SINGAPORE – Tech Day
Monday, March 24th 2014 – 10:30 to 17:00
ICANN – Singapore, Singapore

UNIDENTIFIED: …are you seeing? Okay.

This is Singapore 49, room Olivia, date March 24, 2014, time 10:30 Start ccNSO tech day.

[OPEN MIC]

EBERHARD LISSE: All right then. Good morning. Can everybody in the back sit down so I don't have to call you by name? Okay. Welcome to the 23rd Tech Day. We are not calling it ccNSO Tech Day anymore. We are, as you all know, are trying to widen the audience, so current events apparently have overtaken us, and we scheduled a second meeting about the NTIA invitation for proposals, but to our meeting, but let's see whether this improves in the afternoon.

This time we have got quite a packed agenda. I actually had to ask Christina to get us some more time because we were so fully booked that I didn't want to cut the lunch time off. So we had to start half an hour early, and as usual, I must shout at everybody who is here for the other one, so I'm not here being late on things like this. In the beginning, we do three things about measurements.

The ATLAS project will be demonstrated by Jaap. Edward Cocker from the University of Auckland, I think, in New Zealand, has got

measurement… He had a measurement project as part of his PhD project, so he's going to present that. And then Ondrej to have a similar thing going on and so it might be a cool thing that these three topics are next to each other, and maybe we can get some discussion out of this.

Then Duane Wessels is going to talk about the detdns API, which is something that has been coming up recently which I find quite [?] I don't fully understand it, being really not a programmer myself. And then we have dynamic zone configuration from dot EU. Then [?], the Malayan presentation to us to have some mobile DNS validation.

I have no idea what this is about, but it's probably going to be cool. And then Hashul Hasan, who we all know from DNSSEC [?] enter into the [?] times. He's now working at Microsoft, he's going to talk a bit about their collaboration initiative.

In the afternoon, we have a round table organized by Leo Vegoda about IPv6 in particular, as far as mobiles in Asia is concerned. So out of APNIC, the speakers are from APNIC and from China and so on. I haven't really been involved in the organization of that, but it's going to be interesting for those people who are interested in IPv6.

Then the South Africans are talking about their Mark Validation System. And then Steven Deerhake is going to talk a little bit about what happened to dot AS. And then I'm going to talk a little bit about what happened to dot NA when I took the software that he developed and played with it a little bit. And then Paul Ebersman is going to give us a systematic overview, which is a little bit longer, about DNS attacks and what we can do about it.

And the host presentation will talk a little bit about the abuse management system that dot SG has developed. So first presenter, Jay Daley, is going to give us some wise remarks. So first presenter is going to be Jaap, he's going to speak about ATLAS. Let me just take one. That's an ATLAS probe, it's basically a little WIFI router that I can buy for about 40 Euro off the shelf in [?]. I don't know really what this is, yeah.

He will explain it. The point is, I've got one of these last time. I put it on my system at the house, and it now talks to the system and we get very good measurements. I still haven't figured out how to look at my own measurements, but that's probably not the point.

JAAP AKKERHUIS: Okay. Good morning. I actually had a nightmare about this talk, literally, this morning. And nobody was here, and people who were here were not listening. Let's hope this is going to be better. I am actually not talking about my own stuff, I'm talking here to people from the ICANN cc. And they're actually doing all of these things, but nobody was available, so I got volunteered to explain this.

And to be honest, I don't really know a lot about it, other than what… The slides are from a workshop, and I compressed them a bit so we can fit them into 20 minutes slot we have here. There are a lot of reference to other stuff, so hear that. Just quick word about RIPE for the people who don't what RIPE is.

RIPE is actually the first RII, and doing the Internet number registries, and the [?] later on the APNIC. And the second one is better known in this region. So basically it's not a legal entity, it's just a community. And

they actually set policies for the – for the Internet registries. [Meets twice a year?] at the moment, and work is done in working groups and…

The real entity, the entity which is doing the work is [?] secretaries now known as the Network Coordination Center. A non-profit organization located in Amsterdam, I think around 12,000 members. A lot of them actually in the eastern part of Europe. But they do a lot of… A part from handing out IP numbers, they do AS numbers and that type of stuff, to do a lot of services for their members as well.

And ATLAS is actually one of the newer ones. So ATLAS, they still in the deployment and as you can see, all those little green dots are ATLAS probes. Places where these things are… This is an ATLAS probe. And the little thingy, it's not a wireless card. I mean, I always thought it was, but it isn't. This little thingy on the side is activated, SSID, it contains complete system on it, and it seeks to keep something like that, as a [?] and so on.

That's what it is. This is your Ethernet cable. And this is the double USB thingy, and it actually only takes care of the power. This no power on the Internet of this thing. But that's the basic hardware. This is TP-Link, now to modify it by RIPE with a special custom thing on it. And so there are about 5,000 to around now, it's actually more than on the slides.

The slides are too old. And more than 10,000 registered users. And the probes are doing measurements, ping, traceroute, DNS, SSL, and of course done on IPv6 and that's basically the idea. The idea is that you have an enormous amount – that you want to have one global measurement machine which you can use for various measurements, and for your own measurements as well.

The idea to join this machine and network is everybody can do that. It doesn't really matter. That's why you see in America, and Australia, and all parts of the world. All of the [?] is in Europe on this map. So anybody can become a [?] and it is… There are two ways to do that for you. Personal IDs that you want to measure stuff and have a look for your own network, from the outside part of the world. I actually come to see how [?] is connecting to Washington, or similar stuff.

But as an idea of what the other peeps are seeing from your network. So there are now about 5,000 points where you can show, see how the global network is doing. And you can also put in your own measurements stuff. And the measurements are available to everyone, and there are maps of data for public services. There is an API to download the raw data, collected so you can do your data analyze on ATLAS of the time.

Don't have to do things in real time and well. These are the little probe. This is the now the one we hand out, and I've got 10 with me. So I'll see requests coming, I guess. It used to be, the little one on the left is the first version of it. It was actually a special made hardware, and modified hardware, and the only difference between that and this is that this one is just standard hardware. They're actually cheap in Europe than in [?]. I think they're 20 bucks.

But no need to chase the hardware, just having a lot of disk in it is what really make the thing tick. And the thing on the right is what we call ATLAS Anchor. This is a Soekris box which has more capacity, and the idea is also these boxes are connected to each other in kind of a mesh

and [?] and about 1,000, no there are less. I don't know how many there are.

This actually costs money. I mean, if you really want the [?] box, you get one and you really need to do some work for [?]. It takes more than [?] and… I mean, these things hardly takes anything. People have to [net?] their homes as well, and it's really [?] and stuff. And anchors, their well-known targets and they're also targets for these guys.

So if you find… So if you want to probes to well-known places on the world anchor sites, one of them are stable, they're always at the same place. And if that says 41 anchors installed, but I think it's about 60 at this time. And they actually do measurement in between them, but about 200 probes are targeting each anchor and actually probe measures four or five of these things. So it's better getting more stable view of the world. Again, anybody can get these anchors. There is a way to apply for them. I won't go into details on how to do that, but you can find it on the net. It's easy.

I said there is always quite a lot of data available. I mean, one of the things called seismograph. That gives you, from the various anchors, various probes you see here on the right, telling you how long it takes to get to the anchors. There are a lot of details on this map and you can find all of these things online as well. It helps you to analyze the data, so this one, data analyzing software available.

