Wayne MacLaurin: Good morning folks. We're going to get going. I'm sure we'll have more stragglers come in as we spool up. I'm Wayne MacLaurin. I'm the Executive Director of DNS-OARC and I'd like to thank Eva Harte here and the other folks from the ccNSO Tech Day for allowing us to do this joint meeting.

I think yesterday went pretty well; we had some great presentations. For those of you who don't know who we are, I've got a couple slides just to sort of walk through and show you what we do and what not. For those of you who have probably seen this about a half-dozen times now, I apologize, but it's very short, and then we'll get on to the meat of the presentation.

We have just a little over 60 members right now. We were formed about seven years ago in an effort to provide a sort of neutral platform for primarily – originally anyways – the big root operators share data and do sort of some common analysis on DNS. It's grown since then quite considerably, slower than some would like, but certainly steadily anyways, and as you can see, we count a number of the big players as members and we're certainly always looking for more.

We have sort of three pieces of our mandate – operations, analysis and research. Data collections is probably the thing we're best known for. We've got about just over six years of continuous statistics coming out of one of our programs called DSC which we've been collecting from the root and major TLD servers.

We have almost six years now because we're about to do another collection of what's known as "Day in the Life." It was actually one of the original reasons we got spun up was to do this in conjunction with K.C. and the University of California to do a study on what a typical day on the internet looks like from a DNS perspective. So in the course of six years we've collected just over 50 terabytes of compressed data which we have available to all our members to do analysis and some regression stuff.

The other thing we do is we maintain a bunch of lists. Do have a Do Not Probe list of sites of operators who really don't like people exploring what they're doing from the DNS perspective; open recursive resolvers; the TLD Zone File Repository and 10 years of root zone archives.

On the tools side, we have a number of tools that are widely used. UDP source; port randomization; DNSSEC testing; Lookaside validation and a number of others that people have used over the years and created for us.

On the communication side of things we run three major mailing lists that people use actively. Our DNS Operations mailing list has a little over 1300 members, quite active – it goes in waves, depending on what the problem of the day is. We have a jabber room which is open to members – currently have 182 active accounts in that, which some people are actually using today as part of this meeting.

And then we have Secure Ops which is a closed vetted real-time mailing list and jabber session for the folks at Operations allowing them to actually talk about real problems on the net without worrying about the politics and what not of who's involved - quite active and also quite popular.

On the research and analysis side, we're known for a few things, but most recently we did a study for ICANN on the root signing and what the effects would be. So over the course of the entire which was six months, we collected data, studied the various steps in the rollout and produced a report which is linked on there. Most of you have seen these graphs in one form or another at various presentations and talk about things like the change in EDP size and our famous 512-byte problem.

The DITL Reports, again, which we do once a year, last year we picked up about just shy of two terabytes of compressed data. This year we're scheduled for April 12-14 and we expect to pull in probably the same or a little bit more again this year. If people are interested in that, there is a mailing list. If you contact me, I can get you onto it and explain how things get set up and done and contributed and get access to data and that sort of thing.

So for 2011 and beyond, we've got a couple of initiatives this year that we're looking at - obviously the DITL this spring. We're looking at doing some work around the IPV6 day, data collections and analysis around what that's meaning to the internet.

Starting to look at some real-time analysis again using a combination of the data we've collected plus our access to some of the real-time streams of data out there. It seems like kind of neat and interesting things we can do there.

And then on the tools side we've got an initiative underway to rewrite and expand DSC, something that was started doing whistles, started back in October for us at our fall meeting and we're just getting ready to actually start the development on that side.

So if anybody has any questions, that's how you get a hold of me. I'm here for most of the week. You can track me down or grab me in this session. For those who aren't members, I would dearly like to have you guys on as members, especially the ccTLD community. We see a number of you here but there's certainly a much bigger community. We think we can offer some good stuff. So with that we're going to turn over the presentation to Eberhard and enjoy the show.

Eberhard Lisse: Yes, thank you very much. Good morning everybody. As usual we are having our tech day on Monday, but this time as you heard, we were approached – I don't actually know who approached us – but we were approached somehow so we decided to have the OARC meeting which was supposed to happen at this point in time at the same venue. We have the venue available so we thought we'd do this in a two-day thing where on Monday we do a little bit more of the deeper stuff and on Tuesday we do some stuff which is maybe not so deep so that I even can understand it.

And the first topic, of course, is one of OARC topics so I don't want to detract of anything or anything. It's Vincent Levigneron from AFNIC and he has got something to say about DNSSEC key deletion issues.

Vincent Levigneron: Good morning everybody. Yes, this is yet another DNSSEC presentation. There are two members of AFNIC in the room – myself and the well-known Stephan Bortzmeyer. I guess almost everybody knows him. So if you have questions during the week, I have to leave tomorrow but Stephan is still there until Friday, I guess, so feel free to ask him or me if you have a question about this presentation.

What I am going to present you today – it's a key deletion issue we had two times in fact on our signing infrastructure, and (inaudible) more is the plan of the presentation I'm going to give you some key numbers to understand what DNSSEC segment publication process means at AFNIC, and some more close information about specific ethnics, DNSSEC specification, and I tried to show you what we found in Bind with specific and private recount issue. And I have a surprise for the end of this presentation, so stay tuned until the end.

AFNIC is a registry and we operate six ccTLDs – fr, re, pm, tf, wf and yt). Of course, the most important one is fr. Each zone is signed and DNSSEC was introduced in late 2010. Each zone signing key is rolled over every two months and we have chosen to use NSEC3 and opt-out options.

Of course fr zone is the largest one with nearly 2 million domain names and fr zone contains more than 4 million resource records, which is not a huge zone, but it's a big zone. We have no DS records yet; registration of DS should be launched in one month. We have chosen to implement EPP/RFC5910. Implementation is almost finished and we are doing some tests so we should be on time to open this registration to our registrars.

More specifically about DNSSEC, we use OpenDNSSEC for Key Management. We also use AEP Keyper HSM for Key storage mainly.

We use them at the moment in the program, but it has been updated since then. We used Bind 9.7.1-P2 with auto-DNSSEC option set and Bind of course, do all the signature stuff with HSM. And as I told you it has been updated since.

And we also have homemade synchronization script which is used to create Bind key files from ODS data. I'm not sure it's a very common configuration to use a dynamic update – HSM open DNSSEC, but if some of you have the same configuration, we would be very happy to talk with them.

So we see a simplified view of publication system and we have HSM and four servers. The one on your right, I guess it's your right, yes, are called nspublishers and the ones all dynamic edit scripts it hosts key management to them. It's open DNSSEC as I have just told you. It runs contra scripts and is mainly used for Zend Framework generation… It runs script for Zend Framework complete generation.

We also have two internal name servers which run ND. The first one called nsccheck is many uses for testing purposes, but we have no specific DNS check in this server. This architecture was created before DNSSEC in fact, and has many uses for test when we do complete ZF generations. When everything is correct, the data is sent to a zone end server called nsservers and sent nsservers and send notification to hidden primary name server. This is more (inaudible).

While there are more than 4 million records in the fr zone, there are only 17 NSEC3 records and no more than 35 signature records. This is because we use, of course, NSEC3 on the (inaudible). At the moment there is only two Key-Signing Keys – one published and the other is active. We have two or three Zone-Signing Keys at the time. One is published, ready to be used and it's a pre-published key. One is active and used to sign records and if we are just after a key rollover, there is sometimes another key which is a previous active key and is still published while inactive.

Our zone is dynamically updated every hour but once a week there is a complete zonefile generation for it's mainly for administrative purposes. Dynamic Updates is not used for all types of records; it's only used for delegations – NS, A and AAAA. All key and signature stuff is only based on automatic signing Bind capabilities and we don't use Dynamic Updates in this case. We don't send the DNS key with Dynamic Updates.

If you use the very well-known tool that is Viz, this is what you can see if you check for our zone. This is fr zone but it's just something on every (inaudible). It's a common situation with two key-signing keys and two zone-signing keys.

What happens more especially when there is a key deletion? We have decided to use very large timings and when a key becomes inactive, it is deleted only one month later. When a key is deleted, we purge and archive key files three days later. It was just one hour during the first outage we had in November, but is has been increased after that event. I will explain later why. When a key is about to be deleted, we are sure there are no signature records left corresponding to this key.

This is a little focus on a private Bind record and some of you perhaps don't know this record, but I'm sure that people are doing DNSSEC in this room and Dynamic Updates I've already heard about it. This record is described in administration and reference manual in Section 4.9.4 and it's a five-bytes record. Its type is TYPE 65534 and is used to know the state of a signing process.

This is what we can read in the manual: "If the first octet is non zero then the record indicates that the zone needs to be signed with the key matching the record, or that all signatures that match the record should be removed." This is very important to remind that. And in this case, when the first octet is non zero, the final octet indicates when signing is complete on it, of course, non zero value if it's finished. This record causes us some troubles in what we called "TYPE65534 Bug" but there is a more official reference and it has been corrected since.

We had the first DNSSEC outage in November. What happened exactly? During key deletion we had a network issue making our HSM unreachables. This is bad. The error was not well detected by our system and, in fact, this was the first time we met this situation so the publication process didn't stop as expected. Zone was not updated and the key with "delete" status was still present, while inactive, but it was still present in the zone.

OpenDNSSEC to Bind synchronization process which is a homemade script, decided to purge the key files one hour after it was supposedly deleted. Then, of course, Bind couldn't process Dynamic Updates because it was impossible to use key files.

And, of course, each element, if you take them separately, seems obvious, but when all happens at the same time, and you have to find a solution very fast because your zone is broken, it's a big mess and it's a lot of confusion to solve all the problems in once. And, yes, of course unfortunately, we also had the TYPE65534 Bind Bug that we are about to describe but we were so focused on the other parts of the system and the program, we just discovered that just two months later.

The biggest outage we had was in February, one month ago. I'm going to give you a detailed view on what happened exactly. The situation just before the key is deleted from the zone fr, the zone with referring (inaudible) ending with 7. The zone signing key with keytag 43893 is still in the DNSKEY RRset, but, as I explained, there are no more signatures records generated with that key. So this is the only reference we have about this key in the zone.

There are no TYPE65534 records in the zone because there has been a complete zonefile generation a few days before and no key operation since this time, so this is a complete normal situation. Then it's now time for key deletion. Deletion is triggered by Open DNSSEC and some creation is done with a mdc script. So our script, a (inaudible) script is identified and execute in our mdc signed zone. This is what we do for each key state transition and it works. It should work.

But what was expected from our point of view when this key switched from inactive state to delete state in zonefile with the next serial. So DNSKEY with this regard corresponding to keytak 43893 needed to be removed and it was done; it's correct.

DNSKEY Resource Record Set signature needs to be updated and it was so, and serial should be incremented; it has been incremented so it's perfect. SOA signature had to be updated and it was, so everything should be okay. And of course, these four steps should happen in the blink of an eye.

We just forgot something. We just forgot private type that appear in the zone. In fact, this is what it looks like. You need some translation to understand what 08AB750100 means. But what does it mean? It means that signing process with key 43893 is not finished. But finished to do what? Because there is no keys, there is nothing other – what is there? And there was no signature for this record. This is not a very big problem, this lack of signature, but it's not correct.

The problem is the Typemap of the NSEC3 resource record corresponding to the Apex should have been modified. This is not the case. This is not correct, of course, but this doesn't prevent the validating resolvers to validate. There is just one type missing in the Typemap.

But the big issues appears now with the next serial when Bind signs the private record and add TYPE65534 to the Typemap of the previous NSEC3 resource record, but without updating the signature which makes it invalid. At this moment, the fr zone became inaccessible to any validating resolvers. This was a big issue we had.

If we had waited for several hours later, we would have had a new revision of the zone without private record; with a NSEC3 with a Typemap which is correct; with a new signature for this NSEC3. But, of course, we couldn't wait but this behavior was confirmed on our lab testbed.

What we noticed is that during this long period, while there are no visible changes in the zone, Bind does lots of RBT calculations. When this task is over, the zone is valid again. So that's why we didn't notice anything for the other zones because they are very small and this RBT calculation takes seconds, and for fr zone it takes hours.

First of all, we would like to thank you because our monitoring system failed and it's Murphy's Law, of course, and something failed and everything failed. We didn't check NSEC3 resource record and first alerts came from you.

ISC provided a patch very fast and our tests showed us this bug doesn't exist any longer with this patch. That's what we thought, in fact. We also found a bug in Unbound and the patch should be published soon, I guess around the people from (inaudible).

We also had good feedbacks on our search for a zone verification tool. IDNS, for instance, is very promising, while not fast enough. And it was not very easy to file a tool able to deal with a very big zone.

And what happens next is that there will be other issues, unfortunately. Patched version of Bind 9.7.3 had been deployed and it works. We are also modifying our system to have better control over zone changes. The zone is now validated before from DNSSEC point of view before it is sent to our hidden master. It's not yet completely finished but should be operational before the next fr key rollover deletion phase in less than one month. And we had to introduce a new proxy server that controls notifications that goes to our hidden nameserver.

This is the same slide I showed you earlier and this is the same with modification. You cannot (inaudible) a new server that catch notify send from [NS Master] then do some DNSSEC variation and if everything is correct, all those hidden primary server to do some (inaudible).

What have we learned about the program? DNSSEC is still young yet; it works but it's still young. Teams training is essential and you still need a DNSSEC specialist – they are still mandatory when problems occur. We found a few field-proven tools available and zone size is often a problem.

You should keep all zonefile revisions – it would have been impossible to find the bug without that. It's not very of use with the combination we have with Dynamic Update and Automatic Signing, but it's (inaudible). Hopefully we deployed a zone versioning system a few times before the issue and we just missed a version of fr zone. And of course, you need to monitor and monitor and monitor again.

Excuse me – just so we wasn't missing something important. So we should be on the last slide of my presentation, but I add two new slides and a bonus track, so you get three for the price of two. And first of all, I like to excuse me if we were not focused on your very interesting presentations yesterday, because AFNIC decided to do something very special for this meeting and we had a third DNSSEC contention. We are still analyzing it but it seems we have found a new combo. It's too early to give conclusions, but these are some details of the problem.

So you take this with precautions because it's not completely understood, but this is what we can say. In the morning there were key states transitions on zone fr. It worked well. Following Dynamic Updates were well processed. Then Bind decided to modify its private records, again, this one. But at the same time, it seems – it's not completely sure – but it seems we had HSM reachability issue again so the published zone was not correct.

A new TYPE65534 record has been added to the corresponding RRset; SOA resource record has been modified. But two signatures were missing – the one for the TYPE65534 and the one for the SOA resource record. There is already a patch, but not yet applied and we need to check if it fixes the problem, but thank you to ISC for this patch.

Is RRSIG missing for SOA a big problem? Answer is no, it's bad but it could be worse. And, of course, it was worse. We have two NSD nameservers amongst our slaves and in this very special case, NSD did something unexpected and the behavior is the same for the two nameservers. And it seems it also decided to remove all other signatures of the Apex, which is very, very bad.

So if we wait one hour perhaps I could add new slides about NS zone and DNSSEC contention but I wouldn't wait too much. Okay, it's over. You have questions?

Eberhard Lisse: You of course most of all come to upload and updated version in the internal works?

Vincent Levigneron: Yes, I did.

Eberhard Lisse: Okay, thank you very much. This was a very interesting thing, even from our small little zone with 2,500 names and where I do a home-grown script, so I'm not that worried about anymore and everybody can forget to check every line.

My script talks to me a lot and I have a habit which I take very seriously. If I get a message it can't talk to the system, I actually wake up and do something about it. But there are only, as I said, 2,500 names; it's not four million or something. Anyway, any questions? If I wasn't in public, I would swear at you now.

Paul Hoffman: This is Paul Hoffman. The one thing that I found most interesting in your presentation was not, "Oh, this failed," or even "Oh, this failed yesterday," - it was the lack of checking tools. And I saw Stephan's message on the mailing list – it was about two weeks ago that you started asking and got what was clear to me to be insufficient answers, especially because some people were saying, "Just use x,' as if you hadn't known about x. And then you would say, "doesn't work," for this reason or that reason. And then it was like wash, rinse, repeat - people would just say, "Just use this."

It seems to me that – especially because France is what I would call a mid-size zone, not a huge zone, as you say, but certainly not a small one, not like yours – that the lack of tools, even for mid-size, checking tools for mid-size zones, where you could simply say, "Do a key rollover;

wait for 10 minutes until the tool says, 'Yes, that's good'" – the lack of that seems to be a dangerous thing for mid-size zones.

The larger zones we assume they are either making their own tools or people have their own cyanide kits. But the large numbers mid-size zones, I think is – what we've seen so far – is underserved by the checking tools. Do you agree with that? Or in the last couple of weeks since we got sort of an insufficient response for you, what have you done about that?

Stephan Bortzmeyer: About checking tools, yes, it's a real problem. We found a means to reduce the size of the zone to use existing tools, but it's not possible for someone to validate the (inaudible) on file, and we're in touch with a guy from [LLNET], and I can't say something for them but I'm sure they are working on something that could validate big zones and very fast. But I have no information about the date or something like that. It's too early to say it would be very soon.

But we read the code and we found some bottlenecks and I'm pretty sure it's not very difficult to improve it in a short time, but perhaps too short for us, but perhaps it will be okay in several months, I guess.

Eberhard Lisse: Well, I wouldn't go as far as cyanide, but maybe some Valium will help.

Keith Mitchell: Hi, Keith Mitchell, ISC. Not so much a question, really just to say thank you to Vincent and his colleagues for doing all this work and keeping us informed as far as Bind 9 is concerned. And I'm sorry you've had these problems - we're working on them.

Eberhard Lisse: Alright, good. Next one will be Barry Greene. Oh, I'm sorry. In the meantime he can go and set up. That will even be easy. And I get more exercise.

Animesh Chowdhury: Hi. This is Animesh from Neustar. We had found exactly the same issue of dealing with two million record zones, NSSEC, loss of signatures, but what I was able to do was work with DNS Java, to modify the DNS Java, pre-modify the DNS Java, build new DNS Java which will take a little larger (inaudible), like three million, four million records and as I look at this too, this thing will be done in like 10 minutes. So maybe we can talk offline and I could give you my tool and it was okay for my mid-size zones.

Barry Greene: So I was asked – my name is Barry Greene – I'm currently now President of ISC. That happened a couple of months ago if people didn't catch the news. What's interesting is even though ISC is heavy, of course, into DNS and things like that; I'm actually a network and security guy.

So parallel to this is another security track, and like things happen from the rest of the week, we have lots of overload and people have to jump between different rooms. So I was asked to do some session on security and people suggested do a ccTLD security.

What's ironic about this is the people who need to hear what I'm about to say are probably not interested in coming into this room and if you're in this room, you've already done 90% of the thing you need to do to actually run a healthy registry and registrar. So there's an irony to this.

So as I go through there and say, "How do I bring some materials to everybody," what I'm doing is what I've done in other communities, other operations communities where we'll take a topic and all of you who go out there and talk to our peers who need to do better, need some tools, so these are slides that we will put up on both the PDF and the raw source. So if you want the PowerPoint or the Keynote, you can take them and you can rip these slides and use them in your presentations as you go out there and talk to people and say, “You need to do the right thing.” So that’s the context of them. So you’re going to have both of those, so we’ll have some tools available to us.

So on here, this is a quick agenda. You’re going to hear things that you probably haven’t heard before, but you’re probably familiar with. So this is kind of like out of the box; this isn’t your traditional “This is another ccTLD – here’s what you need to do.” We’ve actually had a lot of those.

We’ve actually had a lot of good presentations – ICANN Outreach and things like that. This is not those materials; we don’t need to cover those again. What I’m going to give you is the context of why the operations and security community, banks and financial community is so upset. What is the drivers going on in the internet around us?

Many of you are familiar with this because you get the phone calls from people say, “Hey, we need to clean things up.” But what are some of the driving factors behind it that you may not be familiar with?

So, let’s go through a Criminal Toolkit. Let’s make sure everybody’s out there. Criminals don’t go out there and take one computer and hack into things anymore. You have a whole network array. And some people say you have a cloud of cyber security assets and all of them are pretty much stolen. Now DNS is part of the infrastructure that the criminals put together with this.

So here you got a guy who wakes up in the morning, has a cup of coffee; he fires up and he has a day job. His day job is to go break into computers and build a botnet. So he goes out there and first thing he does is goes, gets a domain – probably uses a stolen credit card to get his domain, things like this, right, and he's going to use it right away.

So in here he's going to prepare his drive-by system, he's going to go out there and stage malware. The malware is stuff that probably anti-virus can't check. He's going to do a spam run out there. The spam run is going to attract people who do this. This is a particular picture using a spam run in China at one time that says, "Hey, come look at this Taiwanese music star. A new music came out," and lots of people went to it.

And when then went to that site, what happened is they go to that particular site and they get infected. And it goes right through your anti-virus system; goes through your firewalls; goes through your IEP because one of the unspoken things that's happened in the security industry today is a lot of what we say is protection – your assets aren't protecting your assets anymore. The bad guys have figured out a way around that.

A lot of these tools, once they're downloaded, the first thing they do is they deactivate your updates so either other criminals or the anti-virus community can't load updates because they now own that asset. It could be your computer in front of you right now.

And, of course, what happens with this is they start connecting up to their cloud system of different assets around the world; connect to a controller. The guy's sitting there having his coffee he goes like, "Ah, look at these hundreds of computers who are connected back to me," and what do they do? They give him a nice little gooey interface that makes many network engineers jealous because they say, "I wish I had that to run my network."

Yet the criminals have these nice gooey tools that tells exactly what the computer can do. And then the botnet herder starts farming out and saying who's going to do what – what am I going to sell it to? I'm going to sell it to spammers, I'm going to sell it to advanced persistent threat people, all different people out there.

So this is the environment out there. So what can you as an organization – like if you were an enterprise – if you were a bank, if you were a factory, if you were a service provider – what do you do? What do you do because you gotta go out there and say that the problem is that guy who is doing this – we cannot lock him up.

People say, “Oh, let’s go lock up the criminal.” The problem is we can’t lock him up and there’s a particular reason why we can’t do this and it’s going to be years before we are in a position in global society to lock him up. And this goes into the second phase of why cyber crime is institutionalized. It’s going to be here for a long time and we in this room can’t do much about it right now. We are victims of it and we have imposed OPEX cost against us to handle this problem out there because of this.

