this is one of the few industries we could be so privileged to have that kind of thing. So can you please give a round of applause for our four contributors.

Eberhard Lisse: I must of course admit that this oversight is not entirely as important as is purely due to my own ignorance and I've already asked him if he's going to be on one of the next Tech Days and he will be of course, be invited to speak there. Now we have just a short feedback on our Masterclass from yesterday and then Luis will terminate the proceedings.

You've got five and a half minutes and then Luis got three and a half minutes because we need to be out of here at exactly 5:00 because the room is booked for the next participants.

Jaromir Talir: Okay, hello, I will try to be as quick as possible. My name is Jaromir Talir. I also work for .cz and I was asked to give you a quick brief overview of yesterday's FRED Masterclass we had in our offices. You have maybe heard today several times that FRED is our Open Source domain registry currently used in fixed countries around the world and

this master class was just a realization of the idea that we with Eberhard had during class the ICANN meeting in Costa Rica.

The idea was just to bring together users and developers of FRED because such a direct face-to-face meeting is better than hundreds of emails. And also because Costa Rica is a country that is using FRED as a registry and the Czech Republic as the next host is using the same system there was also an argument to do this Master class right here in our country. And maybe that's a surprise that Jacques from .ca was not yesterday on our Master class because .ca is hosting the next ICANN meeting and this around all countries using FRED must be followed by some country using FRED.

So you have four months to implement our system for .ca so you have time. This is just a review of participating organizations. There were three countries that are using FRED – Faroe Islands, Angola and Costa Rica. Several countries that are thinking to use FRED – (Inaudible) and Mali, and also some people from countries that are right now using [CoCo] Tools as our main competitor. But maybe after this Master class they will decide to change their opinion – I don't know.

What we did – we just walked through all our features and technologies of FRED, so we described our components, our possibilities for customization, I did some live demonstration of installation procedure; some basic administration tasks and registration procedures – how to register domains within FRED.

And of course the main goal was to exchange experience – either within the countries that are using FRED or we have got some information how

[CoCo] Tools are doing something and what is the main difference between these two Open Source registries.

And the result – attendees were satisfied – that's what they've told me – I don't know – problem we realized that some people couldn't attend because the system is quite heavily thought to be used in Africa countries. Some people were really trying to ask for remote participation which I wasn't able to do. And of course another problem – some people had a problem with visas to get here so either we know that there is a much bigger demand than we saw.

And of course we've got a lot of feedback about some things that should be done differently and what should be more documented and what should we change. So the main result is such a Master class should be repeated sometime in the future.

Maybe some of you have already seen our nice visualization of FRED deployment around the world but maybe you have also seen that this color of the dot maybe says that FRED is deployed much more than we expected. So that's all from me. Thank you for your attention.

Eberhard Lisse: Thank you very much. It was indeed very interesting. As you all know, we use… .na uses CoCo Tools. I even picked up a few things that we can feed into CoCo Tools to make it even better so we don't have to really strongly consider switching just yet. Just a small point – FRED is not used by .ao; it's used by .com.ao. There's a slight difference. That's an organization that is outside of Angola. If you want to register within

Angola it's very difficult to get hold of them but now at least we know how we can register.

My clients for example in Namibia sometimes want to register in Angola and we at least now know how to get hold of .ao. So that was at least something positive.

Luis Diego Espinoza: Well, to me this is too short time to review all the presentations, but it's not my point. First of all I want to thank you to stay here and I want to thank you to stay here, and I want to motivate the people here and motivate the people is here to keep support in this Tech Day. The Tech Day for me is one of the most important meetings during these ICANN meetings because I have a lot of people in these workshops.

Today was very useful – many technologies, some very good presentations. For example, from Richard Lamb – it's a very good presentation about this innovation that this is a demonstration of how can this working group, how can this Tech Day can [support] to the community. For example, this panel just passed, was a very good, very informational and very helpful to decide or to evaluate different products having the people that has the sign, the development of each one of these products.

Now only thank you and have a nice ICANN meeting the rest of today and again, please support this Tech Day and being here, asking questions, doing presentations, thank you.

[End of Transcript]