And furthermore, you can actually do data where you can zoom on the details of the various measurements being drawn. This is an example of how to do that, the data for a specific probe, and for a specific measurement to a specific target. And so, you can also use this for

doing troubleshooting. Why can't I get to Turkey? Or can Turkey get to me? Well, there is some politics there, but a part from that part…

There are quite of things that have already proven to be useful for using the system. When [?] put in yet another cloning to Belgrade, you can see the local times releasing to L-root drop dramatically. And it's also used to figure out where it's a good idea to place your nature…

The second one has been… AfriNIC used it to figure out what was the most popular [?] cast place to put the nature of it, and get [messaged to death?]. And of course you can measure outings everywhere. The last one in Sudan, went on and offline and seem from various perspectives. And this is just all my examples. Furthermore, the tools for doing officialization [sic] of the data and an example with the reachability of the IP version six here. And another interesting idea is to find to how you can better to do P2P routing based on the prefix size. And this is a Hong Kong study trying to collect routing table somewhat.

And the last one is one we did some time ago, and it's time to find out how fragmentation is used with publishing IPv6, especially with respect to DNS. And we found some interesting stuff. This is a follow up study going on, and I'm not involved in all – don't ask me about the details. But there are a lot of places where you can find these things. Another project that we just started, two days ago it came officially out as a batter test, is doing the [?] DNSMON boxes.

Now also done by the ATLAS system and trying to improve the measurements of who the DNS world looks like from the various vantage points. One reason to move this to ATLAS is that DNSMON is always hosted on the old traffic time measurement boxes people had.

And these are getting really old. That system is going to breakdown, is going to be phased out. So that is some DNSMON [?] on the ATLAS system, it's going to placed. And the idea is that there will actually be more features and better features for doing monitoring of DNS.

And there are a lot of the usual interface stuff. You can have your favorite measurements. This is improving, as I said, the trace route visualization happening, and various registries, regional Internet registries in trying to get the probes all over the world. My personal calls get what you want on [?]. And I probably have a need to get one there, this will be very interesting.

Anyway. There is also a mailing list for ATLAS probes, and where you can ask for specific things, report problems, and so on. And here is a small bit about how to use the measurements. There are customized measurements using the targets, and you can set up your measurements only for your own goal. And there is API for creating the measurements, so you can do it automatically. And there is an API for analyzing measurements.

I didn't want to go into all the details because it will take too much time, and I don't really understand some of them, I only understand some of them, that's probably a better way to say it. So, how does it work? Well, there is a kind of credit system. If you have one of these things, you get credits. So you get about 20,000 credits a day.

And if you perform your own measurement, you spent those credits, and so the entire sum cost, as I say, it's what you do with your own measurements. Ping costs 10 credits, trace route 20, and so on. And there is a limit about how many credits you can daily use. The idea

about credit, it says to introduce some fairness to all the people running those probes, wanting to do those measurements.

And also trying to keep the system from overloading at certain places. So, more details. There is a whole page about how to credits, how they are spent. You can even, if you want to do massive measurements for a certain period, you can actually transfer credits from yourself to a friend or the other way around, so you can take credits if you want to, or need to.

Nothing, the system which is actually just starting, but we'll – people are actually using it more and more, there is a built in plugin to check for various things, so you can send yourself alerts when things are going haywire. So you can use it as a cheap way of looking at how your system looks like to the rest of the world, and actually get alerts. The first one was, I think, this is also [?] plugin and more of these things. So, details there.

And, well there is a big community. People are actually sharing the various measurements they do, the programs they use, for doing the measurements to analyze. There is a big GitHub, a set up community where people can [?]… and doing whatever you please. The whole source code of the system, especially of the [?], is available online as well, so you can see what is actually being done.

And just as me, you can be volunteered for being ambassador, where you help to hand them out these things. The culture is [?] to you and you act at conferences like this, and try to [?] them off. And of course, there is also sponsors, quite a lot of sponsors, helping RIPE to finance this stuff, and I would say, the round of the usual suspects, this is the

sponsors of last year. Some are continuing and always welcome to have more sponsors.

Then the community itself also decides what the globe map of the system is. So who sophisticated should we hold the system check? Not everybody has the same demands of trying to, the figuring out how things work. On this moment, it's more being used for research, but it's to also an operational need, especially for the RIPE community, for the operating networks in that RIPE territory, for doing specific measurements.

Furthermore, some of the measurements, actually most of the measurements on this moment, are perfectly available. So you can look into what other people are seeing and [?] and publishing and analyzing. But maybe for competitive reasons, people won't have stuff being done, at least in private, so not everybody knows the company's secrets, or the measurements, or whatever.

And so this is the debate on mailing list about this. And then there is a problem from, how far can you go to test stuff? I mean, one of the things is it would be nice to use the system to test how much the various networks of the world are available for spoofing, and all of the spoofing attacks. But doing that, you probably have to break the protocols trying to figure out, and is it really ethical that I'm doing measurements by spoofing other people's networks? Stuff like that.

So there is some discussion about that as well, how far you can go about what type of measurements you make. And so, the slides will be available, but here are some more links people want to look at, and

articles, and things like that. There is a Twitter channel. If there, oops, any questions. I think I just made it in time. Yeah. Three minutes.

EBERHARD LISSE: We've got five minutes, three to five minutes for questions. We can go and with the microphone, or you can come to do microphone.

UNIDENTIFIED: Thank you. [?] from Costa Rica. Yeah, here. What about the security in these network probes? Is there a chance that somebody can use the data and generate the [?] network to create some kind of attack, or by example, [?] – is there any chance to hack this network?

JAAP AKKERHUIS: I never say no, of course. I mean, there is always a chance to hack, since it's a Linux system. And what the [?] we do is, they automatically update themselves, so if there are problems discovered in the system, they will be automatically patched from RIPE or the antivirus things, it's downloaded.

And trying to use this as a Bot Net, is actually limited by the fact that if you burn your credits, you cannot do anything, you need to start traffic, getting out of these probes. So there are some limitations trying to get off of that. And that's also the reason why people are encouraged to have their measurements online, and to call out to people can help out when a problem is detected, and things like that.

Now, of course, never guarantees, there is always the last [perk?]. It's very active in development. Just yesterday, so somebody was actually

having some problems with one his probes, and it was fixed within an hour. So this is a very active monitoring going on, on these things. And that's why, for instance, also these ideas of how to find [?] networks, has been delayed for a long time, specifically because then you could – you abuse the system for doing attacks, because that's actually an attack if you started to actively measure spoof packets or introduce them to find out.

Quite some work going on in trying to keep the system safe. And you can also unplug them, I mean nobody will… If anybody wants, there is [?] as well. I mean, as I said, I have 10 available, so just approach me and I'll… You have to fill in, I need some data and then I'll activate it and that's it.

STEVEN DEERHAKE: Jaap, I've got a question for you. Back in one of your earlier slides, you showed schematics…

EBERHARD LISSE: Can everybody identify oneself because we have a remote audience who can't see you.

STEVEN DEERHAKE: Sorry, Steven Deerhake, dot AS. Back in one of your earliest slides, you had a schematic showing the first generation probe, the current generation probe. I've been trying to determine what became of one of the first generation probes in American Soma for a number of years

now. And I was wondering if it actually did recover it? Do the first generation probes still work?