This is the traditional view of the world. We have borders; we have laws; we have law enforcement. In this country, right here in the United States, next door we have FBI agents. They’ll come in here and they have laws and they can put handcuffs on you and things like that.

The problem is, in cybercrime, we can’t do that. A lot of time we know exactly who’s doing the stuff. We get pictures of them; they brag about it; they go out there and post pictures of all the money they’re making every day. Some of these gangs are… One gang in a six-month period – you’ll know how bad it is. One gang in a six month period, using one Zeus malware system cleared $50 million. $50 million – that’s one gang. This is the behind-the-scenes – it’s not getting out in the press how much money is being involved with the cyber gangs.

So this is what the traditional world is; this is the cyber world. It's all come together. We don't have an international legal system to go out there and lock somebody up. It's all glued together. Now what we can do is we can go out there and figure out how to go after different threats. We can look at behavior patterns; we can look at how these guys act; we can kind of disrupt what they do.

So in my definition of security, I break the thread up into three different areas where you have the cybercriminal; you got the terrorists and nation state and you got those patriotic, passionate radical people like the Anonymous Crew and things like that. That's category No. 3.

The thing I get more scared about is No. 3. The reason is No. 1, which I'll go through right now – the cybercriminals. Cybercriminals are really well understood because with the cybercriminals, there are key principles in which they operate, right out of the criminology books. When you go and learn about crime and the behaviors of crime, that's right out of the books.

So these are seven different principles that we've observed over the years on cybercriminals. So anybody ever hear of Steven Covey's "7 Habits of Highly Effective People"? This is Seven Habits of Cybercriminals. Follow your principles and you will be successful on the net.

Alright, so in here, Principle No. 1 – don't get caught. They'll do all sorts of things to not get caught. If you are a good registry or registrar and you've got your act clean - you're doing the logging and things like that; you cooperate with law enforcement and service stars around the world – anti-spam institutions – the cybercriminal is stupid to try to touch your property because you'll have all the telemetry there to help them get caught, so they won't touch you. They'll go to somebody else who's not doing it so well. So they'll stay away from assets that are being well-monitored because they don't want to get caught.

Second thing is they don't want to work too hard. They don't want to work too hard. For instance, everybody talks about Kaminski DNS Poisoning. And I was one of the guys during the middle of it – I worked at Juniper Networks at the time – I'm going like, "This isn't a big thing. It's way too noisy." DNS Poisoning is way too noisy. You look at it in back scatter, you see it; if you're a network engineer, you see it; you try to poison, you see it.

A bad guy is not going to use DNS Poisoning cause it's way too noisy. Everybody in the world can see the DNS Poisoning going on because of all these packets flying all over, backscatter. So if you do any sort of backscatter telemetry system, you see it.

It's far easier for a criminal to go out there and violate a registry and registrar – go after a ccTLD. That's easier; less chance of getting caught. So you contract the behavior pattern of what's more likely a risk. The likely risk is go to the top versus trying to go onto resolvers out there.

Next thing they do is they follow the money. It's funny – several years ago Big telco – seven Senior Vice Presidents sitting in front of me. I'm given a special EVC session and I said, "Look, this thing – we don't have a cyber security issue. Nobody's trying to break into you." I said, "The reason nobody's trying to break into you is you're a big telco." They go, "What do you mean?" "Okay, here is an exercise. All seven of you. You guys are big guys and so you're a telco. I know how to take down your network. My specialty is service [virus] security. I can rip your network apart. You have eight hours to put $3 million into this Swiss bank account. Don't worry about the threat. Don't worry about that. Your job is to figure out how do you get money out of your organization over to that Swiss bank account. What do you need to do inside your organization to get the money out?"

And they're sitting there and they just started to toss a discussion for 15 minutes and I stopped the discussion. "See, that's the point. You're a big bureaucracy. People don't know how to get money out of you. You're not targeted because the cybercriminals don't know how to extort you. Once they figure out how to extort you, you will be extorted."

And that's a key principle on the net. They follow the money. If they can get money out of you, you're a victim. If they can't get money out of you, they'll go somewhere else. There's lots of people to victimize on the net.

Principle 4 – if you can't take out the target, take out everything upstream to it. So what you're seeing sitting there is where somebody was actually paid to go out there and take out an organization on the net because they're getting ready to do their announcement for their quarterly results. And the website's getting attacked. And the website owner knew they were going to get attacked and so they're well prepared.

So what'd they do? The bad guys went after the service provider, knocked out the PE router; over 1,000 customers went down. They'll work it upstream because they're getting paid to do a hit. It's like a criminal says, "I'm getting paid to break your leg. I'm going to break your leg and you're going to pay this out. Well, if I can't break your leg, then I'll go and do something else. I'll blow up your car or something."

You got collateral damage; these guys will do that. So they'll move upstream with it. They do everything internationally, so it's really hard to put together a case when I got evidence in France and Italy and Turkey and China and the United States and South America, Latin America, say, in Columbia. And I got different legal systems throughout it. And how do I get the evidence all lined up to take the guy to court? These guys know that – that's why they do it all internationally, they set it up that way.

They'll go after people who won't prosecute cause if you go out there and say, "I'm going to do a civil action against you," then it shows up in the press and then your company says, "Hey, why," – like HBGary – also HBGary's name is dirt in the press because what happens? You get in the press with a particular incident and it's bad news for your company.

And then the final thing is stable [over]threshold. So in stable [over]threshold, this is something we started seeing in 2006. When the service wires started taking out IRC-based botnets and just knocking them down – just knocking out their command control systems – the bad guys realize, "Oh, the service (inaudible) have a threshold." There's a threshold of pain. The threshold of pain as far as service networks is service level agreement. As long as I stay below that service level agreement. So there's no impact to the service level agreement of some of the big carriers, then the service level providers don't wake up. They can't afford to. They will only spend resources on trying to stop the badness if there is a threat to their SLA. Otherwise, they can't afford it – there's too much badness out there.

So the bad guys stay below the threshold of the SLA – everything's good. Its criminal infrastructure's all nice so let's stay below the pain threshold. Now these guys work with each other; these guys collaborate with each other; these guys compete with each other; these guys attack each other. One of the biggest DOS attacks I've ever seen was over 60 gigs sustained over a three-day period in China going against two criminal gangs.

One gang was in north China; one gang was south China. Both of them had illegal gambling sites. Both of them illegal gambling sites – one in north China; one in south China. They take the illegal gambling sites and put them online – they got mad at each other because they didn't understand there were no borders anymore. These are two criminal organizations – big organizations – who traditionally were separated by physical boundaries – north China – south China.

On the net there's no difference; these guys started what they'd started doing IC and PUDP DOS attacks against each other. Just wailing on each other; stressing out the links between China Telecom and CNC, just blowing those things up. All the links were going down. China Telecom – all the Telecom properties are gone, like, "Oh my goodness. What are we going to do?"

And all of it because two criminal gangs decided, "Oh, we're going to go after each other on the net." This is the sort of thing that can have retaliation of it on there. There's also dire consequences. Slammer had two people die in Korea.

When Slammer happened, a Korean infrastructure went down and a DNS infrastructure went down in Korea, the gaming sites, the gaming parlors where a lot of money is all controlled by organized crime – we had basically guys who would go around and collect money. I go around every day and I gotta go around to one guy and I collect the money; go to the next job and collect the money; third guy- collect the money.

So Slammer's going on, it goes one guy collect the money. "Well, I don't have any money because the internet's down." Second guy – collect the money. "Internet's down. It's down so I can't do anything. I have no money." These thugs – they have no idea what the internet is. They're just supposed to collect money. Say, "We're going to make an example of people." Take the guy out; whack him.

Then people said, "Oh, we gotta find the money." Welcome to cybercrime. It's dire consequences. And this stands out for all sorts of criminals. You'd think the crime out in the physical world is all being done over the cyber world. And because we don't have international laws – and this is all international – the likelihood that we're going to be able to stop this with traditional anti-crime measures is very unlikely.

For us who here who operate network infrastructure – critical infrastructure – and we are tools for them to make their money, right? This is going to be a threat that's going to be with us for a couple of years. They have an entire ecosystem. In the past it was all about bragging rights. Now it's a very complicated system with a lot of moving parts.

And this slide is two years out of date because the moving parts I had on here are mule networks. In the United States, especially down in the south, there's this big thing about Secret Shopper networks. Anybody hear of Secret Shopper networks? Secret Shopper networks is how they go out there and get mule operations out.

These guys have it down to a T, right? So this infrastructure of this whole cash flow system going out there from somebody who writes the code and then goes out there and victimizes it throughout, DNS and DNS registries and registrars are part of this ecosystem. It's going through cycles, so we're going to observe cycles.

Like in 2006, we had a peak of D DOS and they were doing things in industry and D DOS went down. D DOS has come back; extraction hacks are coming back. We're seeing cycles that happen now. This is part of an economy. This is why you hear this whole thing that years ago there was a colleague of mine, Rob Thomas – works over at Team Comru. He and I observed this whole economy. There's an economy here for the criminal economy. We called it a miscreant economy is what we did. We don't use miscreant anymore for cyber criminal economy because miscreant doesn't translate very well with multiple languages out of English.

So what do those cybercriminals do when we start going after and whacking them? They do retaliate. Just see what HBGary and Anonymous. HBGary says, “We’re going to go after Anonymous crew,” – what happens – Anonymous Crew goes after them. So you have retaliation that’s going on out there now.

So what do we need to tell our peers who are probably not in this room, who are not taking care of their brand? ccTLDs is a brand. We all talked about this. This whole ICANN bit going on right now is all about the brand and the money behind it and things like that. We all see it here. Those of us here – we’re all kind of more of the geeky crowd here in DNS OARC.

This brand jeopardy is something important because what’s happened in the last couple of years – and this is probably the main point to communicate out the people who are saying, “Well, I don’t have the money to do it.” Or, “I don’t have the time.” Or, “It’s not important to me.”

So any of our peers out in the industry who are saying, “I’m not going to take the time or investment,” this is going to jeopardize you because the power is in the end points. We run a system and an architecture based off of the principle that Paul Baran put out and an end-to-end system, the con and the surf put out, right? So we have this combination of the architecture of the internet that we put out there.

The power is in the end point. We route based on autonomous system numbers. The power is in autonomous system numbers. We don't have countries on the net; we have autonomous systems. And an autonomous system has a constituency they have to take care of. AT&T has to take care of their customers. Verizon has to take care of their customers. BT has to take care of their customers. Merrill Lynch has to take care of their customers.

Each of these autonomous systems has to take care of their customers and constituents. They have control. It's not a technology issue; it's a risk assessment. So when they go out there and identify certain things we go out there and we take control.

Now one of the first things you saw with this is the Russian business network and Mokolo. Mokolo was a case. Mokolo was basically a bunch of providers getting together and saying, "We're not going to route to these guys because look at all the badness sitting over in Mokolo." This was an embargo. It's a bunch of service providers that embargoed. This can move over.

The autonomous system today has the control to go out there and say, "I don't want to go out there and take any traffic from that autonomous system that's sitting over there in some country because there's all this badness over there. Look at their reputation. Can't get all these reputation feeds. Lots of security companies give me reputation feeds." "Oh, I don't want to take anything from that ccTLD. That ccTLD – everything coming from that ccTLD that I see to hit my network is, in fact, trying to infect my customers, trying to infect my employees. There's too much of a risk. Filter it."

When you get a bunch of these together to do it, this is a consequence. Now, what are you going to do when you got a hospital, you got a medical staff, you got education universities, you got your government, you got your banking systems.

We already saw what happened in Egypt. When Egypt went out there and they went out there and unplugged the net, what happened? The pressure of plugging back in wasn't political, it was financial. You basically put the country at a halt financially because you unplugged the thing. Right?

So this is what happens with embargo – it's a real embargo. And this isn't something that say, "Oh, United States Government says, 'Oh, hey, all your servers went. Please turn it back on in the country because the State Department says this is an odd thing because this country is going out there and complaining because nobody's talking to them anymore.'"

But guess what? The United States doesn't have the authority to go out there and tell all autonomous systems to say, "Please remove your filters." They don't have that power, alright? So this is a threat.

So what can ccTLDs do today and this is why the slide's in there. If you have recommendations with this, please add this cause this is probably going to be something will become a DNS OARC pool for us and the community to talk to people out there.

So there's lots of operations security communities. We got a lot of people out there in operations communities, right? DNS OARC is one of them. You know, getting them on the DNS OARC crew is one of the key things on here.

So in here, what can you do? You gotta do something, you gotta participate, right? The community isn't looking from a TLD operator perfection. They're looking for engagement; they're looking for response; they're looking for please work with us. Because the cybercrime issues going on out there are so dynamic, you may think you may be doing the right thing, but the cybercriminals will innovate, they'll come up with a new technique.

So we're not looking for perfection; we're looking for dialog and work, right? And we're establishing communities to help out. So NX domains – how many people here are on NX domains? Alright. If you don't know about NX domains, send an email and apply. We're setting up. I was working with Andre Ludwig.

This is a group that was created as a consequence of Conficker where we're working with registries and registrars to say, for instance, daily you get security communities collecting all the evidence of bad stuff and they're getting all this data and they say, "Here's a particular list of domains that we know are bad. Here's evidence for them." And they send them out and we work on getting them tightened from the top.

So you dialog with anti-virus vendors; forensics engineers; security companies. You've got direct dialog with them under a realm of trust in it. So this is one tool to get into. Another thing is follow all the hard work. ICANN has got this whole team over here –in the other room right now with the SSIC with a whole bunch of documents. Follow the documents.

If you need help with that, email these guys. If you need help with that, get on DNS OARC and this room will probably help you do that. And then the third one on recommendations is building a relationship with the upstream.

They've got this train of trust within these security communities. You've got to build a train of trust. And trust is actually based off of action and response. And the first ones to work with is your upstreams service providers. This is like a basic 101 of network security today. You can't do it alone anymore. The days of saying, "I'm going to build a big fortress around my autonomous system and I'm going to be all good," it doesn't work anymore.

You need to have a tight dialog with your peers that you interconnect with on the net. So you talk to your upstreams cause hopefully you should be dual-honed minimally, if not more, and then you talk to them and you work with them, and you have a dialog with them. And that will establish a chain of trust that will be able to extend out when somebody says, "Hey, how can I work with you? Who do you know?" And then the chain of trust can be established and you have a good working relationship with people.

So this kind of segues into an example of an action. So Erik's going to come up here in a minute and so I'll unplug so he can set up. But in here what I'm going to be looking for is if you have suggestion to add to this – we're going to be putting the PowerPoint up there.

If you see things that say, "Here's other things we should say," the irony here is I'm giving this presentation to all of you. This is kind of like an example of how you can present it. If you're in this room, you're not the problem. Reality check. If you're in this room, you're not the problem because you're already engaged in the dialog. If you have a bunch of stuff on your infrastructure now, then people will come and pull you off to the side and talk to you.

It's the people who are not in this room who are the problem. And that will be something to scratch your heads and think about. How do you get those registries and registrars in this room? How do we get them in here? How do we get the technical operators in this room? That would be the challenge for us. Any questions?

Jay Daley: Hi, I'm Jay Daley from .nz. A couple of things. Firstly, if you want to join NX domains, come and talk to me. I'm one of the admins; we can sort that out. Secondly, I wonder if there are some… While there's a lot of very useful stuff there, and very, very important – some good messages for us – I wonder if there are one or two informational disconnects that we need to resolve to insure that these things get across correctly.

At the very beginning you had a slide that had a draft 10 ccTLD BCP02 on them. If that was ever best practice for a ccTLD, then it was probably in about 1905. It certainly has not been best practice for any time that I've been alive or I think probably the 50 years previous to that.

And so we need to insure that when… There is a problem of our ccTLDs being seen as being patronized at times by others outside the community and that document is a very good example of it. Now, everything else that you said is very, very good and I think very, very useful, so I don't want to just pick on that bit.

And I think that we in the ccTLD community generally have not been very good at explaining ourselves, at getting involved, at doing these types of things and we can do that. But there is a small current around of trying to tell ccTLDs what their own job is and we need to avoid getting into that by accident, just to insure that ccTLDs do actually get properly involved. And I'm certainly one of those and I think most of the rest of the room here are as well who think that we do need to be taking this much more seriously and doing much more about it.

Barry Greene: Yeah, I agree with that. One of the things I saw when I was sitting next to Dwayne - I didn't get to ask him – but I put that up just to see reaction because I said this thing is so old and I've seen other documents. It wasn't clear out. Which is also the approach that says, I don't want a ccTLD. The last one I was involved with was single core – sgTLD, and that was like decades ago.

So in there, this is also a message to the security community because people say, “Oh, they should do this or do that.” I don’t sit in the operator’s shoes for my ccTLD. What I can do is say, “Here’s a threat. Here’s a consequence. Here’s what’s going on. Here’s where the miscommunication happens. Let’s get them the dialog; let’s get them the NX domains on there. So here’s a guy to see on NX domains if you want to find out more of that,” because that is more operation, not educational – NX domains.

We may need to look at from something the DNS OARC as more of an educational dialog form or is something in ICANN already there to do that?

Eberhard Lisse: To use the priority of the microphone, I found this a brilliant presentation, but what Jay’s saying is quite correct. You deal with 255 or 252 different ccTLDs and each of them has their own bilateral relation to the internet, so even the word best practice is quite frowned upon. However, that said, one of the goals of the technical working group is to write sort of a handbook to say what works. And this is perfectly helpful.

Andrew Sullivan: I’m Andrew Sullivan. The last time I was at an OARC meeting, actually, I gave a talk which argued that in fact, we needed to get a civil society on the internet in order to solve some of these problems, so I’m in a certain amount of sympathy with what you’re saying.

At the same time, it worries me when we talk about, “Oh, the ASs are going to get together and they’re going to make up some rules for the rest of the internet. They’re going to start cutting people off,” and all the rest of it. We invented civil society as a culture precisely to prevent mobs from getting together, hanging people. I mean, that’s sort of the basic thing that we want, right?

And our experience with spam blacklists has been pretty bad. It’s very easy to get on one; it’s very hard to get off. And I’m awfully nervous about the idea that what we’re going to do is get ASs together in a sealed list over here somewhere and you gotta apply to join and then we’re going to cut people off. It tastes a little of –what do you call those – they’re like mob rule – and I’m not saying there is a mob rule here, but I’m worried that we don’t have that counterweight of any kind of government or anything like that and I don’t know how we’re going to build that. I’m not saying, “Gee, here’s the solution.” But I really don’t know what to do about it and it’s a scary thing to open that door.

Eberhard Lisse: The buzzword would be transparent manner.

Andrew Sullivan: And a lack of it. So if you look at the security companies out there today, and here I sit in Silicon Valley and I do a lot of innovation work, over the last five years, the ones that were getting the most money from different sources is basically blacklist protection. And one of the things that the companies haven't really latched onto is how easy it is to do an AS selective or it's out there.

So you see them out there and say in enterprises, "How do I protect myself to go down there?" It's scary because there isn't any checks. It's under control of end point and we've seen what's happened before. I agree with you. It's scary. What's why I'm bringing it up, right?

Nigel Roberts: Thank you. My name's Nigel Roberts from the .gg TLD. I'd like to follow on with that a little bit. I'm one of the few people who did put his hand up when you said, "Is anybody on the NX domains list." And none of our domains have ever appeared on that yet. I imagine it's only a matter of time just being a small ccTLD until somebody comes along.

But one thing really worries me about this. You said that these – and I'm not picking on NX domains particularly, but it's one of a number – you say they are evidence-based. Now that's not been my perception. My perception is an email pops into your email box and it says, "01PQZXYZ.gg is a Zeus domain. Please delete it. It's being used for phishing. Please delete it." And that's it.

Now, from a TLD operator's perspective, we operate by rules and the rule of law. We have contracts with whoever is a domain registrant. Now, our risk is that maybe – I don't know – 999 times out of 1000, you are right. This is an unknown, and the bad guy is never going to take us to our civil court in our local jurisdiction to complain that we cut him off.

But if we cut off a genuine person who has been somehow accidentally included, we're the ones who get sued, not the operators or the blacklist. And that's a problem, I think, that we need to overcome.

Eberhard Lisse: Hang on, hang on, hang on. We're transcribing this. And I would like to move it along so I think this will be the last question.

Barry Greene: So those are really good points. This might be something where DNS OARC can step up on the technical side it says what are the requirements? For the security operators to say, "If you want us to do action, here's what you need to provide." Those guidelines are lacking; I agree with you, because you got a contractual obligation. You have the legitimate ones and you gotta be risk averse on it. It's the balance there.

Eberhard Lisse: So we have remote questions and I want to try not to overlook these.

Christina: Thank you. We have an unidentified participant who's wondering, "As a ccTLD operator, we can have many policies, but what happens if the government/laws within the country do not support those policies?

Barry Greene: Contracts. So the answer to that is you have legislation within a country, but within legislation within a country, you can have contractual relationships that go beyond the local legal system. So it's all a matter of business contracts, so it's a desire around that. And this is speaking for somebody who's been working and running service [bars] outside the United States. I just get beyond the… You go above the law by doing a contractual relationship with my customers.

Eberhard Lisse: Alright, Jay just mentioned that the list seems to be evidence-based. Maybe you and Nigel can sort out your differences. The next one is Eric Ziegast from, again from ISC. We set up these two presentations to follow each other for a reason.

Eric Ziegast: Great. I'm going to be a little less animated today. I'll be sitting up front and I'll go through this presentation and it's a little bit about the Conficker sinkhole effort and what we did to help capture some of the domains that were being used by that, how was the domain base attacked and I'm actually reaching out on CCT-NSO day to reach out to some CCs to see if we can actually add some more censors for domain capture.

First I'm going to go ahead and give you a little bit of background about what it is because maybe not all of you know what was going on and how it operated. What we're doing and we're still doing – it's been a long time. I mean, the thing started back at the end of 2008 and here it is 2011 and it's still alive and people are probably getting pretty tired of it. But there are still things that people can do to help and I'll explain that later.

So, Conficker, it's a worm, it's a virus, it's a superbug. It's not your standard old virus. It's one actually amassed a very large number of bots. You can argue about the size of it and the methodologies to use it, but it was impressive. It garnered the press and attention from a lot of people.