JAAP AKKERHUIS: They still work. There is this functionality that is actually to say, they still work and being worked. There is actually no point to replace them by the newer probe.

STEVEN DEERHAKE: Then I'll keep the search going. Thanks. And I would love a current generation one as well.

JAAP AKKERHUIS: Okay.

EBERHARD LISSE: Okay. Next speaker will be Etuate Cocker. Is it for Auckland? This is from the University of Auckland. As part of his PhD work, he's done some measurement stuff too. So he will now tell us all about it. He's a Windows person, he isn't used to a proper computer yet.

ETUATE COCKER: I think I'll [?] Linux, not [?]. Thank you. Perhaps for those of you that have attended some of my presentation, especially in APNIC, the stuff that I'm presenting to you will be a bit of an addition, in the sense that we extended the research to doing measurement with use of TCP protocol.

This research is not only myself, I would like to acknowledge others at the University that have contributed a lot to what I've done, especially my supervisor, Dr. Ulrich Speidel, Firas as well as Nevil, all of them from the University of Auckland. Right. An agenda of what I will discuss, I will first talk about how VoIP, the voice packets are being transmitted over the Internet. The journey of the packets, what it faces when it interacts with infrastructures over the Internet.

Some of the implications of the growth of the Internet and technological progress, and why do we need to worry about it. The beacon network which is what we developed. Some initial observations based on the data that we have already collected to date. Packet travel arrival quality, again this is referring to the delays and how we use delay as an input into determining [?] with use of entropy.

Then I will talk about TCP, which is something new, like I mentioned before, and I will conclude my talk. Right. [?] voice over IP packets. Some of you may be aware of this, this is a very basic diagram, illustrating what occurs when you use Skype over the Internet. When you speak to the speak to the mouthpiece, that's analog voice that the encoder, or the codex, connects it in [?]. Then created as packets, those packets travels over the Internet on multiple routers and links of different capacities, and slow processing speed.

So depending on when the packets actually arrive on the destination, it is then buffered through the obligations buffer, the Skype buffer. When this is buffered, sometimes the packet arrives out of order. Many of that obligation is responsible for reordering it and playing out through

the other end. So what this means is that when we transmit packets, we transmit it on time, sometimes in sub-milliseconds.

Sometimes packets are stuck in queues at routers. Sometimes they are also dropped, separated from other packets, taken for a ride depending on routing protocol that is employed for routing packets from source to destination. They are made to wait in the receiver's buffer until the rest of the crowd turns up, the rest of the packets turns up, before it has been played into the other voice over, the application.

So the status of the Internet right now, we do this in less than [?], we take this screenshot from submarine cables map dot com. 2014, you can see a lot of cables being planned for implementation. Some of them are already [in demand?]. What does this mean? It means that with more links that is available, it means that network engineers and operators will use the technique known as load balancing for routing packets over multiple path.

For us, we think, we believe, that this is a bad technique, in the sense that with more physical links, it makes shorter path, which is a good advantage but however, for voice over IP application, there is always the chance of packet being reordered along the way from source to destination. What we wanted to investigate is when packets are stored into multiple routers as it travels through the Internet, and this is what we call parallel queues.

The obligation of our research, just looking at it, some of it can be applied on remote sensory. I mean, we can measure the quality of links and with these industries, they can use it for the information. So what will be a world without Internet, without real time applications such as

VoIP? E-mail, web, downloads still work because they don't need, they have real time connectivity. More push to talk like communication, because there is no VoIP application.

More voice and video messages. And there will be serious digital divide between remote regions and regions closer to the core of the network. Introducing the, our project which is the International Internet Beacon Experiment. A beacon, we define it as Internet connected machines that are able to exchange synthesized traffic in a highly controlled manner without beacons.

Currently we have 30 beacons in all of these countries. Further beacons are under construction or will be deployed soon. In every experiment, we log the data into each host, each beacon, and then we transmit all of this information to the University of Auckland.

We don't do this, all of our beacons is responsible for reporting to us. This map might not be really as good as the right map, because they have more probes. For us, we rather have control of who does things for us, in the sense that we only give it to people that we know.

In an experiment, every beacon… So we require two beacons to exchange information, and they exchange 10,000 TCP and UDP packets. For UDP packets, we create packets that has a payload of 110 bytes. And each packet is transmitted at every 20 milliseconds. Packets are time stamped and serial numbered, so that we can know whether it was processed, or whether it was received out of order.

At the receiving end, packets are logged with an arrival time stamp, serial number, and all of these other metrics. We run three

experiments a day, and we use the information to derive the characteristics over the Internet such as packet loss, out of order arrivals, clock drift, etc.

Why do we have these many beacons? We want to have long term global trend, not just the local effect. Right? We want to see what the Internet is like after 10 to 15 years from now. We want to a developed world baseline, but also what we'd like to see what it is like in developing countries in the Pacific, or in other regions.

We want to see the long paths that are generally of interest, especially satellite connectivity. A lot of our international traffic passed through hub regions, so what are the effects does each region have on traffic that passes through them? Meaning that we want to see whether, in using a lot of this fiber optic cable, will eventually improve the Internet or it will cause a lot more data loss in the future.

Initial observations, some of the data that we have. This graph is about data that we collected between our Canadian, New Zealand, and the Cook Islands. Looking at it, it does not mean that you have high speed connectivity, you will not have high packet losses. Looking at it, we observe higher packet losses between Canadian and New Zealand beacon, rather than the connection between Canadian and the Cook Islands.

By projecting the data loss and out of order packets, we can see that at times, the packet losses corresponds to out of order arrival of packets, meaning that as more and more of these packets queue up in the links, it means that some of it will be delayed and it will be received out of order in the receiver. Some of the metrics that we use to project the

quality of voice over application, we use mean opinion score, we use jitter measurement, and we want to be able to tell the difference.

So how can we tell the difference? We can be able to compute the delays that it takes for packets to be transmitted from source to destination. However, it does not tell you about the patterns of whether the links is getting more delays, or it's getting better. So this is where entropy comes in. We employed t-entropy, which is a homemade entropy algorithm at the University of Auckland. How it works is that we assign inter-arrival times, that is the delays that we get the packets, to a specific stream like A, B, C.

We can see here that any inter-arrival time that is less than 17 milliseconds, we give it an A, B, and C. So if we get simple like, CCBDCACF, it means that entropy is high, it's really hard to detect any pattern in the arrival time of packets. If we see anything that is all C, it's all constant, it means that we can determine the pattern.

By comparing the jitter with entropy, we can identify two different types of jitter. One is random jitter and the other one is systematic jitter. Random jitter is when the entropy and jitter correspond to each other, meaning that they both rise at the same time, they both degrees at the center. This sort of jitter, you will get it if there is a lot of packets being queued out in routers as it travels from source to destination. Systematic jitter however is different in the sense that it refers to load balancing.

If there is a lot of jitter, but low entropy, high jitter and low entropy, it means that we can still be able to get a pattern, but however the jitter is high. It means that some of the packets were travelling through a faster

link, because of load balancing, some of the packets travels through a slower link.

But putting all of our data in a plot, currently this is what it looks like. If however, it moves into the direction of worry, meaning that we need to think about our routers and links we are introducing into the core of the Internet. This graph is about entropy between Tonga and New Zealand. The thing that I wanted to point out here is that, between Tonga and New Zealand, it has to go through Canada, throughout Canadian beacon, or through routers in the Canadian region.

The data shows that the delays are not between the New Zealand… It is not the link between Tonga and Canada, but rather it's, sorry. It's not between the link between Canada and New Zealand, but rather it's within the satellite. So as more of the publication delays contribute to data, it means that we have a lot of the entropies. We also project the TTL, and the reason why we wanted to get the TTL is to be able to see the changes in [?] packets as it travels through the Internet. Looking at the graph, there are times we see our packets having different route, but we don't see this every day, but it's only on a few occasions.

Now we will talk about the TCP experiments, because we don't have enough time, so I cut a lot of the slides. In this graph, the horizontal red lines, it means that that's time we receive TCP voice packets. The vertical lines, the red vertical lines, means that that's the time you can see on that amount of packets that you have received.

By drawing a line, the blue line, looking at the blue line here, that demonstrates the bytes needed for continuous rate of immediate replay of voice packets, so depending on when the packets have been

received, if we have something on this light blue region, it means that we have buffer under run. It means that we consume a lot of our data early, but the application will have to wait for the next time it receives packets.

By moving this blue line to the green line, we can then be able to project the minimum time that the voice over application will need to wait for it to be able to have us move play out of TCP packets that it has received from the source. Some of the data that we have projected use of TCP. Looking at it, at times it reaches one second, meaning that a voice over application will need to delay its play out of a voice conversation for one second, simply because of the amount of delays that the packet travels from the link from source to destination.

However, it drops but we would like to find out more of whether sending voice over TCP over the long term, will improve this on the quality of the infrastructure, or of the Internet infrastructures. However, some countries like Tonga, luckily they have the [induced?] fiber optic cable, and you can see a sharp drop in the minimum buffer time requirement, meaning that there is contribution of introducing new links into the developing world.

To conclude, I mentioned – I talked about the real time traffic and best effort protocols. I've mentioned a bit of experiments, how the data has been transmitted from our beacons to our repository. There are still a lot of work to be done. If you want to be a host of our research, just let me know after this.

Any questions? You can find out more from this link. Let me just…

EBERHARD LISSE: Don't worry about it. All presentations will become online, so the links will be available on the presentations. Thank you very much. A little bit more deeper than what we usually get here, but very interesting. My first question, and I would like to abuse the prerogative of the Chair, is would you benefit from doing something like this on the ATLAS network with 4,000 probes?

ETUATE COCKER: The difference with our experiment in the ATLAS, I'm not sure if there is a technical guy here in the ATLAS probe, we see the ATLAS as being beneficial in the sense that we can track the path from where the packets was originated and the destination. However, the problem is that there is multiple path, we need to analyze the pattern of where it goes through.

But that's a totally different thing. The other thing is that with the ATLAS, I think they use a lot of the short payload type of packets. For us, we use actual packets for mimicking the experience of voice over IP packets.

LOUIS: Yes. Louis [?] from Costa Rica. Do you introduce some of the wireless behavior on this measures? By example, several wireless, or maybe WiFi, because this infrastructure has a very different behavior, talking about these time of protocols.

ETUATE COCKER: We don't look at specific connection on a local area network. We only measure the effects of packets being transmitted over the Internet. So, we do not look at the local area network, [?] things. Yes?

STEVEN DEREK: Steven Derek, America Soma. As someone who has considerable experience using the Internet from a little speck of land in the middle of the South Pacific, I think I can really appreciate this work. I think it's really important work, and I thank you and your colleagues for it. With regards to your data from Tonga, I take it they're both using satellite link and also cable? Or just the cable now?

ETUATE COCKER: It's just the cable now. I think they use the satellite as the backup, I think.

STEVEN DEREK: Yeah. Similar to what we do. Where is the terminus of that cable? Not to get too far in the weeds.

ETUATE COCKER: Sorry?

STEVEN DEREK: Where is the terminus of the Tonga cable?

ETUATE COCKER: I'm not getting what you mean.

STEVEN DEREK: Where does it end up? Where does…

ETUATE COCKER: Oh, in Fiji.

STEVEN DEREK: In Fiji. And then its Southern Cross?

ETUATE COCKER: Yes, that's right.

STEVEN DEREK: Okay. Because I just run [?] to Hawaii. But I think we would see very similar results based on past experience with just satellite link to San Diego versus this. But I thank you again for your work, and if there is a way to get American Soma involved in this, I would love to.

ETUATE COCKER: Well, in fact, that's something that I wanted to do as well. In trying to get involved some more, so I'll have a chat with you on this.

STEVEN DEREK: Perfect.

EBERHARD LISSE: Any other questions?

UNIDENTIFIED: I understand there is one…

EBERHARD LISSE: Sorry, sorry. Introduce yourself for the remote audience please.

UNIDENTIFIED: Okay. My name is [?]. I'm from ISOC Asia. I understand there is one project called [?] by Stamford University. So, can you differentiate your project? I mean, what is the difference between the [?] project and your project?

ETUATE CROCKER: I'm not aware of the [?] project, however, there are a lot of utilities out there that uses the ICMP protocol for testing the path, the reachability, the stuff that we do like I mentioned before. We try to mimic what it's like, you having continuous voice IP call over long term. With this, we can then see if the Internet is contributing to a lot of the delays or improving things in the future.

EBERHARD LISSE: All right. Thank you very much. That was a very impressive presentation. [Applause] Our next speaker is Ondrej Fillip. He must first update his presentation a little bit.

ONDREJ FILLIP: Okay. Good morning everybody. My name is Ondrej Fillip. I'm from dot cz. I'm from the Czech registry, of the Czech Republic. I would like to talk about quite a new project, it's very recent, so not everything is completed. But something that is also related to the measurement. And this project is called Turris, which is a lot [?] virtual [tower?], and we mean a protective tower, something that protects people.

So what is it about? It's started lasted year, and it took a very long to prepare it. So the code name was the shared cyber defense project, which we now renamed to Turris, which is a more appropriate name. It was developed in CA.NIC Labs. At the beginning we setup four goals, whether we would like to, hello.

First of all, we wanted to do some security research. We wanted to improve the security of the end users, improve the situation, the CPU devices, those routers you have on the small WiFi devices that usually set up one, see if we get pass forward, if we get to update it and everything. And one side of it, we also wanted to do some network measurements.

So, just to explain those goals a little bit more in detail. Currently, we do a lot of security research. We have Honeynet, we check all DNS data so we try to find some anomalies in DNS data flows. So we get a lot of data from the central point of the network, and we want to do a little bit more. So we want to also discover how the situation looks like at the edges of the network, and close to the homes of the people.

So we decide to develop probes that we will send to the end users, something to [?], but we wanted to go a little bit further. We didn't want to just have a passive device that someone carried home, we

wanted to do a router that the data will flow through that. We wanted to distribute it to many networks, as many as possible inside a country, and again to reuse the system, we use for DNS anomaly detection, so to direct some anomalies in the network.

I don't think, related to the security, if we will distribute routers to the people, we would like to offer them some adaptive firewall, which will be based on data we collect, sort of backlist, and something that will react to some situation in the networks, an actual situation. And also to feed the security team with those data. We are running CSIRT.CZ, which is the national CERT team of the Czech Republic. So we are sort of authority which works in this area, for quite a long time.