For some background reading, you can do this afterwards. At a very high level you look at a newspaper article. There's a great Wiki that's maintained by Shadow Server called confickerworkinggroup.org. As people got organized and started contributing that was a great central resource for information. And if you really want a technical description about how the various modes of Conficker worked, there's a great write-up from SRI about all the technical details for communication and infection and such.

This is kind of an example of a war or conflict where you have two sides opposing each other. The security community and the developer or developers behind the software as a community stepped up. In some cases an operator will just kind of get hit and they'll go away and then they'll come back to fight another day.

In this case the developer actually fought back to get around whatever mitigation efforts people came to use and he used – he/she/they – used DNS and that became part of the collateral damage. Now, it kind of seems at a stalemate right now. We haven't won the war; there's no one in jail. But at the same time, we haven't lost; this spot hasn't gone out and struck with force in a severe way which it could have.

But it’s still out there, it’s still active and I guess it’s kind of like polio or something else. It just kind of works in the background waiting to spring up again and flare up. But there’s something that was actually positive that came out of that is that one of the security providers, perhaps for the first time, they started working with each other.

Instead of being a competitive effort and saying, “Oh, I know these bots and I know those bots and my product’s better and I’m not going to let you have my information,” they realized this is bigger than any one security vendor could tackle. So a lot of the people came together and a lot of resources were used, both from the ISP community, from the TLD community, from the vendor community to start working together and figure out what they can do together.

The virus itself is pretty easy to detect. Joe Stewart came up with this wonderful idea of an eye chart. What you see on the bottom is well-known websites, open BSD, Linux, free BSD – those would show up on anyone’s computer if they were infected or not infected.

But the top – one of the things that this virus did is it prevented you from talking to your AV vendor so you couldn’t go back to F Secure or Secure Works or Trend Micro to get updates because it wanted to –

through self-preservation – wanted to keep running. So those images on the top wouldn't load based on which variant you had.

We did a great job with DNS containment. Initially it was like most malware that you see out there – it'll go ahead and take some dynamic DNS somewhere and register it and then all the bots in the network – you know, 1,000 bot - 10,000 bot network, would all come into a domain, it's registered overnight, kind of sits there. And eventually they start using it.

Well, we found the domains, well, other people found the domains and they got stomped on. So they're captured, but it's still, there's some machines infected out there so the operator said, "Hmm, alright, well, let's see, what if I use 500 domains today? That was pretty novel. Maybe they should get a patent. I'd love to see them at the patent office."

And it worked pretty well. I mean, they had a whole bunch of bots that were updating, but then the security reachers, they caught on and realized, "You know what? We could probably register a bunch of these domains and take all the DN to the sinkhole." We actually got a good chance of containing it. There are only a few registries that they're using. I think there's about – the number was about eight at the time.

And eventually there was pretty good containment. We don't know for sure – maybe a few people got through, we may not have been able to see exactly which – there may have been one or two domains in there which helped others update. But for the most part, we think it's working.

Some of the things that we did is worked with a few of the members to register three name servers for every domain. If it were, say, ISC doing all the name servers, well people would just attack ISC, right? But if you put these name servers in different locations and such, there's a better chance that we'd be able to keep some of them running. But then again, take a really large D DOS like I talked about yesterday, hey, maybe you can take them all down.

The name servers point web callback hits to a web server so the software that gets infected on the PC. It continues to try to call home. It's not necessarily recognized or activated yet. Typically when you amass a bot you want to do something with it. You want to do a spam run; you want to go ahead and steal credit card information; you want to use it as a D DOS against someone else.

Well, you have to be able to reprogram yourself to do that, to get your task list and that was what the call home mechanism was using. It used ACTP Fetch It as its update facility. When you have a sinkhole and you actually point the domain to somewhere else and you don't give them any data, it's actually a method of containment and it's also used for detection. But this is every day, 500 domains get registered and used and we have to keep that effort.

So you see that nice little picture. We're just kind of sticking our finger in the dike to keep the dam from breaking and we have to continue to do that while that virus is still a threat. And it will probably continue to be a threat until the people behind it are taken care of.

So on a sinkhole, they come in in some common infrastructure. Over at Security Information Exchange we have a bunch of sensors that are on these web servers and they come on in. Currently we're getting sustained about 8,000 hits per second, kind of respectable. And we feed that out to a whole bunch of researchers that are on our network and they take it and put it into their systems.

There are some public benefit efforts, there are some commercial efforts – basically security vendors who will say, "Hey, your machine is infected with Conficker." And the job is pretty much to go out there and for most people to go mitigate it because you don't want to have a virus on your network and this is a pretty easy way of detecting that you've been infected by this virus.

But we're teaching – again, we haven't captured anyone. The person has learned that, "You know what? They're actually doing pretty effective so let's up the ante a bit. Let's go to 50,000 domains a day." And it included ccTLDs. And for some of the reasons that Barry might have alluded to before, it is wherever there is generally consistent operations within the gTLD operators – cause they have lots of well-established, lots of domains, lots to lose – there are a lot of ccTLDs which may not be staffed or operate at the same level. And they take advantage of that.

While some of the ccTLDs can step up and go ahead and work with us to help capture data, but also try to do containment, there will always be one, 10,100, who knows how many can't really step up and the weakness in the registries and their ability to work with operational security were exposed.

There was another presentation given by Norm. He was at CERA at the time. He actually works at ISC now over at ICANN 35 at Sydney, with some success. We figured out how to say, "Here's a years-worth of domains; go ahead and register them," and hopefully that will provide us some data and it will help us have some success at mitigation.

But if we don't have everyone participating, there will always be that one registry – even if it's just one – the bots will be able to contact and phone home. And so basically we're set up to fail here.

There are other methods for communication by the virus at this point. They can use P-to-P between bots where it automatically probes out at IP addresses. In the meantime, this is a graph that's available online. It actually shows you that there is a decrease in activity of Conficker over time. The top level is Conficker B; the blue is Conficker C and the number of hits that are coming in – B is much more well contained; C is less contained, but does not seem to be spreading. We never saw it activate.

So we're thinking, "Gosh, you know, winning!" But the problem is we're not winning so we're basically at a standoff. We're at a truce. And the ccTLD participation – that was back in 2009, now it's 2011. It has tapered off such that only two sets of domains are actually getting registered every day. So we're getting less data than we used to.

What did the registries learn? The ones who participated with us learned that there is a security community that is out there and you can gain trust with working through to affect positive change in your ccTLD. The registries that didn't participate or chose actively not to participate learned that, "You know what? All those other people are going to do this work anyway and go back and take a look. It's working anyway, right?"

Well, I guess that's a valid argument. But mostly it's an unfunded mandate. Some of us spend some time working on this. I say mostly because some people have stepped up. A lot of the people are collaborating together and it is mostly a public benefit effort.

There are some security products that are out there that are free – Shadow Server and At Lists – you can register to get your information or you can go to paid vendors who are actually taking the data and actually using it as part of their commercial products and enhances their products that are making money.

We used to have a focus of hey, let's do what we can to contain it and well, we can't contain it. And what can we do in the DNS community in like, say, a DNScert or other things to really go after this? And we have to realize that we're never going to be able to contain this. So at least what we can do is keep chasing that long tail of that graph. If we basically get, say, another five ccTLDs to step up and do some auto registration, I think we'll have a pretty good, at least, ability to take the Conficker C and keep chasing after and doing mitigation on a host by host basis by certs, ccerts and ISPs.

So, ccTLDs can help. Basically we can make the daily list of the domains for yesterday, today and tomorrow available from CWG, use a [T-sig AX fr] method; download it; extract the domains for your ccTLD. All you have to do – because we have no idea how you operate your registry –

is say add domain and remove domain. And we're not asking you to say we're going to reserve these domains for a year. It's just basically as the pestilence comes by, you basically lock yourself indoors with those beams and then it will pass by and you can come out of your house again.

So for each domain, just for a day or three days, you'll be preventing people from contacting that on that day or on that instance, but providing data and continuing to provide the data to help mitigation of who's infected. And we have some other scripts that we can make available to you.

We also noticed that at least one ccTLD is stepping up and says, "I actually want to run some of the sinkholes." And in some cases, they may have connections with other providers inside the country that they would have better visibility into the networks rather than having them all at the centralized Conficker sinkholes that are already in operation.

So we actually have some software. It runs very well automated and very efficient. And if we actually had more sinkholes we could actually spread the risk around of retribution.

Specifically we'd like to thank a few. ICANN and Microsoft both stepped up to provide us some funding for some of the hardware that's being used so we really do appreciate that. Otherwise, data would be lost and people wouldn't be able to analyze stuff.

GTISC has done a lot of – Georgia Tech Information and Security Center – has done a lot of work. But also, not only that, but there's actually some funding that they provided under contract to ISC so we can actually help do some of this coordination.

Specifically there's some people – Rick Weston has gone out of his way to register a whole bunch of domains and Dr. Chris Lee provided some analysis and some graphs. Shadow Server does a great job with the Wiki site. And there are a cast of others including leadership. It takes some leadership to actually herd all these security cats together to work together and they are nameless and there are many. But they are well appreciated. Questions?

Eberhard Lisse: Okay, thank you very much. Again from the priority of the microphone, as you all know, I'm a medical doctor and in Kenya at the ICANN meeting I spoke to Olof (inaudible) from .at, and he mentioned that he was having big problems to convince some radiologists in his university to shut down their machines to fix the virus on the Windows.

So I went to my own Radiology Department. It was a drama. Every single CT scanner, every single MRI, every single ultrasound machine was infected. It's quite difficult – first of all, it's difficult to shut these things down because they don't operate for a day, it costs them money and it's difficult to convince a radiologist who knows nothing about these things that they should do this.

But also the human factor was that the guy in charge of IT took it as a personal affront that we figured it out and he didn't. Was a bit of a drama. But it's quite important that we managed to get at least something on the road and any questions in this regard will be taken now. Any comments? You overwhelmed them. Thank you very much.

Eric Ziegast: Feel free to contact me offline.

Eberhard Lisse: Okay, the next speaker will be Brian Cute from the Public Interest Registry. We usually have a host presentation and last time in L.A. we had .US. We had .com do it anyway, so this time I felt we invite .org. Not necessarily their (inaudible) provider, but they're new CEO and we'll just take two minutes to run the presentation.

Brian Cute: Thank you very much. I'm very pleased to be here. Thank you for the invitation. Let me just start off with a couple of opening thoughts. Again, thank you for having me. I'm very thrilled to be speaking about .org and public interest registry. I just became the CEO on February 1; I'm a month into the job.

The other thought I'd like to share with you – I know you have your tech day meeting today, but the ICANN opening ceremonies just took place and there were some very important speeches, Larry Strickling in particular at the Department of Commerce, gave a very important speech just moments ago.

I had the privilege to serve with him on the Accountability and Transparency Review Team as a member of that team that made recommendations to ICANN as to how the organization can increase and improve its policies, its accountability and transparency in the work that it does. I know that all of us in the room care deeply about that subject so I encourage you, when you get out of the work sessions here to read those speeches and take good measure and, as we all do, participate actively to improve ICANN as we all go forward.

So with that note, again, .org. I'd just like to take about 20 minutes or so to give you an overview from where I stand, being new in the organization what I see as .org today and what the vision going forth is and to touch on a couple of technical issues that are very important to all of us in the room.

Do I have the controller? Can we go to the next slide? Excuse me. So the good news, I'm sure all of you know the brand, .org. It's one of the easily recognizable brands on the internet in the domain name space. We have just surpassed 9 million registrations, which is an important benchmark for us and we are driving strong toward the 10 million registration mark.

And at .org and at public interest registry, first and foremost, we are not for profit. The ethics and vision and ethos of our organization is centered around the mission that people who on the internet wish to do good, that we support all those types of registrants on the internet. We care deeply about how we run our business, the values that we project and the registrars and the .org space we believe also are people who intend to do good, both on the internet and the society. And that's the core of our belief system.

Domains, as I mentioned, we just crossed over 9 million. All of us in the room know that increasing competition from new top level domains is coming as soon as ICANN gets the policy process finalized. And so it's important for us to keep the strength of .org growing, as I'm sure it is for all of you in the room, with your respective country code TLDs – that's where our forward focus is.

With respect to our name base, we have very strong renewal rates. That's an important benchmark for us that we believe is a reflection of the type of registrants who take a .org. We also, as a registry, are very active in policing abuse as well. We take that seriously and take seriously the cleanliness of the registrant base in .org as well. So these renewal rates are an important benchmark to us and they're strong, as the slides reflect.

In terms of our base markets, in the second half of 2010, as you can see, the strong majority is in North America and a very important slice – 24% - is in Europe. And we have smaller percentages in Asia, Pacific, Africa and other regions. It is important at Public Interest Registry that the internet is expanded to all markets and to all potential internet users. The smaller slices are regions where we'll be spending more focus in the coming years seeking to grow .org and seeking to do more broadly good things to provide internet access or support the expansion of the internet in all countries around the world.

Here's just a breakout pie chart of the top 10 markets down to a few more specific countries. Not much different from the prior slide in terms of the demographics of the base. And speaking of our values again, to put some specifics on it, these are three things that we value highly for .org.

The first is that it's a trusted domain extension on the internet. It's very important to us and very important to our registrants. I'll speak about DNSSEC implementation shortly. That is one measure of how important we take this notion of trust. As we drive forward as an organization, we'll be looking at new services that can enhance the trust proposition for someone who's visiting a .org website. Trust will be a key driver in everything that we do.

Secondly, there's a notion that .org is just for not-for-profit organizations – that's not, in fact, the case. And we have seen in looking through our customer base and taking demographic profiles that it really is broader than that set of users; it includes online communities.

We're seeing trends, if you look at things like Wikis, where people come together in that kind of beautiful ad hoc way on the internet. People are choosing .orgs very often to do that sort of thing, and that's encouraging.

And also, third, it's a web address to advance a cause. We're seeing, for example, corporate social responsibility efforts when large companies undertake missions to support cancer research or do social good, we're seeing that they like to select .orgs to put those missions out to the world, not use their .com site. We think that's an interesting trend as well and one that's very consistent with our values.

So here's a snapshot as I referred to Wikis in the last slide, but here's a snapshot of some of the registrants that we're seeing in the profile. So, indeed, it's a pretty broad base of registrants who are choosing .orgs and a couple of prominent names, as I'm sure you know American Red Cross, Greenpeace and others who are prominent users of .org – Craig's List in particular.

And the CEO of Craig's List was here, I think yesterday and gave a really interesting talk about why he chose a .org. He made a conscious decision not to go commercial, probably walked away from a lot of money. But it fit with his values and what he wanted Craig's List to be and that's something that makes us smile at .org.

I will take a note since American Red Cross is on the screen. We are hosting Music Night on Tuesday night. I hope you all can come and have a lot of fun and do karaoke with us. But something we'll make mention there – I'm sure top of mind for everybody – is the disaster in Japan. We'll be providing some information of relief organizations that anyone who wants to can contribute to and we hope that you do because they're in dire straits in Japan right now.

From an operational perspective, and I'll touch on some of the technical issues that are of greater interest I believe to you all – IPV4 versus IPV6 - this is a very important issue to all of us and at PIR. As you see, we've got two slides here that show you the respective queries per day for us about IPV4 and IPV6. We are IPB6-enabled at the registry level. We will play our role in trying to encourage all the users and service providers to transition to IPV6. We're going to participate in IPV6 Day which I believe is June 8, coming up with some other large companies. But you'll see Public Interest Registry playing a prominent role encouraging the industry to make that transition that we all need to have happen for our businesses to continue.

With respect to DNSSEC, I just came on board February 1 so I can't take any credit for what PIR's done, but I think you all know PIR took a strong leadership position on the implementation of DNSSEC. This was a key issue. If you look at DNSSEC from an economic perspective and all the different actors who have to implement it to create that end-to-end chain of trust, it's one of those security enhancements that not everybody in the ecosystem has an economic incentive to take on or to implement. It's just one of those things.

But, nevertheless, PIR viewed this as a critical step in enhancing the security and trust of the internet and the DNS. And so PIR, as it's consistent with its values, decided to take a leadership role, decided to push for DNSSEC implementation for .org. We've done that. We realize it's still a slow uphill climb, but the fact that the root has been signed, the fact that .com will be signed sometime this quarter, all very encouraging things.

I've listed here the number of registrars who have gone through OT&E for us – we've got 34. Now, yes, we have a total of 381 who support .org and 34 may seem like a low number, but it's an encouraging data point to us and we're working with the registrars to encourage increased implementation at the registrar level. You'll continue to see .org and PIR proselytizing on the adoption of DNSSEC – really critical for all of us.

And that's the number that we have signed again – low at the outset but we knew we were a first mover and we knew this would be the beginning of a process. So, again, we're encouraged to see companies like AOL and Comcast and Fandango implementing DNSSEC. We think it's a good sign and as we get more momentum in the first half of this year, we'll continue to work with all of our partners to help along that path.

And before going to thank you, I want to note too that in talking about extending .org to developing countries and markets outside of the U.S. and Europe, we do view .org as being a unique TLD for people who have a mission. We do look to expand the internet for the better good of the internet and society. We take a collaborative view of going into markets. We're not looking to compete necessarily head-to-head with ccTLDs. We think we have a unique registrant base and a unique mission and we want to see the growth of the internet more broadly.

So I'm hoping as we begin to focus in in the next few months on what those initiative will be in foreign markets that I have a chance to meet some of you and collaborate and talk about ways we can in tandem help increase internet usage and the betterment of the internet and society. So with that, I'll leave it open for questions and thank you very much.

Eberhard Lisse: Thank you very much for your presentation. Any questions? I was thinking we let you off easy.

Brian Cute: Ah, Mr. Daley.

Jay Daley: Hi, Brian. Jay Daley. How important do you think it is for registries to have – your last slide there that public good/greater purpose to keep them honest and to insure that they serve the internet community rather than serving their own profit or any other motive?

Brian Cute: I think it's terribly important. You know, Jay, that in a prior life I've worked at VeriSign and Network Solutions and I have no qualms with a for-commercial model. We all have our proper role in the ecosystem. PIR is a not-for-profit and those are the values that we espouse.

In terms of being a good actor and having good policies and practices, that's something that I was personally espousing in former life, will continue to espouse. And one example is – and I know this was a controversial issue for some – but the issue of vertical integration. One of the outcomes of that decision was ICANN putting a marker down that registry data should not be abused for the sole commercial gain of an integrated entity.

We think that's an important principle; that's the sort of thing that you might see as trying to push ICANN along. And compliance is important; good policies and practices. Who was it in the opening speeches? I think it was Larry Strickling who talked about trust being at the center of the internet model. Trust is at the center of the internet model. If the user thinks they're getting scammed or they don't have a good opportunity or somehow they're not being well treated by the registry, it's not going to work for any of us. That's really important.

Paul Loggins: Thanks. Paul Loggins, XLS Corporation. The list you briefly partially showed on the slides about registrars that are supporting DS Records Management?

Brian Cute: Yes.

Paul Loggins: Is that a public list that people can obtain? Because for us as integration vendors, it's really useful to be able to contact these people. And if you're one of them, please contact me because we are doing integration and we would really like to know these people that support it.

Brian Cute: The answer is yes and we can get that to you, whatever the best route is. If there's a list of attendees here with email addresses, whatever's best, we're happy just to provide it to you. Oh, it's on our website. There you go. www.pir.org.

Paul Loggins: And the names of ICANN's registrars with what registries they are accredited to is also on the website. Which registry they are credited in? Which ending .org are comments on? For the presentation last week I had to dig it up so I had to pick it apart. But anyway, I have this email address; you can come to me; you can email them then and you can sort it out.

Brian Cute: We'll get that to you or you can visit the site. Either/or.

Eberhard Lisse: Alright, good. Thank you very much.

Brian Cute: Thanks very much.

Eberhard Lisse: Let me just quickly figure out who is next. I must look on the spreadsheet. Next speaker is going to be Theo Kramer from co.za, that's c-o z-a, second level domain. UniForum is the organization doing it. They are re-writing or writing their own implementation and so I decided I call his presentation Trials, Tribulations and EPP.

Theo Kramer: My name is Theo Kramer. I'm from UniForum, South Africa. We are the administrators of co.za and I hope to be able to take you through some of the background that we've been through, where we've come from, some of the problems that we faced, some of the approaches to trying to solve these problems, our implementation.

What I'll do is I'll have a brief chat in terms of the decisions that we made and the architecture – what we've come up with and our policy framework and, of course, the benefits that we hope to achieve out of this exercise.

The co.za registry currently has over 650,000 domains; it's a second level in .za. There's a couple others there as well; I think there's about 18 or so, but we are by far the largest of those registries. We've been running this since 1995. We really started this as a user group, the Unix user group. At the time we were seen to be non-aligned and there was a bit of an ISP meeting and it was decided that they would let us take over running of co.za. In September 1995 we started off with about 400-odd domains and today we sit with over 650,000 domains in the zone.

Our systems really evolved organically. We had a couple of hackers who put the system together as the needs dictated. We evolved our systems and we really bolted bottom up. During the early days we really worked within a regulatory uncertainty; there was no really regulator, but over the years those things have changed and we now have the .za DNA and we've formed a relationship and we believe that we're hitting on the right track together.

With the co.za current system, we have de facto Rars. Effectively anyone can register a domain name with us if they are able to set up the necessary name servers and including some technical things around that – being able to do reverses and things like that.

We operate a post-payment system effectively allowing you to register a domain and park it for a couple of months before we suspend it and before we delete it should we not receive any payment. We have a very simple email interface which makes it really, really easy to register domains but of course, we get a lot of complaints about that as well – people want what they call a proper interface.

Looking at the current system, we've got huge problems. First of all, anyone can register at the moment so, from a scalability point of view, we end up with a huge client base and that it's just become very difficult to handle the larger the zone becomes.

Also from a technical point of view, our systems were designed bottom up. They consist of shell scripts, [Perl] programs, some C programs, all acting together and never really much thought was given to designing a team implementing a top down design which we could scale.

We believe with the current system we could probably grow up to a million, perhaps two million, but beyond that, things will really, really become problematic.

As I've mentioned, we've got other shortcomings. People park domains with us; we have no formal relationship with our registrars and, of course, there's industry demands as well. "Guys, you need to put an interface together." We get that from the community out there – I'm not too sure if they always understand what they're saying, but we will see as we progress.

And, of course, with the .za DNA really coming on tract, we're also starting to find regulatory pressures, both from government and from the .za DNA. And we understand that we have to work together for the best of the internet in South Africa, in our region.