Another thing, the situation with the SOHO routers is just horrible, you probably know that. It's usually better support for new technologies like IPv6 or DNSSEC, if there is any. Many of those devices have huge problems with DNS, which we figured out in DNNSSEC test project where we tested those devices and not many of them are able to transfer proper DNS. There is no support for third party applications, nothing in app store.

It's really a device which is just plugin, switch on, and you forget. Many of them, they don't have any security features, nothing, even the [?] support is rarely implemented. And probably the biggest pain of those devices is there is no automated software upgrades, no system that upgrades those devices. Again, you just forget you have it because it works, it's sort of the Internet for you, that you don't have any – you don't update it, nothing like that.

And last thing, we wanted to do some other measurements, like the spread of IPv6, DNSSEC. And also measure, for example, the reachability for DNS [?] clouds, which are important. The reachability for the Czech users to the Czech DNS of course. What we did, we wanted to distribute about 1,000 probes. SOHO routers, we wanted to give them for free to the people, or technically, because there is some accounting issues or financial issues, we would lease it for one Czech crown which is about 0.03 Euros for three years.

And we wanted to give them, you know, routers that are powerful enough to provide one gigabyte of traffic while analyzing it. So, it's not easy. And we couldn't find any such hardware on the market, so we develop our own hardware. Additional… We wanted to also give some additional features, you know, to make something that the people will like that will be happy to take home and use it. So something, some additional values.

So, from this developed the router Turris, that's the device, I will maybe first show you the picture. That's it. It's a quite fancy box. And CP size, it has very powerful processor. It has dual core of 1.2 gigahertz, power PC processor. It has two gigabytes of memory, which is probably, you know, more than all of those devices have. It has a lot of the flash memory, so it has 256 megabytes of NAND memory where our domain system is.

And there is also 16 megabytes of NOR flash which is for backup. If there is any crash, you can also restart the system from this backup. As I said, it's quite powerful. It has two, sorry five, gigabyte LAN ports, which are connected by two gigabyte channels to the CPU, and one

dedicated LAN port for the connection to the Internet, and that's directly connected to CPU. And as I said, it's able to [?] one gigabyte of traffic easily, and also with the analyze of this traffic.

One more stuff, it has two miniPCIe experts ports, which one is occupied by a WiFi card. WiFi is in Sweden [?], you know, really, one of the latest hardware on the market. It has two USB ports, which of course are really not related to the analyze, but they should do something for the end users. Any many other pin ports, many other connectors to some other stuff. There is also one micro SD slot for an additional of flash if necessary, but it's currently not use for our purposes.

And we have some simple crypto chip to store crypto keys and stuff like that. It has quite low power consumption, even if it's powerful, we consume about 10 to 12 watts, which is not bad for this kind of device. And everything is designed by open source license, so you can, if you're interested in the device, and you can download it, you can manufacture in your home, if you are able to do that.

So as usual, this project is very open. So this is it. Again, maybe we can discuss it later, but the hardware is probably over the scope of this group. But this was completely designed and also manufactured in the Czech Republic. So, that's a problem. It's, of course, quite costly. It's 1,000 PC series, but we really couldn't see any other chance. Nothing available on the market that doesn't have like fan or some mechanical parts.

We wanted to literally make a device that you were able to put into your home, and left it there. So no fan, nothing mechanical and stuff like that. It has one killer feature that became very popular. You can

dim brightness of your lights, it's really something that the people appreciate, especially those students that use the same router in the same room where they sleep, for example, and people like that.

I don't know if you've ever faced that problem, but it's horrible if the light is too bright. So that's first fancy feature that people appreciated. A little bit about software, which is probably even more important for you. It's based on OpenWRT, and again, everything is open source. OpenWRT, sometimes the GUI is a little bit complicated, so we developed our own configuration wizard.

And with that, we also implemented our own NETCONF device, so something we are really hoping we will be able to work a little bit, or maybe sometime in the future to make some [?] configuration of those devices. It has, as I've said, automatic updates, so we update those devices constantly. If there is a need for reboot, for example if the kernel is upgraded, the user can avoid certain time periods when the device could not be rebooted.

So usually say, "Yeah, I'll reboot it during the night when I'm asleep, I don't care." It communicates quite often with essential server, maybe the form of service of course, using some TLS we use on cryptographic keys which are stored in the chip, so it's quite smooth. And if you get the device room, there is one thing that you have to run on it, which is data collectors. This is one data process that has to be random.

And again, it's an open source piece of software so you can actually see what is it analyzing, how it is doing, and so on. And of course, we improve a lot of things in open [?], especially IPv6, DNS support, the – you have to change password immediately when you install it, the same

for WiFi. And one thing, which is not obvious, if you do a DNSSEC, and it is mandatory to have a precise clock, so it's quite strange.

We spend a lot of time just, you know, thinking how to bootstrap the device, if the clock is not okay and so on. So what do we collect from those devices? First of all, firewall logs, there is IP based firewall and we collect all the logs so we see what attacks are going on, but those things were not efficient of course. We collect router logs, things that might be problematic, upgrade status, some software problems. That's more for us to react if something goes wrong. Also the something physical measures, mainly temperature. You know, it's completely in your hardware and we are not a hardware based company so we have to be sure a number of things are okay.

And the main thing is the collection of source data for the IP anomaly detection. So we developed software called [Micro Collect?] or uCollect, if you wish. And it's model of software based on PCAP that collects data from the router. We listen just on the one port, so we don't want it to, of course, the traffic inside the line of the end user. We promise to monitor just the one port, and we don't go deeper than the packet header, so we don't look into the payload which – we stop at the TCP or UDP level.

The software has currently main modules. One is called count, it's very simple. It just counts the statistics of UDP, TCP router, IPv4 versus IPv6. Just basic stuff, mainly not for [?] but for the user to show how many data go through the router, and also it counts how many data are user generated and the router generated because we set certain limits of data that we can use for our measurements.

And the more important module is called buckets. It just, you know, collects, looks at the data, tries to find an anomaly using the similar technique I described, I think, in Toronto. We don't have much time to describe it, but it's very sophisticated technique, how to look at some data stream and try to find even a very tiny anomalies, just a few packets that doesn't match the normal pattern.

So, the probe tries to find those anomalies and post them to the central server. And, of course, the central server collects more data of such kind. It seems that something wrong happens, and we have to look at it. So, that's it. There are some graphics that you can see if you have access to the Toronto server. This is, for example, from my home probe. You can see in my home, I use 25% of IPv6, and there are some anomalies that the probe detected in my network and reported to the central server.

The good thing is that we do not collect who send the data, once we receive the data we just sort them, we don't have to know which probe sent it, it's not important for all data analyze. So, the privacy issues is something, we of course, saw very deeply. The central portal, we communicate with the users on it, there is a forum, there is support, and stuff like that. There is a lot of graphs like that. This is, for example, graphs which shows what kind of traffic and which are filtered out by my firewall. And there is a lot of tutorials because we would like to bring some additional values to the users to make it sweet for them.

So there are tutorials how to change [?] into [?] a file server. The tutorial how to change [?] into print server, and stuff like that. So there is a lot of other things that, so the users can a little bit play with it.

Everything, this is in the end user agreement, which everybody has to sign. So the agreement says, we will give you this for three years, almost for free, and after that, the device is going to be completely you, so you will get it for free.