Also, if you're working with a system which is stagnating, you get a problem with retaining your skills. Guys want to move on, they want to move on to better things. So those are also problems that we've had. And, of course the other side of the coin is that there has been a huge investment in our legacy system by the ISPs out there that they affect our rars, they've created interfaces for it, now need to maintain it and of course, we can't just wish that away either.

So, for better or for worse, there is really only one standard out there and it's called the EPP Standard or the IETF Standard 69 and we understand that we have little choice; that's the way we have to go.

We've done a lot of technical research and benchmarking looking at some of the products that are out there but we've hemmed and hawed, but then we've gone back to what our policies say. There can be a lot of policy research and a lot of policy development. We looked at international biz practices; we looked at our current policies and we realized that, if you're going to put something together, that's really where you have to start. You have to start with your policy. The policy really forms the blueprint of any registry.

Once we developed that we engaged stakeholders, we engaged our ISP community, we engaged our clients, we engaged the .za DNA, being the authority, and we proposed our policy, we discussed the policy and we made the amendment as required as the feedback gave us.

We also realized – and this is in discussions with the .za DNA – that we will probably be moving into a central registry kind of environment and if we do something, hopefully we can make it work for not more than just one registry; we can make it work for other registries as well. There's a couple of other second levels in .za which may just find the software very useful.

One of the things that we wanted to do as well is with this particular project is we wanted to develop a local resources. Got a whole bunch of software engineers in our part of the world; we've got registrars who we'd like to bring up to speed as well and we also are looking at being able to develop that within the African region as well. So that was really our approach.

So we decided, "You know what? We're not actually going to take something off the shelf, we're going to develop our own." So tracking a little bit forward and looking at what we've got right now, and our implementation, what we decided to do was to make sure that we have a good separation of services, of course with the necessary close coupling between the various services and by doing so, we'd be addressing issues of technical scalabity.

So our system basically consists of a separate EPP interface and message validation front end. This is really the front end which… It's a secure front end, it's a secure authenticating front end, where only accredited registrars can interface to and what that interface does is it serves the purposes of authentication, but more than that, it takes very single message that we get, every single EPP message, and validates it according to the schemers, as defined in the various RFCs.

What we also wanted to do, what we also thought was very important - and this for a scalability of cross-registries - was to separate the policy framework from the policy implementation. So what we effectively developed was a policy engine and we developed a separate language which is based on XML in which a policy definition is implemented. And what that results in is really a structured policy definition and it contains an assembly of EPP primitives and the necessary requirements around that to conform to the particular policy of the registry.

With that structured framework we get some extra benefits. Once you have a structured framework which is defined in a particular language, in our case XML, you just look at that and by filling in the necessary blanks, you can automatically generate the full registry policy document.

You can also take that and you can populate it with the necessary test procedures and at the press of a button, you can generate a full test sweep for a particular registry based on the policy.

So that is really our implementation. If you look at our architectural diagram, it basically spells it out. Each one of those blocks is a particular service. We have the message handler which is the front end; we have the registry engine which gets fed the messages from the message handler; we have our administration interface for registry administration; we have a separate WHOIS server; a separate DNS and of course, a separate registrar interface with a management information system providing all the necessary statistics which may be of value to both the registry and any authority associated with the registry and of course, also, the stakeholders. And of course we have a financial system.

Each one of these blocks that you see is scalable across the .za plain. So the system from an architectural point of view, from a hardware point of view, is totally scalable.

The way that the policy framework works is what we've done is we've, as I've mentioned, we've developed a language and that registry policy definition is maintained – you can either do it with my favorite being VI or you can do it with the registry policy (inaudible), and you can create your own policy.

As the diagram shows, at the press of a button you can generate the documentation; you can generate the policy test definition. Again you can feed that policy test definition into what we all our test engine and you can test the registry software with that.

That test engine, we will probably also make available to the registrars because that also forms the basis for a registrar interface. So what we've done is kill two or more birds with one stone.

Little bit of a view of the policy editor – how it looks. Effectively you have a registry policy and you have the various objects – the domain, the host, the contact and the various operations on those objects within the policy framework of the particular registry that you're operating.

One other big thing that we realized with this whole policy engine is that you want to be up on our case; we want to be able to run more than one concurrent policy because we have a legacy registrar as well. And that legacy registrar operated on a post-payment system and it'll probably carry on working on a post-payment system into the foreseeable future, really, a separate policy from the new accredited registrars which will be working on an up-front payment mechanism. So we can run more than one concurrent policy at the same time.

Policy test framework – this is the kind of environment which we've developed and which you can just run the full test suite based against that policy as the previous diagram illustrated with the test mechanism that we had.

So the benefits that we have is we have a structured standards-based registry system assembly. We ended up with an automated registry documentation generation facility, an automated registry test suite generation mechanism. We can run multiple concurrent policies and of course, we now have also the separation of skills. We have the software development techies working on the service level stuff, the real services.

But we can separate out the policy definition to the people who understand the main space. Not only does that give us a nice separation, but it also has the extra benefit of separating ownership. If this goes to other registries, the other registries will own their policy definition. They'll be able to create their own policy definitions using the toolset that we've developed and effectively take ownership of it and effectively a fork in the road.

And of course, one of the other benefits that we've had, we've had no EPP registrars in our part of the world. This has been a road that we've taken together. We've got some test registrars on board. They've learned all about being able to register domains using an EPP interface and yes, so we've built the skill level in our part of the world with that.

We've also put a webpage together which is that particular webpage and you're free to have a look at it. It gives some detail that has the co.za policy document there. It has the quotation document there for registrars and it also has a technical section which gives full examples on how to interface with the co.za registry. Thank you very much. Any questions?

Eberhard Lisse: Okay, thank you very much. One question – when do you think you will go live?

Theo Kramer: Well, what we are currently doing is we currently have a parallel process so we're running the system. We have the legacy system and we have the new system running in parallel. We are not publishing any data on the new system, but we will be running that live for a couple more weeks and we'll just make sure that everything that's happened on the legacy registrar is what's happened on the new registrar.

So in that test period, once that's done and we're happy with that, we'll get some of the test registrars on board and we'll get them to start doing some live operations and once we're happy with that, we expect to hit the button and go live, allowing other registrars to come on board. This will all happen, I believe, over the course of this year.

Jay Daley: Will you be… Jay Daley, sorry. Will you be publishing your registry policy technology specification?

Theo Kramer: The registry technology specification.

Eberhard Lisse: License model.

Theo Kramer: The license model will be… If that's correct, if I understand you correctly…

Jay Daley: Your registry policy language that then translates into XML that then gets interpreted by the rest of your system?

Theo Kramer: That will be available to anyone who wishes to talk to us in terms of registry software.

Jay Daley: So only as registry software? You won't be providing that as a separate protocol specification dependently or for software?

Theo Kramer: I don't think that's something that we've given much thought to at this stage, but that could happen in the future.

Jay Daley: I would be interested in understanding that.

Eberhard Lisse: Will this be Open Source?

Theo Kramer: There are components of it that won't be Open Source but we are willing to talk to the registries should they wish to go that way and we will most certainly make sure that our software can go into escrow so that effectively the registries will have access to the source.

Steve Deerhake: Steve Deerhake, .AS Registry. Question with regards to – like I see contact data. Did you have much of an issue with the quality of the data that you accumulated over the years through your email templates with regards to getting it cleaned up before insertion into your new EPP-based data base?

Theo Kramer: Yes. Short and sweet – yes.

Eberhard Lisse: Alright. I don't want to be cutting any discussion short, but we are running a few minutes behind. And since we have a packed program, I don't really think we want to have a break. The next guy would be Mauricio and after that is Richard Lamb. I don't know him but I know… Oh there he is. There you go. Mauricio is from the .cl and he contacted us on relatively short notice because he had some interesting experiences on his network.

Mauricio Vergara: Good morning. My name is Mauricio Vergara, I'm from .cl and I have an operation I think that most of the people have known for some time and we would like to show what is happening in .cl right now. It's this question that we are having. It began in this year, in the beginning of 2011 in the first week of January and we have a lot of [picks] on our traffic with MX queries that have a similar pattern.

One of the things that we have seen is that all these extra queries have the [user desire to] turn it on with and almost every query is an MX domain. And the other thing that we have seen is that the transaction ID of the DNS query is always lower than 256. So how we see this?

This is the first week that we have this pattern and, as you can see, this is our normal traffic – 66% of queries; 15% of MX, but on that first week we have this initial 80% MX queries. If you see now on what are the responses of these queries – mostly are NXDomains, the red ones seen here, and almost for, I don't know, 75% of all the queries.

You can see that we have our Recursion Desired turned on on almost every query that is very correlated. So what happened? This shows what we have seen on the peaks that we have in the past two months. So in the first week, we have our normal traffic. But then week No. 2 and No. 3, we have peaks running the 12,000 queries per second on the servers that we are managing initially and from (inaudible). We currently have six Anycast clouds – three of them are still contracted to other persons like PCH, Net Notes and S&S from ISC.

So the thing that I'm showing here is the one that we manage – three Anycast clouds just to get it over the rule. On the third week we have a peak of 22,000 queries per second and then the next week it went back to normally. So we think that the thing was over.

But then it started again with a similar rate and for another week came back normally. But we got really scared on the next few weeks when we got almost 60,000 queries per second running in our servers. So we started to have some problems with this and we are still having this issue running right now.

What are the top 10 hitters by country, it's very, very distributed as I will show you. We have like, half a million queries difference asking for this pattern. This will show you the distribution of the queries, what is the source of them. Almost ever query has no pattern at all. If you can graph it on a map, you will see something like this where the red ones are the heavy hitters on this. So it's very distributed, as you can see.

What are they asking for? We have found four common patterns of what they're asking. Maybe some ascii to hex kind of name – we don't know if that's okay. We have seen some dictionary attacks and we have seen some malformed lists that started with some kind of number pattern and then some dictionary again. And some final users may be something like it's not work configured to ask for some ISPs and stuff like that.

So, what we have done since then. We have redistributed our traffic between our Anycast clouds and we have been using a lot of AS-path prepending on [BEP]. We used to disable temporary logging on Bind servers because the performance was not that good.

We changed our last unicast note to an Anycast one in the last two weeks. And we have to improve our band-widths on our main site and queries per second and conntrack monitors to the alerts coming in faster. We tried to contract other TLDs and associate to see if this was happening to other people. We have a few responses from people that was seeing this, not as big as us, but some people tell us that almost everybody saw these, but they stopped seeing that behavior. So, it's pretty weird for us that it still is happening. So we're trying to gather more information about this.

What things have we learned on the way? Well, our international bandwidth in Chile has almost topped it on our main site, so we have to change that. We have an issue in one of our servers that we have a filter built with iptables that let it through or the port 53 but that doesn't work very well with the Linux distribution that we have. So we have to change it to our rawtable to let the packets go through for it. We saw some ISPs that have EDS that stopped some of the traffic so we were not able to answer everything in those servers. And another ISP has some problems managing the small packets flood on border routers.

So what are our conclusions on this? We have think a lot of methods to try to stop this. We think that it's less complicated to keep answering an X Domain to all these people more than just block or sinkhole anyone or a pattern. We have been over-provisioning ourselves - that is the key for us.

We have thought that DNS service providers that rate you on a per query basis – it will be really expensive if we could have one of these. So we think that our model to have a lot of Anycast clouds is much, much better. And we think that ISP contracts that we have must be prepared to give you more bandwidth in case of emergency.

So the final question for these is is this is spam botnet right now occurring, or is it something else? We were starting to think at the beginning that this was some typical botnet sending spam, but over two months, it has been a lot of time and I will show you right now this is still happening right now. This is we have right now a top of 45,000 queries right now. So it hasn't stopped, so we don't know if this is going to roll or if this is going to stop some day, but that is what is happening. So you have any questions, any comments?

Eberhard Lisse: Who hasn't asked a question today yet?

Roy Arends: Roy Arends, Nominet. In the past Mauricio and I have talked about this problem. We see the exact same traffic signature. We have the exact same problem. We're under load as well. It is significant. It is not a threat currently. We also have no idea where this traffic is coming from except from we do have a set of IP addresses, but we have no idea what kind of malware this causes.

One other thing we were thinking about – if it really becomes a threat to our business, we can actually on the upstream, block this traffic based on a very simple signature, based on the fact that the RD bit is that it says to one and based on the fact that the highest bias of the two-bytes identifier is zero.

Of course, that will also stop legal queries, so to speak, but legal queries in general shouldn't have the RD bit sets. Basically this is the last defense we can think of. So maybe this is a tip for others as well and I'd really love to hear if others see the same kind of traffic that Mauricio and I are seeing.

Paul: Hi, Paul (inaudible). If you can go back to slide 10 for a second… Yes, that first column where you have ascii to hex – those are all IP addresses. Yes, so you can just convert them hex to decimals. Probably Peter has a network order. Not that I know why they're used.

Eberhard Lisse: I was just going to ask do you know which ones?

Eric Ziegast: Eric Ziegast, ISC. You mentioned that you manage multiple Anycast clouds. From the clouds that you can monitor, do you see the same source addresses coming into multiple clouds, or do you see a good partitioning of the source addresses between the clouds?

Mauricio Vergara: Where it's standing right now, this, we've kind of seen some distributed things on every cloud but we have found this small… this small group of IPs that keep asking on every cloud. Where else do you study this?

Eric Ziegast: Are you working with people in the security community – the operational, basically people that are not necessarily DNS providers to do some analysis?

Mauricio Vergara: Okay, right now we are working with our local cert because they were searching for some kind of virus on their expanded… they were receiving… They were saying to me the last time that we talked with them that they have found two kinds of patterns on emails – one is a virus that procreates this botnet or something and the other one that sends the (inaudible) out. So that's the only person that we are working with on security.

Eberhard Lisse: Any other questions?

Peter Lauscher: Peter Lauscher, ISC. We've had at least two other ccTLDs I think both of them have actually popped up on DNS Operations that have had similar issues as Chile. I know I forwarded at least one of them to you saying, "You guys should talk," I'm trying to get the other one to speak back up because I believe he still has the same issue – it's exactly the same thing. High levels of MX queries again, ccTLD and so forth. And both of them are also, we secondaried for them as well so I'll try to make sure that they get in touch with you.

Mauricio Vergara: Okay. Thanks.

Eberhard Lisse: Alright, thank you very much. That was quite informative. Richard Lamb is the next and Bill Woodcock. They back at Clearinghouse have come up with something that might assist especially smaller and mid-sized ccTLDs in signing the zones.

Bill Woodcock: Richard and I will be kind of tag-teaming this presentation. Give me a second to get the video going here.

Eberhard Lisse: In the meantime, is Kelly Hardy in the room cause she's supposed to present next… Oh, there she is. Okay.

Bill Woodcock: It's a pleasure to be able to announce this in a roomful of so many friends. I guess we'll start out with Rick talking a little bit about the sort of goals in common that ICANN and…

Rick Lamb: Okay. I'll just get closer. Alright, just a little bit of the genesis of this talk, this presentation and the concept here. I don't think – hopefully it's not a surprise to any of the people in this room that one of ICANN's goals recently has become to accelerate DNSSEC deployment. It's my boss' goal so it's become my goal.

But in doing so, we would like that to happen in a way that still maintains, as the slide says, some level of security and trust in the system. Anybody can run sign zone or tools and do this. It really is the procedures and processes and structure, the security model around this that's important. So that's ICANN's goal.

So in looking at how to achieve this goal, I went through and looked at a number of options and found a number of people in the community that are trusted members that are perfectly positioned to do some of this and I think this is where maybe you could explain some of PCH's goals.

Bill Woodcock: So PCH's goals, of course, as a non-profit are to support critical infrastructure operators like yourselves and to sort of generally increase the global availability of the internet, that is, to bring costs down and reliability up.

So of course, we feel that DNSSEC fits well within the parameters of that general goal. And the way we conduct our mission is to try and stand up infrastructure that people can use and then help them build their own infrastructure through certain knowledge transfer and doing trainings and so forth.

Right now we do 60 to 70 trainings a year around the world but most of those are on the topics of internet exchange point construction and operation, Anycast, IPV6 and so forth. So we're very glad to be adding DNSSEC to that list of things that we try to help people with.

So basically our approach is to have a shared secure signing platform with knowledge transfer. That means we're standing up an operational platform that handles DNSSEC signing of zones for people who want to use it and any ccTLDs that want to use it are free to. And we will help them understand exactly how it works and understand how to replicate any and all functions of it and help them migrate from that shared platform to their own individual platform under their own ownership and control as they see fit.

So this task leverages a lot of existing experience within, obviously, ICANN in the part of Rick who built the root signing system and PCH, as we've been operating as many of you know a large Anycast network for a long time.

So I think one thing that was really important to both of us is that this be a best practices implementation that we do everything in the best way that we knew how so that it would be beyond reproach and so that auditors would be really, really happy with it so that anyone who goes and copies this system will likewise be beyond reproach from auditors within your country.

And lastly, as with all of our services, it's provided on an as-needed basis rather than a fee basis, so this is available to anyone who wants to use it to DNSSEC sign their ccTLD and we are not charging for this service.

Rick Lamb: We'll be getting into this, but one of the key points of this thing is that there's no locking here, that the whole idea here is that we want you to run your own system in the end. So one of the things that, of course, that brings up is modularity and we've built this in a way that uses a lot of separate components so that it is modular. There are different building blocks, different parts of the system can be used and we'll talk about this a little bit later. But, you know, whether you want to hold the KSKs and part of the keys or have some parts of the operation done somewhere else. So that's critical here.

There is a clear transition path like I said, from this platform to your own platform. We even have all kinds of checklists and everything to make this a process that will be flawless.

Bill Woodcock: So obviously the first benefit is that it sort of gives you immediate DNSSEC signing of your whole zone without having to go and do anything more than just sort of telling us yes and following a short checklist of kind of paperwork steps. It gives you security that's on par with that of the root signing process. It is technologically the same process.

The main difference is that whereas the root process is all performed within data centers in the United States, this one is spread between Switzerland, the United States and Singapore so that the process will not be subject to the legal dictates of any one country. And the countries were chosen obviously to be as neutral as possible and so that there would not be too much legal leverage that can be applied across all three.

Obviously a big benefit of this is that it offloads the expense of HSM as an operational facility and so forth. As we do with the Anycast system, we can build a single large system that's shared by many, many organizations and so the cost of that winds up being amortized across the benefit to many and obviously the actual sort of dollar cost to users of this system is zero. So this is in some ways a benefit to us that it's amortized many ways, but a benefit to you that you don't wind up having to go and build a bunch of special rooms and put a bunch of special equipment into them.

The benefit of the best practices environment is that if we're doing knowledge transfer, it's really important that we make the upfront investment to make sure that the knowledge we're transferring is, in fact, best practices knowledge and not just whatever is expedient to get the job done quickly. So for that reason, we've been building this system up over time and taking advantage of the three years that Rick built the root system over and that we provided Anycast back and for that root system.

And then, of course, as we keep reiterating, as you gain confidence in using the system and processes, you can transition over to taking over operational responsibility for yourself.

Rick Lamb: Just to clarify I just need to get certain things straight. That was… We worked with Bill designing cast system for the IN Net test bed. So the technology here is based on a lot of the work that was done on the root system, but we also did have this other long-winded, three-year effort to have this IN Net test bed as well.

So again I want to emphasize there's this clear path coming and going so if someone wants to work with Bill or work with PCH, we have this clear set of documents that will say, "Under the control of the ccTLD house, here's how we can get that ccTLD signed and published by PCH." In this case, of course, the keys would be held by PCH and this is all in high security, HSM stuff. No one ever gets the private keys. You've heard all this. Blah, blah, blah. So it's completely secure, but more importantly, is also the transition away from.

So when people have come up with their own operations and feel, especially the amount of experiences and built up an interest in doing this, we'll work, or Bill will work with cooperatively with that ccTLD to actually go through a checklist, clear steps in transferring material. It doesn't require, necessarily, if any of you guys have thought about how you do this, it doesn't require that private keys be shared or anything like that.

And in this case, of course, the ccTLD will be controlling both KSK and ZSK. And, of course, any transferred relevant information to maintain audit records, again, there's this whole paper tiger here that has to be fought and we would make sure that the same paperwork would be there as well. And there are, of course, variations of this as people may think where maybe the KSK is held by the ccTLD and maybe the ZSK and those mundane operations are run somewhere else. So there are variations just based on the modular architecture.

Bill Woodcock: So the signer platform is sort of big and complex as these things go, relative to a simple software-based one, but straightforward enough to explain. It's again based on the root signings design, uses Bind signing tools. The KSKs and ZSKs are kept in HSM's Meade FIPS 140-2 Level 4 which is the highest level of physical security, hardware security of the signing appliances.

We have fully redundant offline KSK facilities in San Jose and Singapore. The Singapore side is not yet built out, but we anticipate that it will be before the next ICANN meeting which will be there in June and fully redundant online ZSK facilities, the portion that actually does the signing in San Jose and Zurich. And again, the Zurich facility is not yet built out but we anticipate that it will be by June.

It's a Bump-in-the-Wire operational model. We suck in unsigned zone data from the ccTLD administrator. We perform the signing operation and spit out signed zones on the other side. The transition plan for moving from this to a platform of your own if you choose to do so or when you choose to do so consists of knowledge-transfer workshops – we'll send people out to help you understand exactly how we do it and what the operational choices and administer of choices you might want to make are so that you understand all your options.

This is very similar to the kind of workshops that we do for countries that want to get internet exchange points set up, for instance. Then there are these checklists that explain in a step-by-step way what each of the tasks that need to happen are and in what order and how they interlock with each other. And again, it's a complete solution so we've got all of the key management, all of the paperwork, all of that sort of stuff and it's all Open Source.

So, for instance, if you need legal documents to appease the auditors in your country, we've got them and you're welcome to them and you can change them however you like and reuse them and republish them and so forth. Standard Creative Commons License.

So the diverse locations, as I said – San Jose, Zurich and Singapore. And that's backed up by our global Anycast network so on the publication side, once the signed zone comes out, if you want to use our Anycast network for publishing that signed zone, you can see roughly, depending on how many are online at any given day, somewhere between, say, 65 and 85 locations around the world with Anycast servers that can publish the data.

So sorry, they're very small type on this slide so the PDFs may be more useful to you if you're interested in this later. Basically, what you've got is four pieces of operational hardware here. There are two that hold the offline KSKs and two that hold the online ZSKs. The top cluster there in Zurich has just the online ZSKs; the middle cluster in San Jose has two sets of hardware – one for the offline KSK and one for the online ZSK, and the sort of deep backup in Singapore has only the offline KSK, but in a real emergency that took down both San Jose and Zurich, it could be converted into an online signing facility by repurposing the role of that HSM.

So the HSM itself, as I said, is a FIPS 140-2 Level 4 hardware signing module. That and the signing server, the little machine that is taking the zone file in and doing the signing operation, those are sitting together in a Class 5 IPS security container which is essentially a safe that is built with thermal transfer and fiber in and out so it is able to run operational servers inside the safe while still being locked up tight to control physical access.

That, in turn, is inside SCIF which provides electromagnetic isolation or RF isolation and another layer of lockdown for physical security and that in turn is within a Tier 4 data center or some other mechanism for providing two additional layers of physical security against unauthorized access. So again, this is replicating the way the root zone is done and the requirements on ICANN for the root zone were sort of developed, I believe, collaboratively between ICANN and the U.S. Federal Government who again wanted to be sort of above reproach, so…

Rick Lamb: Following Standards, NIF Standards that are published and very well known.

Bill Woodcock: Yeah, so basically, we're just trying to follow best practices here, even if the best practices seem kind of ridiculously over the top in physical security. So let me sort of show you the little animated picture of how this works. The ccTLD hidden master is sending IXFRs; the IXFRs get buffered up and moved over into the signer which cranks away and puts a DNSSEC signature on it. Then that gets moved over to our outbound master which distributes it to the Anycast servers and the other authoritative slaves. Yeah, that was two hours of poking around on my laptop. So there's the static version that one could actually print of that.

So basically we can't be signing everything all the time because we have a kind of serialized pipeline in the HSM, right? The HSM can only be doing one actual signature at any given moment…

Rick Lamb: Or a very limited number.

Bill Woodcock: Yeah. Very limited number. In practical terms, that is the bottleneck – the HSM. The physical HSMs, the hardware HSMs have a lower through put than a software process running on a big, fast server. So the way we're doing this is we're accumulating a few minutes worth of ISFRs together before we do a signature over the zone as a whole. And then we move that out into our outbound master and convert it back over into ISFRs out to the Anycast nodes.

Now this is not quite what I would call state of the art today. Probably state of the art today is what Brazil is doing, but that required that they wrote all their own code. We are not a software development shop, so we are again just trying to use best practices and replicate what's done in the root right now and I think both for us and the root, it's a question of waiting until the generally used tools out there support signing on an ISFR basis rather than a whole zone basis at which point we will overhaul that core module of the system to sign on an ISFR basis which will be more efficient in terms of use of the HSM and at that point the signing latency between when we receive an ISFR in and when we send an assigned record out *via* ISFR, that latency will be reduced.

The other thing we can do, of course, as load picks up on the system is we can parallelize HSMs and we are indeed planning to do that as load picks up.

Rick Lamb: All these things are modifiable but currently the time frame we're looking at at these various operations are the batteries and HSMs are good for five years, so good time to refresh the HSMs then. And as far as the key ceremonies that we have to do to pre-generate again, following the same sort of fashion that we had done for the root, is to pre-generate these DNS Key RR sets, pre-signed DNS Key RR sets – and we do that once a year – pre-generate them in a key ceremony at one of those locations – Singapore or Zurich.

And then maximum ZSK role frequency – this is going to depend on a lot of things like the SO expiration of the zone, various other things like maximum DTL, but six months. And hanging around various cryptographers in their various meetings, you know. This is always a question – How long is a 1024 bit key good for? Some people will say, "Ah, don't worry about a 1024 bit key; they're good for another 10-20 years." But once in a while you'll get somebody who goes, "Oh, no, six months and that thing's gone." So anyway, it can't hurt. That's how we're picking those numbers.

Alright, I'll go right into key management. As stated earlier – I know we're running out of time – but key management is one of the critical parts of any one of these operations; it's not just signing the zone, so we have automated signature process that automatically signs updates, does ESK rollovers, does an integrity checking before publication as well. We've seen examples of that and we've seen how important that is but this does validation as well of the zone.

Of course, tons of real-time monitoring. I mean just, you know, monitor, monitor. I have to give credit to the .fr folk on this, they're tops. But, I mean, it's a very important lesson. Email alerts, of course, to the TLD operators – all of that stuff is in place. SKS generations – done offline. And by default, those are the primers we have been operating with to a 48 bit KSK and 24 ZSK and sec 3 opt-out.

Bill Woodcock: Okay, so mindful of time, I'm going to just kind of skip this side which is a little repetitive and go on to the live demo, which is our chance to embarrass ourselves in public. Okay, so really quickly here. Okay, minus two minutes. We're very grateful to Mohammad Al Zarooni who is in the audience here. You can see .ae here.

This is using the Swedish DNS check utility to evaluate the validity of the signatures. This is using our signing system but not using our Anycast so we have not pushed the sign zones out into our Anycast network; this is just looking at the outbound masters for the DNSSEC system. And here it is again for the .emirate IDM also showing that the DNSSEC comes up all green. And then this is the administrator view in our system, so this is what Mohammad will see when he logs in to look at statistics. And you see that because this is sort of in alpha testing stage for that domain, we've been playing around with the different parameters so this is the remaining validity of the signature on the zone and this is again for .ae.

And this graph shows log scale – the number of domains in the zone which I believe is about 350,000; the number of signed records in the zone which is about six or so; and this is the latency; this scale is being read on this side so it's clustering around a minute. So it's somewhere between 30 seconds and two or three minutes between when we receive an unsigned update in one side and when we push the signed update out the other side.

And then here's exactly the same thing for .emirate, and again you're seeing clustering. We weren't messing around with this one quite as much so clusters are a little closer together, and you can see again log scale on this side. So I think that's it.

These are the test phases to use this. Basically first, we assign the zone, you verify the validity on our signing system as you can see with two UAE zones right now. Second, we push that out to the Anycast servers and you let anybody who wants to, gets it. Third, you coordinate with your other authoritative slaves so they can grab it as well. And lastly, you put a DS record in the root and you'd be live.

Eberhard Lisse: Thank you very much. Again very interesting. Just one question. With my little zone of 2,500 and one signed domain, how do I push it from me to you?

Bill Woodcock: We just slave it. You just…

Eberhard Lisse: You're slaving it already anyway. How do I push the signing from my little provisioning slip onto a real one? We discussed this offline I think.

Bill Woodcock: Yeah, but yes, there's a ton of paperwork about that…

Rick Lamb: There's a process for that as well, to share, so that you don't have to give me your private keys.

Russ Mundy: Hi, Russ Mundy, Sparta. In your training course in particular, Bill, that goes with this, is there going to be information that makes the point that you should give protection and control and accuracy to the content of the zone roughly equivalent with what you're doing with the cryptographic mechanisms? Because many people see this and they tend to forget all about their focus which really needs to be the content. That's why DNSSEC exists.

Bill Woodcock: Yeah, obviously the big weakness in all of this is, regardless of how good we make the cryptographic processes and how good we make the audit trail and how good we make the physical security, if the ccTLD administrator allows their own system to be hacked or allows one of their own staff to be compromised or something like that and they hand us incorrect zone data, we will sign the incorrect zone data and hand it on. So that would sort of belie the whole purpose of DNSSEC but ultimately there are some things that simply have to be under the control of the ccTLD administrator, regardless of how much they might want outsource. So, yes, we will emphasize that very strongly.

George Michaelson: George Michaelson, APNIC. I think this is a really nice initiative, guys. I think you deserve a big hand clap on this one. The key escrow thing – stuff I've been involved with, I find people get unbelievably touchy when you get to the edges of the key escrow. So if you come up with some messaging on why this is the right thing to do and why this is an acceptable and safe and rational thing to do, I think we're all going to benefit from that because the kind of negative press you get when you talk about generating private keys for people is kind of weird. I expect you're going to have to think hard about some of that.

I really, really like that you've already put some of the infrastructure offshore international. I think that's a very strong message. Some of the human factors failures that I hear from other people, "Oh, yeah, we went to the alternate key box but we haven't updated the key set there so when we signed with that it was immediately wrong."

I think you're probably going to find that you'll still have exposure to those things cause you just cannot always cover for some of those process failures, but you've got a really nice system here.

Bill Woodcock: I think part of the way we're trying to address that is by being entirely transparent in our processes, by publishing all of these checklists and so forth and documenting our own following of our documented procedure so that if we forget to do something, one of our friends will remind us. That's what I hope.

Rick Lamb: And I would like to talk to you about how that messaging might be made. I think technically most of us in the room understand that there's safety here. Private keys are generated inside a box; no one can see therm. There's no way that they can ever be extracted. But how to get that across is a difficult thing, yeah.

Simon McCalla: Hi Bill, Simon McCalla from Nominet. There's going to be a lot of folks talking about signing services this week, I suspect. How do you see this service fitting in with some of the other offerings that are around?

Bill Woodcock: Obviously there's only going to be one key for any one valid ZSK at any given time for any given domain. We don't care whether we hold that or they hold that or someone else holds that. What we would really like to see if the transition to omnipresent DNSSEC be as quick and smooth as possible. We are a non-profit; we are not in commercial competition with anybody, so we wish everyone else who is doing this in any way the best of luck and we will happily support anybody who's doing this whether they're commercial, non–commercial or whatever.

Rick Lamb: And is ICANN there to help anyone who's interested in doing any of this stuff. I mean there's no favorites here at all.

Eberhard Lisse: Alright, but that will be the last question because we're running a little bit late but we've got a buffer.

Jay Daley: Jay Daley from .nz. One of the things I think that ccTLDs are beginning to recognize as they're implementing their own signing platforms is that probably the larger resource and skills issues tend to be with registrars and I think that a number of ccTLDs are considering how they might do things for their registrars such as running a signing platform for them as an interim measure.

And so the registrars then develop their own. Is that something that you thought about doing or involving through the ccTLDs but not independently with the registrars?

Bill Woodcock: Yeah, actually, we're already doing a pilot with (inaudible), which is Patrik Faltstrom's sort of full service registrar which is 100% DNSSEC-enabled, and they're already on our Anycast platform so they are not using our HSMs; they're using their own signing mechanism, but they're integrated for the DNS records and the Anycast and we're sort of willing to explore that with anyone else who's interested and I think Patrik is interested in seeing best practices – I mean, all of you who know him know that he spent his whole live pushing best practices through the ITF.

So I think we're very interested in seeing that level of integration. We're a slightly over 20 person shop so there is so much that we can handle at any given time. And so we're trying to get this out the door smoothly. And then we can take that feedback from the ccTLDs that that's what they need and begin working in that direction.

Eberhard Lisse: Alright, thank you very much. As we usually often do, we're going to have next a little presentation about what CoCCA, one of the more common registry platforms is coming up with. And after that the competitor will have the same chance to give us their point of view. Kelly Hardy.

Kelly Hardy: Good afternoon, ladies and gentlemen. I'm making this presentation on behalf of Garth Miller who could not be here today and it is on the updates in policy and technology for CoCAA.

[background conversation]

Eberhard Lisse: How do I do this with PowerPoint? Huh? How do I swap this with PowerPoint? Presenter view…

[background conversation]

Kelly Hardy: Thank you. I'll make this quick; I know you guys are thinking about lunch. Sorry about that, guys. I'm a policy person, not a tech person.

Okay, so the objectives of CoCAA v. 3.1 are making installing and upgrading a registry database simple as installing and upgrading your favorite office suite. Think of Open Office. Easy one-click installer for OSX, Centos and Open 2 available soon. We are installing the popular EPP IDN registry system as a bundle with [Post-gres] in less than 10 minutes and automated notifications, critical security patches and upgrades the latest version can be automated.

A recent policy and technology developments include allowing registrars to publish reseller info in the WHOIS, full historical abstracts purchased by public or accessed by law enforcement regulators, automated validation of registrant contact details by registry to retain activation, automatic trademark validation against CHIP databases; two new contact types – agent and DNS administrator; proxy registrations and DNSSEC. And now we've got a demo of our release candidate which will be available on the 16th and this is as easy to install as a Windows program and will go much more smoothly than the PowerPoint install.

This is our installation demo. You choose a folder to install the CoCAA EPP registry software. You can optionally register CoCAA EPP registry as a service and that way it will automatically be started every time the machine is started. Enter your IP address or a local host. Your password is eight characters long with two upper case letters. You do not want to use an existing key certificate. Apply your same password. And click through until you enter a user name and password. The software is now installing. Thank you, guys, for your patience.

The CoCAA philosophy is that you can retain local network engineers, build your own data center or co-locate, share registry software development and maintenance expenses with other users, fail over and disaster recovery is simple with a common software platform. You can choose the best DNS solution for your project and there are many options.

CoCAA and DNSSEC is v. 4.0, includes DNSSEC to be released at AFTLD in Ghana. We've decided against just signing the zones without adding the required registrar and registrant functionality. CoCAA v. 4 will allow registrants to nominate DNS administrators; registrars and registrants often don't manage the NS servers for a domain.

For more information, visit CoCAA.org.nz. Software v. 3.1 and new installers will be formally released on March 16. Our software is almost finished installing.

Eberhard Lisse: While we wait for progress to stop, I used to in L.A. three years ago, we had a shell script to do this for us and we're suffering quite a bit. I don't need to reinstall CoCAA all the time, but if I were to install it on a new system, that is much, much, much, much, much more easier. So as we all know, CoCAA is Open Source for ccTLDs only and not for gTLDs. So if anybody intends to implement it, that's a quite easy way of doing it. It has moved on now.

Kelly Hardy: And as simply as that, the software is installed.

Eberhard Lisse: And as you all know, CoCAA is a Java application so you exit it with your browser. It fires it up just now. It must accept the license.

Kelly Hardy: Thank you, guys, for bearing with me on my first PowerPoint presentation. Apologies for the delays. Okay, you accept the licensing. And you are completely ready to sign in and begin use.

Eberhard Lisse: This is the normal interface that even our version that we run in production shows, so it's actually… I run it on Ubuntu. Once installer becomes ready for Ubuntu, this is something that I really want to have in my armory to enable me, if I have a fatal hardware crash that I can even if I have to move to a different site with my hourly backup, bring this up immediately. I think this is a very helpful new step. Alright, any questions? Come on, guys. Alright, Kelly, thank you very much.

Kelly Hardy: Okay, thank you, guys.

Eberhard Lisse: And now Jaromir is a programmer. I went to visit them in Prague once I think four years ago and when he tried to help me how to get a FRED to run and I must say I was at that stage unable to manage it, but FRED has come a long way. It runs in .cz and some other ccTLDs and it works in production for a relatively large zone – 500,000 names. 700,000. And what is in particular impressive that since they're using this, they have sort of doubled their size within about two years or so. So this is a product that is also open, totally Open Source GPL and can handle quite a significant load and is therefore a product that some ccTLDs might look at and Jaromir will now tell us about the latest developments.

Jaromir Talir: So hello everybody. My name is Jaromir Talir. I'm a Technical Manager of .cz nic and following CoCAA's presentation, I was asked to give some presentation about our solution about a FRED registry system. So my presentation is about last year changes and features that we added and I hope it will be quite fast.

I will start with some description of FRED system and then I will shortly mention a new project, MojeID which actually drove almost all changes that we had to make for registry system. And then I will describe three new features or three major changes that we did and some plans in the end.

As Eberhard mentioned, the FRED is also open source registry. Open source for ccTLDs and also for gTLDs and it's based on open source tools like postgresql, apache, omniorb system, running in production for more than three years. It has a common set of features like EPP, automatic zonefile generation, WHOIS, DNSSEC. And key features are modularity of the system so different registries can install just different pieces of small pieces of the system and speed, so the main part is developed in c/c++.

Where is FRED now deployed? There are six production deployments around the world to Czech Republic, Angola and Tanzania. In the past and last year, three more deployments in the Faroe Islands, Costa Rica and Estonia and there are also some testing deployments. Right now I know about Albania and maybe there are some others. And the interesting thing is that these deployments almost everywhere slightly modify the system to suit their needs.

For example, in Angola and Costa Rica they don't have registrar systems or the way of managing registry, so they have just one registrar and they created their own web frontend for registrations with support of some validation of registration applications and so the FRED is used almost like a backend system like database system.

For example in Estonia where they recently started to use FRED, they did some slight EPP modifications like addition of some new arguments for create domain and update domain comments. So I think that the system is quite flexible for making modifications.

Our recent project MojeID was described by Andre during last ccNSO Tech Day meeting in Cartegena so I will just mention that this project was the answer for a question about what we can do with our contact database if we can somehow extend from domain registry to identity registry. And what we did was that we deployed open ID server over contact database with our domain owners and domain administrators and we are featured to be used just for validated contacts.

We do some basic validation like sending a mass email and a letter to contact data and we also do some extended validation by checking identity card in our office. And so the main changes in the FRED registry were done because of these new projects during the last year.

The first enhancement is in an auditing component of our system because at the beginning, more than three years ago, we started with a simple requirement to look at EPP requests and we solved this really simply just with one plain table for storing requests, but after a few days we had to really hundreds of gigabytes of data in our lock system and we started to have problems how to keep such amount of data and also this system was tied only to EPP and we wanted to look more interfaces like new OpenID interface and WHOIS for example.

So what we did is that we completely separated logging infrastructure to dedicated component with dedicated database and we redesigned database structure using partitioning mechanism and postgresql so right now we are able to keep our database table quite constant size. And we are able to maintain the huge amount of data much easily. And we also used different model for storing requests so we are now able to look data from any registry interface like WHOIS, administration and also OpenID requests. And we can easily add another link or features to some new interfaces.

The second component that we enhanced is messaging component because always contacts in registry needs to be informed somehow such events like domain is going to expire or some data are changed. And previously we had just emails notification and some simple PDF generator for snail mail letters that were unfortunately printed, completed and shipped manually from our client center. And we wanted a new communication channel to be used like SMS communication tool or domain owners. And we wanted to automate this snail mail expedition.

So, again we refactored the system of this component and we proposed some general model of message and communication channels and sending engines and we implemented two sending engines for SMS and snail mail letters because we could use HTTP interfaces to some companies in our country that are able to do it on behalf of us for sending SMS messages and also for completing snail mail letters.

What are some of the things to use this new framework back for our email messaging component but still I see there's a lot of promising possibilities like to use this framework for messaging systems like notifying about expiration of jabber or something like that.

And the last change or the last feature is about the contact data validation. So we had to add some new states to contact data and registry, like if it's validated in basic intermediate or full way. And we introduced also some new mechanisms for authorized changes, so when a contact asks to be validated, we first send some pings using email or SMS or a letter, and when the user fills these pings correctly, then the requests is processed and a new state of contact is set.

So those are the changes and you can see that a lot of these changes is under the hood, so we try to, with our experience, redesign the system to be more modular and more extensible. And from the features that we plan to add during the first half of this year is redesign a little bit invoicing model to be able to offer also postpaid, not just prepaid invoicing model.

And the second feature that we want to add is some annual reminder emails to contacts in registry so mainly the reason is so we will have some new communication channels to domain owners to inform them in one way about data that we store in the registry and maybe some of our plans and some of our new projects and like this. So that's all. If you have any questions for me?

Eberhard Lisse: Thank you very much. I didn't know you were increasing your market penetration, but one thing that I forgot to say is that a large number of domain names in the Czech Republic are signed, so huge proportions…

Jaromir Talir: Fifteen percent more than 100,000s.

Eberhard Lisse: Yeah, and not only that's important in itself, but also that FRED is designed to handle this intrinsically is quite a good sign. Any questions? Please. Anyway, we are seven minutes early. Oh, Stephane?

Stephane Bortzmeyer: Stephane Bortzmeyer. You mentioned IDN but you didn't give any timeline.

Jaromir Talir: IDN is sort of not too much well accepted in our country, so we don't have any plans to extend support for IDN in FRED. It's almost prepared. The only thing what is missing is checking some coding themes. It's possible to register IDN domain in FRED, but there are things that must be done to be used in direction, like configuring some card sets for different zones and the checking against these card sets. So right now we don't have any plans to implement these few missing features. So to say that we are completely full IDN and compliant.

Eberhard Lisse: On a point which is related to IDN but different, at the GNSO meeting two days ago, there was a discussion about internal summarized data with discretion allowing people who have (inaudible) access of the name to use their real names in registration. Is it supported by FRED or do you plan to do it for your colleagues with a carat on the letter on the name?

Jaromir Talir: Right now we don't have any plans for that way, so maybe in the future.

Lutz Donnerhacke: Lutz Donnerhacke. I have a question about domain transfers. If you have so much domains in your zone signed, you must have experience with inter-registrar domain transfers, especially in the very beginning of signing such a lot of zones. We have heard that a lot of zones share the same KSKs and it would be very interesting to see how to solve this problem.

Jaromir Talir: Yeah, exactly, it's about maybe 80% of those signed domains is using one KSK and actually I don't have any information because these domains are from some registrar which is also the poster and I personally don't have any information from them about their experience. With the transfers, we had some sort of experience that it happened that the domain owner transferred the domain or changed the name servers from these registrars and forgot to change also DNS keys so from the point of view of DNSSEC, they completely went out and we solved this by updating our EPP ITIL bed that when changing name server sets and you don't mention the change in DNS keys so we delete completely the DNS information just to be aware that this situation will not happen.

So it is our way of solving this problem. I have no information how many times it was used or whether it was ever used or not so the registrars wanted this feature so we implemented them.

Eberhard Lisse: Alright, thank you very much. We actually managed to finish two minutes before the time so I'll let you all go for lunch now. Please be back on time exactly sharp at 2:00 because then we'll start this update to Bind 10 and then we've got some very cool presentation about some measurement stuff from Les Cottrell. I haven't figured it out myself just yet.

Eberhard Lisse: Alright, so good afternoon. Welcome back. I hope you all enjoyed your postprandial depression now. And if somebody falls asleep, I will make sure they wake up again. But it's obviously not the speaker and we welcome Larissa Shapiro from ISC who will talk to us about Bind 10. I find this in particular appropriate because a few meetings back – how many meetings back – six, seven meetings back – we started a little bit of an initiative to support development of Bind 10 and some bigger ccTLDs and I recall fondly the Germans and the Canadians apparently pitched in. So now we can hear a little bit about what's happening.