It has to be your main gateway to the Internet and you cannot switch it off. There are, of course, so limits if necessary. We are allowed to test reachability of third party sites, that's for the network monitoring. And the uCollect process is mandatory. Other than that, everything is possible. End users are allowed to get root access so he or she can modify it, play with it, and we are very liberal in that so an end user can completely modify the system and even some hardware modules.

There is one PC express free slot, so that can be any other addition to the device, anything else. Privacy issues. Something that probably comes to my mind immediately when we talk about something like that. So many are described in the agreement. We keep separate database for accounts for the list of users and data. We never connect those databases. And again, if possible, the data are itemized, so nothing we need for the analysis is lost. We have strict retention policies. We just retain the data for a very short period, just two days.

And after that, we just keep status signal outputs, nothing like that. We cooperated very heavily with the local personal data protection authority about that, because we know that it's a very dangerous area, because widely with that. And because everything is open source, because we declare everything in advance. Surprisingly we got the only positive award in Big Brother award in the Czech Republic, which was quite surprising for us, but it was very nice.

And again, we never go into the panel, so we just stay in the packet headers so the limit is TCP UDP header for us. Current status, we distributed just 10%, so we are at the beginning so don't expect any results today. We are working on improving the detection methods. We are learning. It's a big project at the beginning and we are learning many things, so we also working on an operating system improvements. We got some feedback from users, this is not good, we need to improve this and that.

And again, we are working on tutorials. So how to change Turris into FAT server, or DLNA service or something that streams videos to television, VPN and concentrator, because again, not many root [?] VPN, so very good for people that understands what that means actually.

So what's going to be the future, we are waiting for first results. After we will spread the [?] network, we will collect some data, I hope I will be able to report something else to you. We plan to make another batch of about 800 of them this year. Since, VDSL is the main technology in our country, we would like to make a media interface, because it's not very, it's not perfect if you have VDSL modem and then the router, so we have two boxes in our homes so we would like to make this better for the users.

As many improvements as possible, of course, operating system, the probe. And also something which is very far, very distant future, hopefully we will make something, more universal system for more routers. We hope we will be able to push the firmware into some other hardware, of course without the probe, because this is resources consuming thing.

But just very, tune system, I know. This is the last slide. With this automatic updates that just works, and has also the amplifier. And maybe make some market for third parties application. We also we already discussed with some of antivirus companies, that they are attracted by this project.

And again, to create something for the users. So for example, how to connect your camera into the router so you can look into your home if you are not there, for example, or something like that. Some home automation. So, that's it. I'm sorry I couldn't bring it today here because it was too distant travel for me, but I will definitely bring one piece to London and also to RIPE meeting if there were some Europeans here, so you can touch it.

It is very nice. It has very fancy [?] works, and it is a really nice piece of hardware. So that's it. Thank you very much. [Applause]

EBERHARD LISSE: Sounds like you must wing it until your next meeting in Japan, so they can [?]. Any questions? We are a little bit behind so I'm only going to allow one question. Warren?

WARREN KUMARI: Warren Kumari. If anyone wants to phone, come find me. That's not a question.

ONDREJ FILLIP: Thank you for the question.

EBERHARD LISSE: Is it a nice phone?

LOUIS: Louis [?] from Costa Rica. Ondrej, did you check other hardware available in the market before building your own hardware? For example, a device like…

ONDREJ FILLIP: Yeah, we checked them but the problem was usually memory. Even CPU is much slower than we use. The memory, they have like 128 megabytes and we wanted more. So we told those guys, we said, "Can you manufacture 1,000 pieces for us with bigger memory?" For example. And they just said, no, 1,000 pieces is not enough for us, so we will not make a special batch for you.

So that's why we decided to develop our own hardware.

LOUIS: One more. What about the cost?

ONDREJ FILLIP: It's pretty costly. As I said, it's manufactured in Europe; designed, manufactured in Europe, so a cost of one piece is roughly… We never count it precisely, but it's roughly 150 Euros, so it is a quite costly device. But it's very powerful. It's really like a PC. Yeah.

EBERHARD LISSE: As I usually say in this context, you're a non-profit.

ONDREJ FILLIP: Yes.

EBREHARD LISSE: Thank you very much. Next up is Duane Wessels.

DUANE WESSELS: It's on now? Okay. So my name is Duane Wessels. I'm here to talk about the getdns API. This is a joint project between my company, Verisign, and NLnet Labs. And I apologize. I didn't get the memo about drooping ccNSO from the cover of the slides here, so that's on here. But if anyone wants an update, I have an update without ccNSO on there.

So this is about an API to make DNS queries. For anyone who has done any kind of programming on Unix, usually you would use a function called get host by name, or maybe get data info, one of those kind of things. And so, those interfaces had been showing their age recently, and as we all know, the DNS has really evolving with lots of new features. And so this effort grew out of that, as a way to give people a new way to make DNS queries from applications.

The API that has been implemented was actually specified by our friend, Paul Hoffman, who first documented the whole thing and then released his API specification under a Creative Commons license, and then Verisign and NLnet Labs were the first to implement against that specification. And one of the nice features about this API is that you can

use it both as a stub resolver, so sort of a dumb stub resolver that sends queries to a recursive, or it actually can do all of the recursion itself.

It can function as its own recursive name server. This specification was really designed with applications like web browsers in mind where you have sort of very large applications that are doing a lot of, well, a lot of DNS lookups and they need a high level of asynchronicity. They need to be able to do DNS lookup sort of in the background.

Another design goal was to leverage all of the new features like DNSSEC and [DANE?] and other things. The current specification of this document, of Paul's document, has been sort of updated based on our first attempt at the implementation. Oh, that's nice. So this is sort of the mission statement of the project. I don't know why it's sideways. Hopefully the rest of the slides won't be sideways, but the goal here is to make it really easy for non-DNS experts to access all the powerful new features of DNS, and do it in an asynchronous manner.

Aw, man. You're kidding me. Does anybody know? So like I said, this was a joint implementation. This is released under new BSD open source license. The first public release was made just a couple of weeks ago in conjunction with IETF, that's the 0.1.0 release. Here are some URLs where you can get the code, you can fork it on github, you can go to the website and read the documentation and subscribe to mailing lists and things like that.

So like I said, one of the major features was the need for asynchronous DNS calls, and in order to do that, the API lets you sort of hook in with a number of event libraries. Libevent is one that's very popular, but there is also a couple of others that are supported. And when you sort of fire

up the library, you tell it what event processing library you're using in the backend, and it all sort of works magically for you.

We have full support for DNSSEC. We get that, of course, because this API is built on top of Unbound from NLnet Labs. There is some good attention paid to IDN handling. And as you can see, here is the list of sort of supported platforms at this point in time. Linux, MacOS, FreeBSD, there is a rough implementation for Apple iOS devices, and coming soon are Windows and Android.

So I alluded earlier to the fact that with this one API you get these two modes of operation. You can have the library work as a stub resolver, where it will simply forward queries to some other recursive name server on your network, or perhaps not on your network. In the stub resolver mode, the application – there is sort of no cache built in with the stub resolver mode, so the application can choose to do its own caching if it chooses.