Larissa Shapiro: Hi, I'm Larissa Shapiro. I'm the Product Manager at ISC and I'm here to talk a little bit about Bind 10. I will make one brief caveat. I am the Product Manager;

I'm not highly technical so if you have questions that I can't answer, I will refer you to a colleague.

Okay, you've seen this slide already from Eric, so I will give you the 30-second view. ISC is the Internet Systems Consortium. We are a non-profit public benefit corporation. We do a lot of different things. Most people know that we do Bind. Bind 9 is very widely deployed today as name server, but we are doing a ground up rewrite which is Bind 10. We also offer INC DHCP which is a widely deployed open source DHCP solution.

We offer professional services that support our software. We have public benefit programs that include our Hosted@ program which hosts a lot of open first projects. Of course F Root - many people in this room are aware of. We do a lot of work with DNSSEC and IPV6 protocols and my colleague, Eric, has been speaking here already about SIE and our security work. We're doing increasingly a lot of security work.

Oh, and I want to point out the little parrot guy on the right-hand side. That's our Bind 10 mascot and we did a contest out on the internet and we had members of our Steering Committee and ourselves went through and selected some nominees and then we had an election online and this guy was selected and his name is Bundy, Bundy the Parrot. And it was very important to our Program Manger that I point him out to you.

I really don't need this slide in this room because you guys will agree with me that DNS is important. But at ISC we think DNS is really important because it's what everything uses to connect to everything on the internet and because slow DNS really sucks and because with IPV6 coming, you can't actually just type IP addresses anymore, although most people never really did that. And because insincere DNS increases ways for people to attack other internet systems.

So the point of this is to say that we really need really good DNS and that's where we're going with Bind 10. DNS has come a long way since the inception of Bind 9. There really wasn't a lot of competition for Bind 9 when it first came around and Bind 9 is more than 10 years old. It's very much a monolithic phenomenon and it does still run the domain name server indicates that about 80% of the name servers in the world run Bind 9, quite possibly including yours, statistically. Although maybe not in this room; I don't know.

Anyway, it's a good piece of software but it's getting really crafty and other implementations have brought us a lot of new ideas and there are a lot of new concerns that DNS needs to respond to. I've listed a few of those here. Our colleagues who do NSD and Unbound are very effective at dealing with servers with small numbered zones and also I think the separation of the authoritative and recursive has been really influential.

And then we've looked a lot of PowerDNS because people want to hook up SQL. And then you've got all your applications that include Bind that offer various management layers that we're quite interested in. This is all leading up to where we're going with Bind 10. And there's also a lot of needs in the DNS world that haven't really been filled yet, plus more that I didn't mention.

So that's really where Bind 10 comes in. Bind 10 was… The idea was conceived as was mentioned earlier, several years ago and certainly Paul Vixie and Joelle Damas were very instrumental with some other folks in this room in coming up with the initial start-up funds for Bind 10. And it's a ground-up re-write of Bind and it's meant to be a solution for diverse DNS requirements. The project is now entering its third year; we're about to release the release for the end of the second year. And it's really a collaborative exercise with our sponsors in a way that is somewhat new for us.

So here's the five-year plan for Bind 10. In the first year we did the initial architectural work as well as the first deployment of the authoritative server. And in the second year, which is just winding up, we've done a lot more work on the authoritative server and we've also done the initial recursive server.

And in year three we're planning to become ready for production in both the authoritative and the recursive servers. But you'll notice that year four is what we call the “Drop-in Bind 9 Replacement.” And what that really means operationally is that there’s probably a lot of aspects that if you currently have Bind 9 in your network, you should be able to run Bind 19, but it’s really a stand-alone thing right now. You can’t just drop it and it won’t necessarily have all the hooks by the time we get to the third year deliverables. So by the end of year four, we should be there.

And then in year five, we’ll get to what we call the really fun stuff which basically means expanding the concept of what DNS includes and we’re going to do clustering support and that’s where we’ll get into the plug-in model and really trying to support people being able to make their own modules for Bind 10. We’re really hoping that that will be a big movement in our community.

A few of the Bind 10 architectural goals. Really key to the way Bind 10 works is the modularity. So there’s isolated processes – authoritative, recursive, statistics – there’s a whole lot of demons and they talk to each other through the “boss” process. And well-defined APIs in libraries that are also well documented and this is part of making the product extensible.

The full run-time control – the idea is that you should not have to restart; you shouldn’t have all your configuration in name (inaudible) anymore unless you want to or need to; it should all be dynamically generatable and on the fly. We're completely redoing the command line language. It's meant to be flexible and intuitive which should mean that it is good and works well. And actually we're working on a demonstration version of that to come out earlier in the next year and we'd really like to get people to try it and give us feedback before we do the full implementation.

So those of you who are interested in Bind 10, that information should be coming to the dev list and the announce list when we get that demo ready and we'd really like your feedback on that cause we'd like to make the command line tool that you want. Here's some more architectural goals. We're trying to really be customizable, but what we mean by out of the box, is that it should be not so difficult. So you can do a really elaborate customization and get down to a really granular situation, but there should be quite a few sort of one-button options, authoritative only, recursive only, slave or master only, dynamic DNS on or off and then we should have an agnostic way for you to use a variety of SQL backends.

And then this is what I referred to before. We've got the goal that we can customize through code changes that there will be an easy way for people to add modules and then hopefully share them in an open source marketplace and make them available. And that we really would like the APIs and everything to be sufficiently hackable that sys admins can add their own modules with a minimum of pain. That's a real goal of ours.

Some more architectural goals. So scalability obviously when Bind 8 came out it was… you could call it s single core. The system supported single core but it was just CPU. And then by 9 we support multiple cores (4 or 6). Bind 10 will support the reality of modern hardware – up to hundreds of cores and multiple machines clustered together.

Robustness – this relates in many ways to the modular idea. The idea is that reducing serious software bugs, what I really think is important here is that in the modular system we can reduce featuring… We had an issue a while back where there was a packet of death that could come around that we realized in Bind 10 wouldn't occur because there's not fate sharing and so you would only knock over one module. And I think it's not like we could always be less impacted, but I think that the more that the modules can stand on their own, the stronger the system is.

This is something that I'm really proud of about Bind 10. We've really made a shift inside ISC in terms of how we're developing in Bind 10. We've always been an open source company and we're still a managed open source company but we're really developing Bind 10 in the open. Almost all of our development conversation happens on the public developer list and is documented on the public track site. The addresses are there. Anyone can join the list, anyone can read the track site and anyone can read the source code at any time. It's in a [GIT] repository which you can access through bind10.isc.org.

And if you're just looking to follow the project, there's also an announcement list. The developer list is somewhat high traffic; the guys really use it every day to discuss issues.

The other thing that is really exciting about how Bind 10's working is that some of our sponsors provided us with developers, and so we're working as a collaborative team. We've got the ISC core staff – that's our program manager, Shane Kerr, myself, our account manager, Norm Mitchie and then our developers. And then we've got developers from JPRS. One of them is here today.

We've got developers and now also testers from CNNIC and we've got one developer from CZ.NIC. This has been a wonderful thing because we really appreciate the diverse opinions and also we get a lot of input from the operational requirements of the organizations that are involved and it's just been a really good experience.

The other thing is we're working on more community engagement in the open source community. We've applied for the first time to be part of the Google Summer of Code Project for Bind 10, so we're just trying to be as open and inclusive as possible while maintaining our usual adherence to our software principles and our quality principles.

Here's where we are status-wise. We are very nearly at the end of year two. Technically the Bind 10 years go by April 1. The project was started on an April 1 and so we go April 1 to April 1. So the release will be next Tuesday, March 22 and we'll actually be doing it from Prague because the CZ.NIC is hosting a Bind 10 developer summit for us during the week before ITF. So we're actually setting everything up ahead and then we'll flip the switch from Prague.

And we're doing developmental releases every six weeks. In this last year we've moved to a [SCRUM] development model and we're doing frequent developmental releases and we're really hoping that in this year we can increase the amount of early participation and feedback beta testers, if you will, who can look at those releases ongoing and give us feedback. We realize that through the first two years things have been in a pretty experimental state, but we're moving along to the point where that may become more logical.

So some of the features which are in the year two release – the biggie is the functioning authoritative server with Bind 9-like in memory performance database performance. That's actually a red/black tree implementation. And the functioning resolver – it's not a validating resolver. The authoritative side does DNSSEC now; the recursive side will be doing DNSSEC very soon.

And we have an SQL backend implementation as well. That was actually a year one thing but we've taken some of the quirks out in year two. And then we've got a statistics gathering channel with command line access that's been contributed by our colleagues at J. Paris.

And then, just in terms of moving forward, the year three goals we actually have separated out the production ready authoritative server as the goal for September of this year, 2011 and the production ready validating resolver as the goal of the end of the third year, March 2012. And I've got the URLs at the end there just so you know where to go.

Okay, this was my colleague, Shane's, idea. He said, "You should explain Bind 10 as this, sort of as being like early Mozilla." So you know how you got the Mozilla browser somewhere in the 90s, I'm forgetting what year now cause I'm losing track, but it sort of worked but you could see it was going somewhere really fantastic, you know?

And the thing with Bind 10 is, like that project, we're not intending it to be a single piece of software delivered by a company and then managed as a product. I say this as the product manager. But truly we mean it to be a community ecosystem phenomenon. We would like that so much to be a tool for the community of DNS users and cared for by ISC. But it's not ISC's; it's ours. And I mean ours, the people in this room.

So the software is getting ready for the next stage where we really want early adopters to come on board and start testing and really using the software. And we've got some folks already on board for that and I'd really like to talk to anyone who's interested. And we will be launching a formal test program over the next couple of months. That's critical for the next stage of the project.

How you can support Bind 10. We are a small and non-profit company. We really do need new sponsors. Quite a few of our sponsors are here and we can't thank you enough. Bind 10 needs more sponsors. We really do. And we need money. But we also need other forms of sponsorship. We actually have one sponsor who, due to their organizational bylaws doesn't give us money, but they give us an engineer. We gladly accept offers of staffing, testing, put a box in your lab, whatever it is. Or give me a list of your requirements; tell me everything that bugs you about Bind 9. I would love it. I love to hear that list. Sanity checks – I'm not kidding. And all of this stuff is essential to building the product that you need and you want.

I really mean this one – thank you to our sponsors. And many of you are here; I'm going to just read it out. Our founding and patron sponsors as well as our new sponsors, that's CIRA, JPRS, CN.NIC, Afilias, AFNIC, DENIC, Nominet, Ripe, Co.za, Registro.br, CZ.NIC and our newest sponsor is Google. We cannot fund this work without you; we are very honored to do it with you.

So a couple of upcoming Bind 10 opportunities to learn more. As I mentioned, March 22 is the year two developmental release. On March 30… Actually there's one more. If you're coming to IETF, there will be a Bind 10 presentation and a more technical one at the IEPG that Shane and I will do.

And then there's also going to be a box, and this is actually one about Bind 10 and we're also doing a ground-up rewrite of our DHCP. I know this isn't an entirely DHCP crowd, but some of you might be interested. It's going to use the same framework as Bind 10. So if you deal with both in your network, it could be interesting. And we really need support for that.

We're going to have a birds of a feather session at IETF to talk about that project. And then we're going to do a web seminar in April about the third year of Bind 10 and where we're going and soliciting feedback, so there's more information on the website. You can register for the webinar on the website. That's it. I'm open to questions.

Eberhard Lisse: Okay, thank you very much. I'm abusing the privilege first of course.

Jay Daley: Hi, Larissa. Jay Daley from .nz. Have you any idea what type of custom modules people might be writing for Bind 10?

Larissa Shapiro: Yeah, I think one thing that people are interested in is storing other forms of operational data in the DNS. There's organization store administrative data in their DNS. Also there's been some discussion of different forms of reporting, monitoring and reporting stuff. We're working on our own monitoring and reporting materials, but I think people are interested in that kind of thing as well. And truly, I don't know. I hope there will be things that I can't imagine coming out of that. We'll see.

Eberhard Lisse: You mentioned sqlite backend. MySQL, [Post]-SQL?

Larissa Shapiro: There will be. There will be both of those. The list of what we're going to implement for sure I think includes those and Berkeley database. But the idea is that the backend should be agnostic enough so that people can also apply other databases as they need to.

Eberhard Lisse: I'm asking specifically because FRED and CoCAA tools do zone generation so if we can do this with SQL it would be really cool. I don't really want to have to write to an sqlite.

Larissa Shapiro: You would prefer [post gres]?

Eberhard Lisse: No, I don't but CoCAA runs on [post gres]. And FRED runs on [post gres]. So a generic interface that you can plug in MySQL or post-SQL in. I think PowerDNS also runs on [post gres] so if you went to replace Power DNS which is the obvious aim as a competitor, you want to be as compatible as possible to make transition as painless as possible. Any other questions?

Larissa Shapiro: No. Alright, thank you very much.

Eberhard Lisse: Okay, at least, I, being a small ccTLD, I can lend moral support.

Larissa Shapiro: I appreciate that.

Eberhard Lisse: Okay, the next one is Steve Gibbard. He's now working for Nominum and he's going to talk a little bit about DNS infrastructure.

Steve Gibbard: Okay, you guys should be able to hear me now I think. So what I'm going to talk about here is partly building network infrastructure to support critical DNS infrastructure like ccTLDs and partly when to build your own infrastructure or when to use somebody else's. And got a little intimidated after sitting through all the great research talks yesterday, so I'm just going to stick a disclaimer on this right now and say there's no research here. This is just purely a how-to talk aimed at the ccTLD crowd, more so than the OARC crowd, I think.

I am going to focus on network infrastructure here and not go too much into what happens on the server side cause this is my area and I think a lot of you in this room can talk about the server stuff a lot better than I can.

So DNS, as you all know, is critical infrastructure. Without the DNS nothing else works. And authoritative DNS, well really any DNS, but authoritative DNS in this case, ought to be at least as reliable as the network that's hosting all the services that you want people to be able to get to because if your DNS is down, then having your network still be up gives you nothing. And as you all know, DNS is a hierarchy and so you need not only the DNS servers for everything that you're trying to work, but everything in the hierarchy above it if you're going to be able to find it.

So when you're thinking about how to plan this out, you're operating a domain, you want to know where to put your servers, your reliability is going to be best close to the authoritative servers. No matter how reliable the network gets, it's still going to break occasionally and the more stuff you have between your end users and the servers that are trying to access, the more likely it is they're not going to be able to reach that. Plus when you get stuff closer to them, you end up with faster response times and that's always good.

So in the case of ccTLDs, they're location-based and somewhat obvious where they should be reliable, assuming they're actually being used in the countries where they're intended to be used and not being sold as substitute gTLDs or something.

So you want them to be useable in the countries where they're based. There may be other countries that do a lot of trade with those countries or that are neighboring countries or something like that where they have to work. People outside may not care much if a country code half-way around the world goes down, for the most part. If users in one place lose access to a ccTLD half-way around the world, probably most of them won't notice. So it's useful to be conscious of where network partitions may happen. If there's a network partition it's good if your local communications keep working.

There are a lot fewer satellite connected regions of the world now than there were when I first started talking about this several years ago. There's been a lot of places that were previously satellite connected that have been building fiber in. But for places that are connected to the outside world by satellite, sometimes they'll see that satellite activity go down for a brief period several times a day when something goes wrong. And so if you're in a place like that and you're relying on outside infrastructure, that can get really annoying.

If you have a region that's connected by fiber, then outages get a lot rarer because the fiber is a lot more reliable, but if it's a single fiber coming into the region, and that breaks, the outage may last a lot longer than it would for a satellite connection to go down and back up.

Local phone calls tend to work without international connectivity. It can be somewhere in some poorly connected developing country and call across town and it's pretty much going to work. But then you go try to make an international call and that can be difficult. There are a lot of places, however, where even though you can make local phone calls across town, the moment you start relying on internet infrastructure, then anything going across town starts depending on the international links and becoming as unreliable as the international communications, and that's something worth avoiding.

So notable incidents that are a few years old now. Sri Lanka – 2004, there was this fiber cut in the Columbo harbor and it got described as an outage of "internet and long distance phone service," implying that the local phone service still worked and seems kind of embarrassing from the internet perspective. And in their case they had their ccTLD hosted locally but they didn't have a local root server and it looks like they do now. So, hopefully, if that were to happen again, then at least some of their local connectivity would keep working.

Burma/Myanmar in 2007 – that was a governmental cutting off activity but apparently stuff kept working inside and things in the .MM domain worked inside but not outside. There was a lot of discussion recently about the situation in Egypt where the government cut them off and what infrastructure could have been built to keep Egypt's internal activity going in the face of that. I guess a government shut-off is a little harder to plan than something caused by a technical problem because you don't know how far the government is going to go to cut things off. But this is another case where the more infrastructure you have locally as well as perhaps externally out of reach, can help survive that sort of thing if that's a goal.

Root servers are now pretty well scattered around the world. I started talking about this stuff years ago and there were some pretty vast areas of the world that the root servers were not covering well. And the root servers have done a really good job of doing a wide build-out since then, although if we look at this, there's still a bunch of Africa and parts of South America and a few other populated places that are still pretty dependent on root server infrastructure far away. So there's still some work to be done there. And if you're in one of these regions where there isn't good server connectivity, then I think a lot of the root servers out there would be happy to work with you on it.

So now we go on to what do you want to build for your own ccTLD? And what are the goals, how do you build it, what's the topology look like, what sort of redundancy should you have? We covered a bit earlier on who are you trying to serve and whether it is just your local users or whether you do have markets elsewhere that need good reliable access to your ccTLD.

And then something else to consider is how is your region connected – everything going through a central exchange point in which case you've got an easy location in which to put servers to serve all your local users or do you have a bunch of different ISPs with different kind of activity off to other parts of the world and very little connectivity, in which case you may need to go put something in each ISP or go find some common point upstream where you can put infrastructure that they can all reach.

And then whose infrastructure do you use? You can build your own and I'm going to later in this go into some technical stuff about how you would go about building your own infrastructure, but you may not need to do that. There's a lot of infrastructure already operating for the benefit of ccTLD operators and others – ISC, PCH, RIPE and a few others provide that for free. And then there's several commercial Anycast operations, including Nominum, who I work for.

And then there are a zillion free unicast options out there. If you look at the current list of servers that are doing authoritative DNS service for countries from top level domains, it's about 600 and something of them and a lot of those are various universities and people like Randy Bush and things like that who have some servers that are operating a few of those on.

So if you want to get a good comprehensive infrastructure with some redundancy so if you lose one of your providers everything else will keep working and cover both good global coverage and the areas that are useful to you, you can mix a bunch of these other sources together and get a pretty good large scale global build without having to go deploy too much of your own.

You may find that none of the easy outsourcing options cover the areas that are most important to you. So you may still need to build some of your own infrastructure to fill in that gap and then use somebody else's global footprint. So there's a bunch of ways to do this depending on what your needs are, but generally you don't have to do a whole bunch of your own infrastructure to make this work.

But what if you are going to build your own infrastructure? Where do you put your servers? Ideally you'd have some central location like an exchange point that all the ISPs are really well connected to and you can stick stuff there and have it be able to reach everybody. And some places are better provisioned for that than others.

As an alternative you may have a bunch of different competing ISPs put a server in each ISP then each ISP's users can get to something local and you can deal with replication and that's something where performance isn't terribly important.

If you're somewhere like Latin America where there's relatively little internal connectivity compared to the amount of connectivity that goes to Miami then putting stuff in Miami where everybody can get to it reasonably quickly, has some advantages.

And then there's a question of how do you serve the rest of the world. And partly you're going to want to serve the rest of the world because people outside may want to still get to your local content. And partly if you put servers of your own elsewhere in the world, especially if you're doing things in an Anycast setup, then perhaps you can sink some DOS traffic directed at your servers somewhere where you have more capacity and where it's going to have less of an impact on your core users.

So if you have a very limited bandwidth international link connecting your core region with the rest of the world, there may be something to be said for doing a two-location Anycast – one in your region and one at the other end of your international link just to soak up that outside traffic.

So there's an issue of do you do unicast or Anycast services? Anycast services are a bunch of DNS servers or a bunch of servers of any sort really, with the same IP address scattered around in a bunch of places. A unicast service is a single server with a single IP address or a bunch of servers behind a (inaudible) with a single IP address in one location.

And this is mostly an issue of scale. If you're doing two or three servers or something like that, you want users to be able to hit a couple of different IP addresses in case there's a routing problem with one of those IP addresses. So you want each of those globally visible. If you're doing 60 or 70 servers or something like that, then well you can only get 13 NS records in Deer Zone and that's going to be a problem. So at that point you need to do anycast, you need to relay them in order to make them all useable.

So even if you are doing Anycast, having several different service addresses in different places is good for reliability. Let's say you have two Anycast clouds with different IP addresses, and then you bind both of those IP addresses to the same server somewhere and that server develops a problem where it's still announcing its routes but it's not responding to queries that come in. And in that case, anybody hitting that IP address is going to see that it's down, they're going to fail over to the other IP address; they're going to see that that's down, too.

So if you have multiple IP addresses being used for your DNS servers it's good to distribute those in different Anycasts, so you'll have one set of servers that has one IP address and another set of servers somewhere else with another IP address. Unicast configuration in contract is pretty trivial. You take a server, plug it into a network and it'll work, and at that point you just have to make sure that you have enough capacity and that you don't have all of your different Unicast servers sitting in the same place where they'll be taken out by the same network incident.

So let's say you are going to build an Anycast network – there's a few things to keep in mind, and as far as keeping your traffic low-call (inaudible) and making Anycast apology makes some sense. Backbone engineers are often really good at keeping traffic local. You don't too often anymore see trace routes where you're trying to get from say San Francisco to San Jose and it goes through Virginia on the way there, which 15 years ago was actually pretty common.

We still see though that Anycast DNS operators are often not so good at this and do end up in a situation where you've got a bunch of servers very close to where the queries are coming from and yet the queries all get routed off to somewhere completely different. And that usually happens because people start plugging servers into you know, they have say their San Francisco node and their Virginia node, and their London node; and they plug each of them into a different network, hoping to get some network diversity out of it, and end up running into local preference issues where network backbones usually prefer to send traffic to their customers rather than to a peer even if the peer is closer. So you end up with stuff going off in the wrong direction.

So I'm sorry to pick on VeriSign here but there are four local j-route servers in the Bay Area according to www.rootservers.org. And I just a couple days ago was trying to hit those local servers from four different network points I had access to in the Bay Area, one of which is actually this room. And I wasn't able to get to them. My home connection had upstream connectivity from Level 3 and my traffic to J route ended up on Level 3's network and then ended up in Seoul.

From my office we have upstream connectivity from Global Crossing and we, our traffic to J route gets on Global Crossing's network and ends up in Mumbai. And I have a server co-located with PCH which has peering with Asia Netcom and I ended up in Taipei. And then from this room I had, the closest luck I had in any of this testing, this meeting appears to be connected via NTT and it ended up in Seattle. So I assume this has been this way for enough years that I assume VeriSign has some goal in mind with this. But if you are trying to get nice, fast answers, that's something to watch out for.

So you can get around this. Anycast can keep traffic local. What you have to do is get consistent transit from global ISPs. So you put stuff in a bunch of different locations scattered around the world and you get the same set of transit providers everywhere, and then they'll do hot potato routing to you, handing it off at the closest point and your traffic will stay local. You do want to be careful with this that you don't just get your transit from one global ISP everywhere, because then if that network does have a problem you will become unreachable; and if you've got enough locations you may be able to say "Okay, well half our locations are going to get transit from one, half are going to get it from another, and then we should be able to provide reasonably close responses" no matter which of those transit providers they had.

Sometimes you need to put a DNS server in a location where your global transit providers aren't available, and if you're doing peering with your Anycast system as well as getting transit that can work pretty well. You just, you peer in that location and hope that if your transit provider isn't available there you won't have too many customers in that location who are using your transit provider. So if you peer with pretty much everybody in the region you'll do reasonably well with keeping traffic local.

You need to be careful there as well, in that if you start peering with somebody in one region you need to peer with that network in other regions as well. Several years ago I was running the PCH Anycast DNS network and we ran into an issue there where we had peering with I think it was Telecom Italia in the US but not Europe, which was I believe Telecom Italia's preferred way of doing things at that point because they were trying to protect their European market from outsiders. And we started looking at where our queries from various parts of the world were hitting our network, and we noticed that there was this big ring around the Mediterranean in which both North Africa and Southern Europe where Telecom Italia was pretty much the dominant transit provider for just about everybody.

And all of our traffic from basically every country in that ring were all ending up at one of our locations I think in New York which looked like a performance problem given that we had a bunch of locations much closer in Europe. And so we went back to Telecom Italia and said “Oh, wait a minute, this isn’t quite working. If we’re going to peer with you anywhere it needs to be everywhere. Look where your DNS traffic is going, look what this is doing to your performance.” And they fixed it and traffic started working the way we expected it to be.

So if you’re going to do that be very careful that if you peer with somebody you peer with them in all areas where you overlap. And don’t take transit from non-global providers. I found on the PCH network that fairly often we would install something somewhere and a bunch of local networks would come in and go “Please, we want to be your transit provider! Look, we’ve had this great connectivity to the rest of the world!” and the moment you do that you start seeing really suboptimal traffic flows and that can be bad. We had one case where somebody in Perth in Western Australia decided they wanted to be our transit provider even after we said no and started announcing us out to the rest of the world via somebody who was a customer of a customer of a customer of [Status'] network in the US. And suddenly anything in the US that hit [Status] was going to Western Australia and performance was really bad.

So always insist that your non-global providers that you're connecting to treat you like a peer, not like a transit customer. So this is what it should look like. You have a bunch of cities scattered around. In this case we have Hong Kong, San Jose, Ashburn, London, Sao Paulo and Mumbai. And we have the same set of transit providers in Hong Kong, San Jose, Ashburn and London, who we're calling Transit A and Transit B in this example. Those transit providers probably don't exist in Sao Paulo and Mumbai so we haven't gotten transit from them there.

So our Mumbai node has peering with an Indian peer who, because their Indian, all their traffic is from India, they're sending all their traffic to Mumbai – that works pretty well. Our Sao Paulo node has peering with a South American peer who also has connectivity up into the Eastern US and they end up hitting us in Ashburn as well because we overlap in both locations. And again, so in this case the South American peer will never end up sending traffic farther away than Ashburn and hopefully to Sao Paulo, because hopefully they've got their metrics set up so that that's closer.

The Indian peer should never have to send traffic further away from Mumbai but those sites won't negatively impact the rest of the world.

And if you do this, these are plots of connectivity, plots of which country's query sources ended up on which global servers from the PCH network from probably four or five years ago, and that shows that if you do this there's a Paolo Alto node that got most of its traffic from the US and Latin America, and a little bit from Australia and Asia; there's an Ashburn node that got North and South America and some of Africa and it looks like it got Turkey for some reason. I forgot what the story was in that case.

And then the London node got most of Europe and a bunch of Africa; Hong Kong node got a bunch of Asia, so basically the distribution that we were hoping for there. If you're going to do this, if you're going to do Anycast networking you need to do some work with routing protocols. You're going to need to do your upstream peering with BGP in order to announce the routes and you want to keep the routing very consistent, again to keep traffic going into the right locations.

So having a single global AS can help keep things consistent. You want to be careful about your BGP attributes. I saw a case fairly recently where there was a multi-vendor network with two different router vendors who were doing different things by default to the BGP [meds], and because there wasn't a route map resetting one vendor to the other's default or anything like that, that was causing non-optimal behavior in the upstream providers. So be careful about BGP attributes.

And then don’t propagate the Anycast routes between sites. It’s very tempting to go, you know, try to build this (inaudible) connect bone and get you know, OSPF running everywhere and announce your servers from one site into your OSPF which then translates, which then moves it over to your other sites, and then you end up with one problem in one side that can just spread across your IGP and cause problems for your stuff and the rest of the world, too. So the more separate you can keep things the more you can isolate problems.

Internally you’re also going to have to do some routing because you want a site to be able to drop the routes when the local servers go down. So you can use iBGP to talk between your servers and your routers; you can use whatever your favorite IGP is, OSPF, ISIS, whatever. And you can, there are a couple of ways to originate the routes on servers. You can use Quagga or [Bird] as dynamic routing tools that you run right on the servers and you could also use a load balancer that monitors whether the servers are available and injects routes as need.

If you’re trying to pick a routing protocol to use internally for this I think you’ll find that OSPF has a lot wider support. There’s a lot of load balancers for instance that do OSPF but not much also. Once you start dealing with IGPs that aren’t intended for routing straight to the outside though you lose a lot of filtering capability and a lot of control, so BGP can help out with that quite a bit. Most of the at least smaller Anycast DNS infrastructures I’ve seen have just used routing software on the servers to announce the routes out to the routers. You can also use dedicated load balancers in front of the servers and those can monitor the server availability, do the routing protocol stuff, spread stuff among local servers.

You don't really need those to do load balancing because if you announce a bunch of routes that look identical into most routers, they'll just go round robin or hash-based load balancing or something between them and will work perfectly well for load balancing. But the load balancers have other features that can be useful like rate limiting connections from a source that tries to send you too much stuff or something like that. So those are sometimes useful.

Again, if you're mixing load balancers and Quagga or something like that, be careful about your routing attributes because it may be that it starts the routes in entirely different ways and one of them looks invisible as long as any of the other ones up anywhere.

So then of course, any time you're building infrastructure you want redundancy. More servers are better than fewer as long as you don't get so many that you can't manage them. As far as getting redundancy between your own DNS infrastructure and somebody else's who you're outsourcing to, there's no contradiction there. Just use both and then you get the best of both worlds. You get your local control over the stuff that you really care about and you get somebody else dealing with the stuff that's further away.

You do need to do a fair amount of monitoring – check your zone serial numbers and all your servers. If you're using Anycast, you probably want to monitor your individual unicast management addresses on the servers and make sure they're getting up to date information rather than just what Anycast server your monitor system has to see. And then also to tell you if something's way off in your Anycast routing can be useful to check response times from a bunch of locations. And there are a bunch of commercial services like [Pingdom] for instance, that will do monitoring from a bunch of places around the world and give you some information on that.

If you're doing – I'm being told time here. I think I got two slides left or something – if you're setting up an Anycast network, there are some things you need to get set up. You need your servers running Quagga or BIRD if you're not running load balancers. You need BGP capable routers obviously; you need IP transit from consistent providers in all the sites; co-location space everywhere you want to put them.

And then the really sticky issue is if you want Anycast routes that you can announce and withdraw in each location as your service becomes available or unavailable, you need a /24 of address space per site for this if you're using multiple transit providers. And that can be a little tricky to get, although at least in the ARIN region, if you qualify as critical infrastructure, that gets easier.

What should this look like when you're done? I've picked on Nepal's .np domain as this example for their set of name servers. As you can see, they've got ISC; they've got AP.NIC; they've got PCH and then they've got a bunch of local stuff… oh, and RIPE, and then they've got a bunch of local stuff to make sure that their local users are really well served. So some really old paper for further reading and I think that's it.

Eberhard Lisse: Thank you very much. Lots of information to digest during your presentation but of course we uploaded things. One thing in developing countries in Africa in particular, it doesn't matter whether you have got a system next to the other country. All local ISPs don't PSO. Virtually all of them talk to South Africa before they talk to each other. All the routes go to Africa. Sometimes if there is a problem, it goes to Europe before it goes to…So there it makes more sense to put actually the F-root in Java is actually closer for us than putting one in our own things. But if we have one in our core location we can convince them to peer much easier, I would think. Any questions? We're running a little bit behind the time.

Les Cottrell, from the Stanford Linear Accelerator Center, will tell us about pingER.

Les Cottrell: I'm Les Cottrell. I'm from Stanford Linear Accelerator. It's not too far away. It's my first ICANN meeting so I'm a novice, so don't be too mean on me. I'm going to talk about network monitoring so as Monty Python says, "It's something completely different," but I think you'll be interested anyway.

So I'm going to talk about a very simple way of monitoring using ping which exists on every machine that we pretty much ship these days so it's pretty easy to do. But we've been doing it since 1995 so we've got a long history of showing what the internet has been looking like.

The work has been led by SLAC, but it's also been done by other high end physics establishments, in particular Fermilab in Illinois and the International Center for Theoretical Physics in Trieste. The other group that we're working with is in Pakistan and they've been very helpful and they're funded by the Higher Education Commission in Pakistan.

So what I'm going to do first of all is kind of give you the mechanism by which you make the measurements and then go on to show the results of the measurements as seen over the years and also then finally wind up with some case studies. One of the is on Japan, which is kind of interesting today because of the earthquake. Another one is on Egypt and the Middle East after a cable cut in 2008 and then there's another one which I don't remember at the moment but as soon as I see the screen I'll remember it.

So if you can advance the next slide. So I will go over what we measure, the coverage, what we find, also relations to other things in particular, the human development index and then some case studies. So if we could go on to the next slide.

So the way it's done is there is a monitoring host which is shown by the laptop at the left, just to show that it doesn't have to be anything specific, which sends 10 pings to a server somewhere in the world. The server then responds with the 10 pings coming back, assuming everything's working. Then the data is saved locally on the machine making the measurements and then they're uploaded daily to a repository – there's actually three copies of the repository – one in Pakistan, one at Stanford and one in Fermilab.

From this we measure the round-trip time and the losses and from that we can get a lot of information. So if we go to the next slide, I can give you an idea.

So this shows you the coverage. It's a bit dark. I don't know if you can actually see it. We have actually 70 sites making measurements in 23 countries. There are four in Africa – it's been very hard to get sites in Africa. We also have what's called beacons which means that every monitoring site is first to monitor the beacons and then there's roughly 740 remote sites which are monitored. They don't need any software, they just have to respond to pings. We cover 50 of the African countries and the countries that we measure to are contained 99% of the world's population. So as I said, we measure RTT and from that we can get jitter. We measure loss; we also measure unreachability – I'll come back to that in a minute.

And then we can derive from these measurements the throughput, the mean opinion score which is used for voice-over IP and for other things like that and the directness of the links. Next slide.

So the simplest thing is to measure the round-trip time so that the left-hand axis here shows you the round-trip time in milliseconds and the bottom axis shows you various hosts in Pakistan. The light green shows you the average round-trip time and the dark green shows the maximum round-trip time. So you can see that there can be big differences – some hosts have very little difference between the two and that shows that they're not very congested. Other hosts have very big differences and that kind of gives you an idea of the congestion.

One can also measure the jitter which we do by measuring the inter-packet delay variation. So if we can go to the next slide. So this is the losses. The bottom axis shows you the time starting in 1998 coming through to 2010. The left-hand axis shows you the percentage lost.

Now losses are very good in the sense they also give you a good estimate which does not depend typically upon the distance that you're measuring to because typically the problems are at the end nodes, at the edges. So it's distance independent.

As you can see, the left-hand axis is exponential and you can see the lines if you draw roughly straight lines as drawn on the graph. Obviously there are a lot of wiggles in them. But if you've put an average to them they'd be pretty straight. So what you're seeing is an exponential improvement and roughly a factor of 112 years improvement in the losses.

The best countries or best regions are almost obviously, North America, East Asia, Europe and Australasia which you're now seeing on our ridge less than .1% packet losses measured by ping. The worst are greater than 1% and they are Africa and Central Asia. So we go to the next slide.

Here's another measure. This one, if you send the 10 pings, you send in 10 pings separated by one second, if none of them respond, we designate the host as being unreachable at that time interval. In this case we chose to make the measurements from a reliable host which was at SLAC and we measured it to Pakistan. Along the bottom axis you can see the host and along the left-hand axis you can see the percentage unreachability. And you can see there's enormous differences. Some of the host which is to the right have very large unreachability; they're just not available for large percentages of the time.

If you look at the graphs to the bottom left, they are smoke ping type graphs and you can see some of the hosts that just were not available at all. The black shows you we got no response. So what is happening there is that there are big problems with power in Pakistan, there's a lack of oil – there is some oil in the eastern province of Baluchistan – but in general there's no oil to create the power itself and their budgets are very low. So what happens is the universities have power outages and they do not have UPS and so this is what we see when we're monitoring them.

So we can also take the losses and the round-trip time and then knowing that most of the traffic certainly for the last 10 years or so has been TCP Reno and knowing how the congestion algorithm works in TCP Reno, we can derive a rough estimate of a throughput which is roughly eight times number of bits in a byte times 1460, the typical length of a large packet, divided by the round trip time and the square root of the loss.

This formula was devised by [Matthew Sitel]. So again on the left-hand axis you can see the derive throughput and on the bottom axis you can see the time starting in 1998. And what you see is that for Europe and eastern Asia and Australasia, you've got very good throughput. That's the upper lines. Rather than showing all the wiggles, in this case I just fitted exponential fits and the left-hand axis you'll notice is a lock scale.

Then Australasia is catching up and so is eastern Asia catching up. If you look at Russia and America and then take the throughput today which you'll see is roughly 1,000 kilobits per second or megabits per second, they are roughly five to six years behind American and the Middle East. Southeast Asia is nine years behind and if you go to Africa and then I extrapolate back, you'll see it goes back to1992, so Africa is roughly 18 years behind where we are in the western world.

So if you think where you were in, say, 1992 using modems and things like that, that is the situation that you're seeing in Africa. Also you'll notice that bottom line which is orange and is marked towards the right, actually it's red, but labeled Africa, is not only behind, its slope is less than the slope of the other line. So it's falling further and further behind so that within 10 years, it will be 150 times worse than, say, to Europe.

Another thing we can derive is something known as the mean opinion score which is where you sit somebody down and you let them talk to somebody else on the phone and slowly you turn up the noise and the round trip time and things like that and they tell you an answer, "Oh, this is perfect," or "This is awful," "This is like talking on my cell phone to someone who isn't speaking the same language as me," or whatever.

They come up with numbers where five is perfect and one is just you can't hear anything. And again you see here – this is the MOS score – one through four and a half in this case. The five is obliterated in this case – and you'll see going back to 1998 that actually most regions now have pretty good connectivity. South Asia is not so great, although we regularly have Skype meetings with people in Pakistan and they work reasonably well. But you'll see that Africa is still in very poor shape.

We've also looked at how does this correlate with social activity? And what we see here is if we look at the blue box to the top left on the left-hand edge, this was when there was a vacation. It was one of the religious holidays in Pakistan. So you can see the utilization was very, very variable. The left-hand axis shows you the round-trip time going from 300 milliseconds up to 700 milliseconds, very, very variable. Then when all the students went home and everybody went home for their holidays, the round-trip time became more stable. Then they came back again and you can see various other things happening. There's days when they went on strike, which is when it's very low. The strike is the second blue box to the left. You can see then when they all went on strike for various reasons, things got bad. So you can actually correlate what's going on. You know what to do if you want to fix the throughput, just send the students home and don't do any work.

Anyhow, next slide. So another thing we can look at is the directness of the connection. You all know the speed of light in fiber is roughly two-thirds velocity light in a vacuum. So if we take 300,000 kilometers per second as the velocity of light, then the round-trip delay in kilometers we've given as being equal to Alpha times the minimum round-trip time times 100. So if we know the distance between the two hosts, then we can derive Alpha from it. And in our case we do know where the hosts are, so by measuring the round-trip time, knowing the distance between the hosts, we can get an estimate of Alpha.

Large values of Alpha, approaching one, mean that you have a very direct connection. Small values of Alpha, typically .2 or maybe .3 or .4 or .5, mean the connection is kind of roundabout. If we can go to the next slide.

We can see again this is to Pakistan and the direct links here, if you look at the Karachi and Lahore, Karachi and Islamabad, and Karachi and Peshawar, they all have very large values of Alpha. So they are directly connected to one another. The one on the right-hand side, which is Islamabad to Quetta, if you look at the map at the bottom right, you'll see Islamabad towards the top and Quetta is on the east side of Pakistan. And you would think that the traffic would go from Islamabad to Lahore and then across to Quetta. It turns out the map is incorrect. We were given the map but actually the link from Quetta to Lahore has not been put in, hence the traffic goes all the way down to Karachi and then down to Quetta, which explains why the Alpha is very low in that case. So we can begin to get an idea what the traffic actually does as opposed to what we've been told it was supposed to do. Next slide.

Okay, so we can also look at how does our measurements, particularly in this case, the derived throughput, which is the bottom axis, which you can see is again a log scale, goes 100 kilobits per second, a megabit per second, 10 megabits per second. Plotted against the United Nations Human Development Index, which is a measurement of the length of people's life, their knowledge base, in particular what fraction of the population travel – primary, tertiary, secondary education – and a decent standard living, as measured by the gross domestic product per capita.

So you get a number from the UNDP Human Development Index between zero and one. You also get another number for the throughput. And what we've done is done is for various countries. The size of the bubbles are the populations of the country, the colors of the bubbles are where that country is by region. And you can see that there is actually a clear correlation between the UNDP Human Development Index and the throughput, which kind of gives you the idea that two things going on.

One is if you have a better Human Development Index, you're probably going to have a better internet and vice versa. If you have a better internet, you're probably going to be better able to develop things and have a better standard of living. Next slide.

Now we're focusing a little bit on Africa. The left-hand axis here of the bar chart to the left is the minimum round-trip time. In other words, we sent 10 pings and we looked at the minimum of those 10 pings. The bottom axis are the various countries which I've labeled as you can see.

And what you see – there's a big gap. Some of the countries – the dark blue – which I connected by geostationary satellites – are to the left and they have very large round-trip times, about 450 milliseconds. The countries to the right are connected by terrestrial links and have round-trip times typically less than 350 milliseconds. If you looked at Africa at the beginning of 2009, you see the map on the right. The red shows you the countries in Africa which at that time had links that went by, well, it showed you the minimum round-trip time of greater than 400 milliseconds – this is dark red. And you can see that central Africa and eastern Africa at the time had very poor connectivity and almost certainly were running *via* satellite links.

If we go to the next slide, you can see what happened. In 2008 there was one fiber that went round Africa. It was the SAT3 fiber and it was very expensive. They priced the cost of connections based upon the prices to connect up *via* satellite, by geostationary satellite. It was a consortium and there was very little competition. And then in 2010 the World Cup happened and there was a mad scramble for a lot of people to put fiber through Africa, both down the east and west coasts from companies like Seacom, TEAMs, the East African whatever it is system for – Subsea Fiber System, to put fiber down and they succeeded. And so what we see now as we go to the next slide…

So what we see is the impact here. As the sites were able to move their routing from the geostationary satellites to this terrestrial connections, what we saw was dramatically reduced round-trip times typically from 700 milliseconds – this is infinite sight, mind you – to 350 milliseconds seen immediately. So what you see here at the bottom is a smoke ping graph. The background shows you the losses; yellow means there's 10% losses; light blue means there's no losses.

You can see that to the left, the round-trip time is roughly 700 milliseconds and then they dropped it to about 230 milliseconds, then they had to back out a bit and only one of the directions was using the fiber links. And now they're pretty stable at 320 milliseconds. And not only that – the losses were reduced. And the throughputs went up because, as I mentioned earlier, the throughputs are inversely proportionate to the round-trip times. Next slide. I'll skip through the next one fairly quickly.

These show you other places. The top one shows you Zambia. You can see on August 20 it dropped. They had some problems and they actually had a lot more losses for a while. The next one down shows you Tanzania; big drop. And then the final one shows you SLAC to Uganda and bear in mind Uganda is not on the coast. You have to come through Kenya first in order to get to Uganda. And not only have they got Uganda, they've got to Rwanda which needs to go through Kenya, Uganda and then get to Rwanda. Next slide.

So I'm going to go on to a little bit about some case studies. There were two big fiber cuts in the Mediterranean – one in December, 2008 and on in January, 2008. The graph at the bottom right is kind of a contour plot where the countries are labeled to the right and along the bottom is the date and the colors are supposed to show you what the throughput was like. Dark red colors show you good throughput; blue is deep ocean so to speak - it is very poor throughput.

And you can see that when the fiber cut happened, everybody sudden drops into the ocean so to speak and you get this dark blue color. And it not only affected Egypt – the cut, I believe, is off Egypt – it also affected many other countries as far apart as Pakistan, Bangladesh and India, as well, of course, as the Middle East.

So if you look at the graph on the left, you can see that actually the current activity did not disappear entirely but it actually went up from

200 to 400 milliseconds as there was much more contention because they were using backup links. If we go to the next slide.