Alternatively, when you start up the library, you can tell it to function as a recursive name server, and it will do all the work itself of chasing down referrals and getting all the data it needs. In the DNS API, there is this thing called the context, and so the context controls which of these modes it operates under. If you want it to have both modes in your application, for example, you would create two contexts, one as a stub and one as a recursive. That's certainly allowed to do.

Oh, and one other nice feature is that, if you tell it to function as a stub, but you tell it you want DNSSEC, it will, sort of as a special case, do its own iteration maybe without sort of telling you that it's doing it, but it will do its own iteration to chase down all the DNSSEC signatures and

chains of trust. So be default, in the stub mode, it doesn't enable DNSSEC validation, although there are a number of ways that you can do that. The reason for that was really because, if you're functioning as a stub, there is sort of no trust between you and your recursive name server, unless you do something like T-SIG or SIG zero or something like that.

And so the system was made so that is probably not going to happen in most cases. So if you really need DNSSEC, you probably want to work in the recursive mode. You can also turn on DNSSEC, either on a per request basis, or in a per context basis. So here is a few examples of functions in the library. The get DNS underscore general function call is the most, as you can guess, the most general way to access the library. Here you can query for any type of data.

And as you can see, the function arguments here are, the context that you've already created, the name that you're interested in, the type. There is an extensions argument which is where you would specify various E-DNS zero or similar types of extensions. There is a port transaction user argument, a transaction ID, and then a callback function.

So the callback function is where you would get the response back from the library. This is sort of similar example. This is where you're asking for an address, either an IPv6 or IPv4, or both types of addresses, and as you can see, these two functions look very similar. The main difference is just the emission of the type argument here. So when you use this interface, the library, you know, it knows what types to query for and it handles it automatically.

This is the flip side, this is querying for the name. You know, give it the address and get the name back, very similar. Again, there is a specific function where you can get SRV records from the DNS using this form. So we had a lot of good feedback on this after announcing it at the IETF. There have already been a number of forks on the github repository. There was a package built for Redhat Linux within a day of the release.

Like I said, it's on Mac. In discussion with a number of parties, we've had really good feedback from these companies listed here. There is a release 0.1.1 plan this week. It will have no dot js bindings and Python bindings are also in the works. There is a hackathon coming up in a couple of months.

I have to confess I don't know what TNW stands for. Does anybody know what TNW stands for? I think it's in London, right? Or is it in the Netherlands? Yeah. Okay, so. Jaap knows about it.

JAAP AKKERHUIS: It's quite of a bigger Internet event. It takes all week, so.

DUANE WESSELS: Okay. So, if you need those details, I can help you track those down. I just don't remember off the top of my head what TNW is. So a few things to still be implemented. There will be support for the MDNS and NetBIOS namespaces. There is not yet support for suffix search lists. There is not support for keeping TCP connections up.

And there are a few EDNS OPT extensions that aren't implemented yet, but all of these are being tracked in a read me file.

JAAP AKKERHUIS: TNW stands for The Next Web…

DUANE WESSELS: The Next Web, okay. Great, thanks. So every effort is made to limit the number of dependencies, but there are a few dependencies for using this library. Of course, there is libldns and libunbound and those, in turn, require openssl. There is libexpat and libidn for, obviously, IDN names.

Here is a list of some of the core team contributors. I won't read them all to you, but these are the people at Verisign and NLnet Labs and other places that have worked on it so far. And if you have questions, I'd be happy to take some here, and there is also websites and the code repository.

EBERHARD LISSE: All right. Thank you very much. It's always nice to hear something on the leading edge. Any questions? Peter Janseen? No, Peter Janseen is the next speaker. Any other questions?

UNIDENTIFIED: Is this project related to Unbound? The utility Unbound?

DUANE WESSELS: Well, it's related in the sense that this project uses some code from libunbound, but it's – they're separate software packages. And it's related in a sense that some people work on both projects.

UNIDENTIFIED: Unbound is a really good tool. I'm using it as well for rezoning stuff.

DUANE WESSELS: Yes. Yes.

EBERHARD LISSE: Jay?

JAY DALEY: Jay Daley, dot NZ. I may have missed this, but has there been outreach to the browser manufacturers?

DUANE WESSELS: Yes. From the very beginning, they were involved in the specification, and now that it's out, they're certainly aware of it. And we're sort of hoping that events like this hackathon will sort of get them inspired to start experimenting with it and coding with it in the browsers. Yeah.

JAY DALEY: All right. Thanks.

EBERHARD LISSE: All right. Thank you very much. Just quickly upload another presentation. And then Peter Janseen from EURid is up next.

PETER JANSEEN: Good morning or noon. My name is Peter Janseen from the dot EU ccTLD registry. A bit different topic than what was on the agenda up today. I'm going to talk about what I call, or what we call, dynamic domain name zone provisioning.

First of all, DNS servers are everywhere in the world. I just tried to make a small grouping, let's call it like that, of what name servers are actually used for, on the Internet. On the top you have the non-authoritative name servers, which is actually the name servers that are being used by end users to actually find the answer of what they are looking for, IP address of website, where to send email, and things like that.

And below that you have three [?] that are actually authoritative name servers, root servers, TLD name servers and then the next levels. Basically they're all the same except for they would have some, what I would call, semantic differences. Name servers has only one relatively small zone, infrequently updated, but it's very reasonably high query loads. TLD name servers is basically the same, although that's a zone file depending on the TLD. Obviously, it might be a bit bigger than the root zone.

Mostly frequently updated, and also mostly having a lot of high query loads. And then the second level and next level domain names, which are mostly smaller zones and mostly infrequently updated. But the point is that there are a lot of these zones being added and removed on a day to day basis. And again, depending on what the popularity of the domain name is, it might be having a high or relatively low query load.

That said, let's dive into a bit of concrete examples of how you might provision a DNS zone. And I just had an example here of some parts of a configuration of BIND, one of the name servers that is out there. The first part is what we call a master name server, the second part is what we call a slave name server. Basically what it says there, the part of the configuration is either the master or the slave, so either you have the contents locally or you get them from a master.

And it specifies where the files sits on the file system and who can transfer and things like that. The zone file itself, would look something like this, with NS records, A records, and things like that. On the other extreme of the configuration possibilities of DNS, I just list a part of what power DNS users as a configuration possibility. Basically power DNS allows you to actually store the configuration of your zones in a MY SQL database, and what you see there is where the My SQL database, which user password should I use, what should I use to actually collect through the database.

And basically the whole zone and all the records, what you just saw in BIND, a zone file syntax [?] will actually sit in tables in a SQL database. So basically, the two extremes are either its content config files and then you're faced with generation of these config files, distributing these config files to the different name servers. Or you have a database backend in your name servers, and then it's basically all about inserting and updating entries in database tables, and obviously making sure that all the different databases are replicated or are in sync with each other.

Basically the whole question is, is it configuration or is it data? Configuration typically sits in configuration files, while data typically sits

in an online database. Now, a zone being added or removed, and is it configuration or is it data, that's basically what we're talking about. So if you look at dynamic provisioning of DNS, I listed some of the goals which probably are, to some extent, important for all people out of there.

Performance, obviously, is always an issue. The ease of provisioning is a general remark. It should be relatively simple to do this. You should have some implementation flexibility, which is another word for [?]. You should be able to swamp in one name server with something else without too much hassle. Security is obviously a hot topic, both the protocol itself as well as what I call entry points. The more software and the more modules, and the more things you have, on your name servers, or on your servers in general, the more venues that are open to attack.