So then we had the recent internet shut-downs. This is the top graph here shows you SLAC to the National Authority for Remote Sensing and Space Sciences which is in Cairo. And you can see that the last connectivity we can measure was on 11:30 p.m., January 27 and at midnight, 30 minutes later, it had gone away. And then it returned to service at 1:00, February 7, 2011.

The one below shows you Libya. You can see that Libya dropped connectivity on February 19 and then came back again. And then at 20:00 on March 9, it's gone away since then. The next slide.

So then we also measured Japan which, of course, is very topical at the moment, and we monitored six hosts in Japan. There's one at RIKEN which is actually a research establishment just outside Tokyo. There's one at KK which is a physics lab about 60 miles N.E. of Tokyo at – I've forgotten where it is. Anyway, it's just northeast of Tokyo. Then there's another one at Osaka and another one at Okinawa. Unfortunately, we didn't have on at Tohoku or at Sendai. You can see the location of the earthquake itself by the red.

So what we saw was the average… First of all, none of the hosts went down; they all kept running. They were all accessible. But some of the hosts, as you can see, took a large jump in round-trip time and then stayed at the new value. So we started to look at that data in more detail. If we go to the next slide.

So in this case we looked at a monitoring host within Japan. The first measurements were from SLAC, so they were from the west coast of California. So then we looked from within Japan, from the research establishment at RIKEN and it turned out all the Japanese hosts had a constant RTT; none of them had this big step. So then we looked at the RIKEN host which seen from California had taken a big step up in round-trip time from other monitoring nodes around the world. And there was no effect seen from Africa, East Asia, Europe, Latin America or the Middle East.

But there was a big effect as seen from some sites in North America and from Canada – the U.S. and Canada. India – it was mixed. There was a node known as [C-DAC] which is one of the internet providers in Mumbai, saw no effect. [Puno], another node did see an effect going from 380 to 460 milliseconds in VSNL in Mumbai so a small increase. Sri Lanka saw no increase.

And then where we have about 20 monitoring hosts in Pakistan, we really found out it depended upon the internet service provider. If you were using the Pakistan educational research network, then you would see the effect. But if you were using somebody else, you didn't see the effect.

So then we went back to looking in the U.S. and what we found out was if the route went westbound from the U.S. or from SLAC to Sunnyvale and then crossed over to Japan, then there was no effect. But if it went eastbound from SLAC to ESNet to the Avenue of Americas in New York, then there were big increases.

So since reading all the blogs that have been going on, it appears that there'd been some fiber disruptions in Japan and I believe what must have happened is that the congestion as people move off the broken fibers has affected some of the links, but by no means affected all of the links. Next slide.

So that's it. The pingER measurement engine, by the way, was IPv6 back in 2003. The analysis, however, and the presentation stuff is not IP v6 capable yet. We have to get that fixed up. There are some references here which you can go and look to in more detail. There's a pingER home site; there's an Annual Report we put out each year. And if you want to hear more about the case studies, there's a link down at the bottom.

So that's it. Any questions or any thoughts?

Eberhard Lisse: Cool. You can just do this with ping? Cool. I'm living in Africa and we're going to talk offline because you'll get a monitoring host with us too.

Stephan Bortzmeyer: Stephan Bortzmeyer, AFNIC. How do you, in the cases of hosts that stop accepting STMP request or firewalls that suddenly stop to… because when you have long-term measurements over several years, those firewalls are modified and ICMP calls are very often blocked.

Les Cottrell: Yes, that's a major problem. We have to keep on track of that. We keep records each day of hosts which are no longer responding and have not been responding for 60 days. If they're not responding for 60 days, then we typically try and choose web servers and the reason we choose web servers is we know if port A is open. So then we can actually see if port A is still open and if ping is failing, we know they've blocked ICMP.

In that case, assuming that we don't have enough hosts in that region and that country to characterize the reason properly, we try and find another host in that region in that country which is responding to pings. So the hosts do change over time and as long as we have enough hosts it doesn't become a problem. But, of course, if there's only one host in the country and we have to change that host to another host that is in a very definite region, for example, then we could get some strange effects which have nothing to do with internet performance. So we do have to be careful with those things. It's a good point.

Lutz Donnerhacke: Regarding the registries, the (inaudible) registries added a new field to the databases. It's called pingable, so we have it on the system and it has a system object and we have a field and it's called pingable; it's guaranteed to be not blocked and to respond to pings, in order to find out connectivity to the (inaudible) system is working. Do you feel that such testing facilities should be more instructed or more pressured by policies like ICANN to say we do need a correct way to check that everything is okay? Should it be included in policies?

Les Cottrell: I think I got your question. I think you're saying that there's some nodes which guarantee that they will have pingER up and whether or not they will respond to pings? Correct me if I missed it. And whether or not this should become a policy that somehow should continue to respond to pings. Is that what you said?

I'm 100% behind that. It would make our life a lot easier. But the reality is I don't know whether that's ever going to happen. It's certainly extremely important being able to find things. I think most people when trying to troubleshoot a problem, they will probably use ping as a first line of approach. And when it fails it can be very misleading, especially if you tell a novice user, "Why don't you use ping and see if it works," and then they tell you, "Oh, it doesn't work," and it doesn't work – nothing to do with the network, but to do with blocking pings. It is very disturbing but we've accepted that.

By the way, if anybody wants to become a pingER node, we'd be delighted to have some more monitoring nodes.

Eberhard Lisse: Alright, thank you very much. I think an applause is deserved. Okay, Japp will tell us an update about NSD and then we're about done.

Japp Akkerhuis: Hi. I'm Japp Akkerhuis and I will try to be quick, although I have 19 slides.

Eberhard Lisse: That means Jay will be quick. The point is we don't have to vacate the premises on time because the meeting afterwards has been canceled. It's just I still want to go to the other side to listen to the other presentations. But we don't have a time constraint that we have to be out here by a certain deadline.

Japp Akkerhuis: Like everybody else, we're also thinking about doing new version of NSD, the authoritative name server. And what we want to do here is first give an overview of what the general ideas are behind the NSD then list main features of the NSD version and then what we are going to do for NSD No. 4 and I'll give a vaporware example.

First about NSD versions. The version numbering comes from 1.2.3. The version number is completely built off the… the version number is completely… normally internal completely changed in energy. The second is a major upgrade, added features, stuff like that. And minor number is for fixes and very minor changes so that's basically the theory behind the version number.

What are the principle characteristics of NSD is that we are actually authoritative only and basically geared to its root servers and TLDs. It all started because the root servers were slowly migrating to the same version of Bind and we have to create something from scratch to prevent cross, to prevent that a single packet of data will bring down all the root servers at the same time.

One of the other characteristics is that there's just enough in documentation where we really expect the user to be technical competent so we don't have to explain every little detail about how DNS works. And the other thing is because it's keep it very simple and we don't want any extra features creeping in and bloating the code. We are firm believers in doing one job and doing that single job well. And we also didn't bother with other classes unless we had to. We used just the internet class because the rest is not used anyway.

One of the main things we really want is resilience against high load so we should really perform well under high load without being crashing or deteriorating performance if we do it slowly.

The other one – it's completely built from scratch; completely independent code so we introduced our own [Bird] and don't take the wrongs from other people. And that's trying to keep in the hope we don't make the same (inaudible) as other people.

So the resilience against high loads we kind of did by having the answers we compiled, having it recorded. Also we assumed the data is rather static so the [sole] is not changing every five minutes. Remember this is for the root servers and it might not be true in the future. But the TLDs is also more for more static data than anything else. They might add a lot but in (inaudible) it is not a lot. One of the static points was that we will sacrifice memory for speed so we are memory hungry.

NSD 1.0 – that was the first one, just a server. The answers were pre-compiled in a database. The servers doesn't know anything about DNS. It's just completely ignore it. But pre-compiled software was done and user interface was very Spartan; there was no configuration. Period. It works or it doesn't. And that was basically the idea. And the sole configuration, if you really wanted to do something with XFR and variable stuff like that. There's little to no XFR tech support. If you really wanted to XFR (inaudilbe0, if you really wanted to learn it as a slave go to file from somewhere else. Find another way to do it. And we basically support the basic IFCs and the other two mentioned here.

What you see is that we actually expect a lot of clues from the users and that's why configuration – if people want to go config the zone, you had to know what you were doing. Well, things changed when NSC 2 came out. That is first version which was started out to be DNSSEC ready and it had to change internal database and also we couldn't compile that much interface because of the sorting the generation of NSEC, the server had to know something about DNS. There was less ignored about DNS protocol itself.

And we actually now delivered our own AXFR module for getting all the files from the other sites. I mean if you really want people to do it with one technology set, go to Bind. Use that.

[background conversation]

Japp Akkerhuis: Okay, I thought it was being dropped. And for all this we started to do this just local verification file and so make it easier for the user. And the next step still in the 2.0 is slightly more about dynamic behavior. You go to transfer into the server or out to the server we supported TSIG and we also delivered some small counterprogram for a less Spartan user interface and it's actually somewhat complexity in internally and for user and might user can get away with slightly less clue and because some things were done for them. But still big memory hog. It will always be.

NSD 3.0 – This is actually the current one that we're running now and there's a lot of extra stuff to make life easier for the user and so we now list Notify; we used actually timers in SOA for doing the slaving. So we also do full DNSSEC in this one and NSEC3 currently into it and there's more DNS meta support, DNAME is in it and the counter program has been expanded and that's right, internal complexity – you get special for getting that XFR stuff and so now it's actually three modules doing IPC. We're still on the same server speed and no internal database change.

What really did change is that it's even easier for users to use it, so less clue required there, apart from doing the complicated files because that too of course.

There's some ideas we have about NSD4. At first, we have a logo so we can make t-shirts out of it. And the other thing is lots of zones and that's one of the goals. We should do at least a couple of hundred K zones at the same time and that's the question we get a lot. So special high speed want to use it as well, not just ccTLDs. So we think about doing sole configuration templates because specialized ISP more or less to say only the name changes and whatever.

There are a couple easy ways to get around so you can automatically generate and there are much easier sole configurations. What we are also going to do is again a change of the internal database and that is two aspects to that. We are actually able to speed up the server with changing data model from what we're doing now.

And we're also thinking about more people assessing so we can even getting faster. And one of the things to do is to store the NSEC3 hashes and there are a lot of other things we're thinking about.

But we, of course, will happen is internal complexity will grow enormously but complexity moves to the compiler subsystem and so not to serve itself. We try to keep that as ignorant as possible because ignorance is bliss.

And as we come to our well (inaudible) status we are showing signals to the various parts, actually have a real controller port so it can much better control what is going on. And as we all actually hide the complexity way more than... So the users can be again clueless. We will also try to allow for more dynamic behavior. Reconfiguration on the fly and reloading some and maybe some dynamic updates, but we only want to do that if we can get away with it. It really makes things more complicated than you want.

What we also want to do is improve the TCP support with all those bigger and bigger packages you much easier to go this be than anything else and then some ideas to allow a lot of TCP streams. But what we actually don't want is all these other features are hampering the target audience, the root servers and the ccTLDs because that is still our point where we're working at. So we've got to lose some features or say no to people.

We will actually say no to some features and requests because it will slow down things. And whenever this will become non-vaporware, well, at the end of this year we hope to have the first version done. But we're still open for suggestions. We're still kind of playing with it. And I'm around here all week so if people want to have some specific stuff, talk to me and I'll compile the list and then we're going to say no.

We have done some testing and so for real marching slides, here we go. The new memory layout had to put in NSD 3.2.7 and we tested that with three scenarios, sort of like a growth in volume and one zone, 500 delegations. That's it boys and girls. Medium sized TLD. One zone, 1M delegations. And high speed type of workload – 100K zone is 10 delegations/zones. Standard stuff.

And test setup – well, we had the numbers of the test setup but basically we did about a million queries, randomized. So it's roughly 100 qsp per second will come to 64 mbit queries stream. So if you see that happening you know how many queries too. The assumptions are all domains called more or less the same; we don't look at NXDOMAIN.

We also don't do DNSSEC with those tests. But then it doesn't really matter that that's for the server; it doesn't really know about it.

Here's what we can do for the root server. The red one is our standard echo test – basically shows you you go up to the STAC and the only thing we do is echo the package back and you don't do any processing. So in case your ID the maximum speed you can get stuck in and out of your machine which is now 10,000, something like that.

Then we have here the first version of NSD4 so it's new memory and this is NSD3 and this is a version of Bind. And you see the difference counts for themselves. So you see we actually win some space there. Roughly about 30%.

The same with if you've lots of zones, or a big song – sorry, we're a TLD now - and you get about the same figures. So you can easily do that with about 50 queries per second – you get that. And then we have being an ISP lots of zones, and we are holding it pretty good at. And these are these real bumps which we don't – we have an explanation for it but it's so complicated you don't want to know. But that's the deal with how… We have time but you still don't want to know. And I forget how this really… complicated. But where you see is that this looks promising still for holding up. So we actually surpassed our original calls and we can do all these things.

More marketing slides. This is a message we… So it's way more dramatic – 95% returns. This is brutal by NSD and the echo demon. So you see in efforts we actually bump up 30% using these points. As I said I'll be around to talk and note the little line on the bottom. I'll take checks. Any questions?

Eberhard Lisse: Also quite a good presentation I must say.

Mauricio Vergara: Hi, Mauricio from .cl. Do you have any plans to introduce more logging in your software?

Jaap Akkerhuis: Well, the first one didn't have any logging at all and the reason for that is you've got TCP don't do it that way. Although it doesn't tell you a lot about what's going on inside the server. And there are some plans to do some improving logging but logging slows down the server. So you're very reluctant to introduce logging for standard use. We do have an awful lot of debugging and other flex but you don't want to run them continuously. If you have a specific need for doing logging, let us know what exactly it is and we might be able to introduce something there.

But we actually do some logging now but it's very little and we have to do a lot of hand waving to get it out but we might do something. Especially now if the (inaudible) channel, it is probably easy to switch it on or off in doing a real run so you can have quick log, see if things are okay; okay, switch it off again because you want to get maximum speed.

Matt Pounsett: Matt Pounsett from Afilias. One of the things we find with 3 is that it's also very IO hungry. We've got a couple of our zones that are updating 6,000 records per minute, things like that. And we have to do the IXFR right back to disc; recompile the database then reload the entire zone off disc into memory. And if you have a few of those running on a machine because you have multiple zones, then it really bogs down the IO on the machine. Are there any plans to change how that architecture work?

Jaap Akkerhuis: Yes. There are plans for changing that. And I was talking about a compilation. The things will be compiled; there will be multiple processes running and doing stuff on the fly and then when they're off of the port of the new database and so it can be done in parallel and… If it's really becoming a problem now, one suggestion is we can do the compilation on different machine, but that is more work, of course.

Mahmet: Mahmet (inaudible), NSD user as well. The logging question – also do you have any plans for some kind of statistical tool, deployment that can actually show us what's like the query load on the server, etc.?

Jaap Akkerhuis: We haven't thought that much about logging yet but whether… we are very careful to try to keep the server as stupid as possible because it helps. As I said if there's specific things you really want to be logged and if we can squeeze it in, we probably will. So make wishes known early.

Mahmet: Were these tests done on IPv4 only?

Jaap Akkerhuis: I actually don't know.

Mahmet: Is there any reason why no DNSSEC-enabled tests? You mentioned that those tests weren't done with any kind of DNSSEC queries, if I'm not mistaken.

Jaap Akkerhuis: They're done without DNSSEC, no site zone or whatever and it was just a very quick run to see whether it was in need of improvement.

Mahmet: I personally think that it would be a good idea to include in official tests with you guys as well as with Bind 10 and ISC with zones that are signed because the actual world is right now signed so…

Jaap Akkerhuis: Yeah, yeah. This is not meant to be serious test or whatever. This morning I actually heard what we did as another experiment to do more compilation and doing that on the fly to have more compiled answers – gives you another 30% but it really eats memory. It's probably not something we really want because maybe the memory chip, the files will be very happy.

Paul Verhoef: Paul Verhoef, (inaudible). A lot of setups actually using primaries for providing DNSSEC to another server to then sign it later on and then push it forward. So one of the things that's really missing in NSD for that deployment where you have a hidden name server going towards the sign or going towards public name servers is the somewhat more automatic NSD patching and the IXFR from differences. So that's really a big miss. And the other thing is please use GIT instead of SEN.

Jaap Akkerhuis: I didn't get all of that but the way AXFR are done, the (inaudible) will be gone, it won't be there anymore. That's why it's called 4; it won't look like NSD3. It's almost the annoying way NSD3 works won't be there because it's completely different. And that's especially because people complain about why are you doing… I think you were one of the complainers. What probably will happen it seems not all of these things are needed for the root servers that we probably will have v3 have a longer life than what we normally do which is basically one year after a major version comes out, the rest is killed, but since this is so completely different and not all of our users are actually in need of this, we might have longer support for v. 3.

Dmitry Kohmanyuk: Hello, Dmitry Kohmanyuk here, Hostmaster, Ukrainian registry; also user of NSD, quite happy. I have a very simple request so that we'll start donating to you guys. It used to be that the times, temps and log were in Unix UDC format. I realize everybody wears the clocks that show that. It would be very nice to have an option to log in UDC, ISO time format, like a year, month, stuff like that, not the Unix UDC time format. So that's the only thing I would say and thank you for your presentation.

Jaap Akkerhuis: I didn't catch everything of that but you mean time stamps which log in?

Dmitry Kohmanyuk: And on second since first of January.

Jaap Akkerhuis: One of the reasons of doing that again, if you don't have to conform to human available stuff then saves cycles. But it's already now an option to do it in real time, in human editable things. But if you really want to squeeze everything out of the machine you want it to be stupid as possible.

Dmitry Kohmanyuk: I don't understand that because that's not time critical so you write the intent to translate it. I don't understand that. If you want to squeeze all (inaudible) out of a name server, you… You trim all the fat as they say in this country.

Jaap Akkerhuis:: I don't know what the source code trim all the fat (inaudible), and I don't know what the original… what the current line count is for NSD. Anybody could find out. But I remember NSD 1.0 when it came out. It's only 8000 lines of code, that was it, compared to the 300,000 of somebody else.

Eberhard Lisse: Alright, no more questions? Thank you very much. We most certainly must do this again. Okay, as usual we have closing argument so to say and as usual Jay Daley will do it and because for this time it's a joint thing with OARC, Wayne and Jay will do it together.

Jay Daley: Sorry, sorry, but I wasn't really (inaudible).

Wolfgang Nagele: I'd just like to do a call. My name is Wolfgang Nagele. I work for the RIPE NCC and AFNIC had a presentation yesterday here about their

DNSSEC failure and we also have to admit that we had a failure in E164 ARPA which is the (inaudible) domain. And in light of that we decided that we want to put forward an approach in terms of how we can verify DNSSEC before we are publishing a bogus data. And so what I'm basically calling for here is interested DNS operators to join me after we close here to have a brief discussion about what requirements those DNS operators would have for such a verification mechanism. So anybody willing to chime in there, just stay here after we have closed.

Wayne MacLaurin: Well, first of all I want to thank everybody who presented and everybody who showed up. It's been a really interested couple days. A couple things we observed – first of all this is by far the best attended OARC meeting we've had in a very long time, so thank you. We're up about probably almost 50% from where we usually are, so that's great. Interesting bunch of topics so that was great.

In particular, I want to thank the guys from .CZNIC for showing up with the GPU presentation yesterday and to echo the comment of why – why not? I think that's great. I think as a community we need to do more of that kind of stuff. If you look at things like the problem that the guys in Chile are having with the weird patterns in DNS they're seeing, I think there's an opportunity for somebody, hopefully us, to take a look at that data and start playing mythbusters.

And let's take a few ideas of what it might be and let's see if we can prove or disprove it. We have a wealth of information through not just the individual registries, but as a whole. I think that's an interesting opportunity and something we're going to pursue. It doesn't necessarily have a particular goal, but it could produce some really interesting research and more importantly, maybe even find the answers to some of these questions.

So that's it for me. If anybody has any questions about OARC who don't know us, I'm around and I'd be happy to answer any questions, either here or offline. I hope to see more of you in the future. And with that, I'll let Jay close it off. Thank you very much.

Jay Daley: Okay, thank you. For those of you who want to laugh, very shortly we'll be starting the DNSSEC for beginners session in another room and I will be playing Dr. Evil again. So as Wayne said, this has been a great meeting. We've got 100 people in the room or something like that and I think we've had a very productive session.

For some years now the focus really through OARC has been on measurement and analysis and through the ccTLD community has been on greater collaboration and I think that this meeting has enhanced both of those quite nicely. From the ccTLD perspective, certainly the measurement analysis helps with our capacity planning, server placement strategies, service planning, those sorts of things, and I think huge progress has been made.

At the same time, though, there are two meta benefits, or larger benefits which I think are important for us to note, one of which is this cross-industry collaboration so that we're not just being ccTLDs or just being security companies, just DNS experts or whatever. We're working together towards some greater common cause.

And the other one is just bringing a degree of scientific method to DNS and then on to the technical community. For me, the answer to the question, "Why try to implement signing on GPUs," was to test the hypothesis. We're never going to know whether it works or doesn't work until somebody writes some code to test it. We can eliminate it as a hypothesis or we can learn something from it.

We have both at the same time a remarkable wealth of experience in the room and knowledge, but that can also be our undoing at times because there are people here who can work out the implication of something technical in about quarter of a second. And yet, sometimes unless we actually build something, run it and test it, we'll never know whether those things work.

And so what I see coming from these meetings and what I think is going to continue coming from these meetings and this community is that scientific approach where we develop our hypothesis, we build things, we test them and we act rationally and we get away from our judgmental side of things at just deciding one technology is better than another. And I think by doing this in a more structured way, we're learning things much better, we're developing and we're going to make much better, verifiable progress.

So thank you for coming; it's been fantastic. And I think we will be talking again about whether we can do this again because it's been such a good meeting for all of us. There are now two sessions to follow off to if you want to. There's the Law Enforcement session which has been going off all day where Eberhard will be leading off to that. And there is the DNSSEC for Beginners which I'm sure most of you know, but if you want to come along for a laugh and throw things at Dr. Evil, I shall be heading off too. Thank you all for coming.

Eberhard Lisse: Okay, the Law Enforcement session is in the Elizabethan and it's A to Z which is just around the corner. Alright, thank you very much, everybody, for coming.

[End of Transcript]