So the less, the better. These days you have to mention the fact that you have to support DNSSEC, so that's why it's there. And in general, the simplicity of the software setup, the easier it is, the less, the simpler it is, the easier it is for everybody. If you look config file based provisioning, the flow then will be generate a config files, or a set of config files; distribute them to the name servers, and I just listed there a few possibilities like Chef, Puppet, CFEngine and rsync. And what have you note a zillion other possibilities that people have used in this world to actually make this happen.

And then obviously once these files sit on the respective name servers, you have to have a mechanism to actually force these name servers to reload the config files, or load the config files. And again, I mentioned a

few of the things like cron, gearman, Signals, socket interfaces, and all sorts of things.

I would say the challenges there are, if you have different name server implementations, most probably the layout of the config files are different per name server implementation. If you look at how NSD does the [?]… They're all sort of the same, but also sort of different. The reload command is probably different per name server implementation as well.

You need to have support to be able to reload config files without stopping and answering queries, because otherwise, your whole name server infrastructure stops answering queries when you want to reload config files. And again, if you do this, there is an extra layer of software for distribution and command.

Again, at the other extreme, if you look at the database backend provisioning, it's all about inserting and updating records into database tables. The challenges there are, well, you have to have databases running, which was not necessary if you are config based, file based, sorry. You have to replicate somehow.

Performance might be an issue because you have other software running. Databases software running databases are taking a lot of CPU and IO to actually run. And interoperability might be an issue because they're aren't so many name servers out there that actually support SQL database backend.

On sites that… If you look at what has been running for years now, it's something called dynamic DNS. So how do you actually update the

content of a zone file? And this is in preparation to actually go where we want to go. You have a name server running, which is the master, which gets on the upper top left of the dynamic update message, saying, you know, at this resource [?]…

The master server will send out, notifies to its slaves, the slaves will receive that and it will come back and request an AXFR or an IXFR, so incremental or zone file transfer. All this is in band, it's self-managed, there is no other software. It's interoperable because it's ITF RFC standardized. And basically the updates happen without downtime, so it this is actually all good things, in my view.

So if you want to look at this in a dynamic DNS provisioning way, and this is sort of a messy slide, I must admit. So, if you look at the configuration data zones, it's basically, you can look at it as data in a meta zone. So on the one hand, you have the IP address to aid like for a host name, wwww.some domain.eu as an example. It might sort of say, well, this is a key value pair and actually the configuration is basically a set of key value pairs as well.

So basically you might say, well the provisioning of zones might be considered as data in a meta zone. So why not use a DNS channel itself to communicate this configuration data to the DNS servers? The advantages will be the name servers at startup, they don't have a configuration that basically the only thing that they have is, well, who should I trust? Well, basically that should be a TSIG key, ACL set of something, something like that.

So if you look at, and this is impressive. So this PowerPoint, it actually works, even on a Mac. This is, okay. So if you remember the dynamic

DNS slides, two slides ago, if you would put this to a higher level and actually do DNS provisioning over DNS, you would have a provisioning, a master provisioner setting somewhere, that would receive a dynamic update message that would say, for instance, add some domain.eu as a zone.

And by the way, NS1 is the master; NS2 and NS3 are the slaves, and all sorts of extra configuration. The master would receive that and actually send out notifies to slaves, slaves would actually respond and transfer in the configuration from the master, and will actually self-configure to become, in this case, NS2 and NS3, a slave for the domain, some domain.eu. And config that the NS1 master should be the NS1 master for this domain.

And then the rest would have dynamic update messages as always would happen. So basically, it's a two-phase process. First of all, dynamic update messages adding and removing and configuring zones, and then the rest would just happen with adding and removing resource packets to the zone files itself.

That said, the dynamic update protocol, it's called like that, the dynamic DNS messages are standardized. This is one small extract of an RRC. Basically, this is a resource record which has some config fields, or fields. What is important here is that, to ask normally, is the Internet class… We looked at this and said, well, let's put another class there, for the moment we call this the control class, which for obvious reasons we have used 2a in hexadecimal as the number, the magic number.

For those that don't realize, it's 42 in decimal, and as you know, 42 is the answer. So why not use that as a class for the configuration? All the

rest stays the same. Then you would have to have some new resource record times. I just listed a few of these, and I'm coming back to that in the next slide. One would be, for instance, a master resource record type that actually allows you to specify the configuration for who the master is.

That would be some flags that would be an IP address, that would be a port number, that will be the name of a T-SIGN key. The other example I gave is NTRFC, which I, for the moment, can't remember what the acronym is for. It's a numerical transfer count, yeah. Okay. So it's the number of tries that a slave should try to contact its master to transfer its zone. And again, it's one of the parameters that you put in a bind config file, for instance.

Again, you have to invent a resource record type for that to communicate that from the master to the slave. Basically, what we have been doing is in our labs, we have something called the YADIFA. The YADIFA is an open source, clean implementation of a name server. So next to the obvious choices that you have in the world like [?], and we came up with another version of another implementation of a name server.

We added some extra stuff to this. Most importantly, the second that, the dynamic provisioning extension. We added some command line tools, some libraries, some API. And the important part of this that we defined the protocol of a preliminary protocol to actually allow configuring zone files over DNS. So, we're well within time, I would say. The whole idea is that you would have a standards based, interoperable, DNS message based, dynamic zone provisioning protocol.

And if anybody is interested in talking about this, there are some things there where you can contact us and to discuss this. Thank you.

EBERHARD LISSE: Thank you very much. [?].

JAY DALEY: Jay Daly, dot nz. Thank you Peter, that was great. What are your plans for documenting and publishing this, and seeing if we can make this into any formal standard?

PETER JANSEEN: Yeah. We have been looking at this, and our first step was to have some running codes, which we have now. It's not in a very stable version yet. We have been doing some demos. We have been playing around with it. But the whole point is to come up with something which is relatively stable.

And then actually, one of the end goals probably should be to have a set of RFC, another set of RFCs out of this, because if you want to have something which is interoperable with the rest of the world, it would be best that everybody would be implementing the same standard. So we are planning on spending some time on drafting some RFCs. And whoever is interested this to actually help out, I would be most interested in hearing from you.

JAY DALEY: Thank you.

EBERHARD LISSE: Any other questions? Okay. Thank you very much Peter. You can give him a hand if you want. [Applause] And the next one is Hasnul Hasan from dot my if I'm not mistaken. They have got some mobile DNS validation, which Don Hallander recommended. So let's hear about it.

HASNUL HASAN: Hello. I'm Hasnul. I'm from MYNIC in Malaysia. So today basically, I'm going to talk about our experiences for the past year, and how we actually did things to improve the security of our mobile DNS validation.

So in 2013, we had two major incidents. One happened in July, and the other happened in October. In July, basically, the hacker was targeting our own systems. So, there was a lot of lessons learned. And so we [?] improve our security. And in October, because we had to improve our system security, they targeted our end users.

So that's where the interesting method comes about, when we talk about the security of the DNS. So, in October tragedy, what actually happened, when we actually improved our security, they went on and basically attacked our users. Basically what happens is basically they start to attack the email account of the users, because what happens is, in our system, if you forgot your password, the standard modus operandus is basically leveraging our email.

So you say you put in your email, and then basically we'll send you a reset password capability, and then basically you can reset your password. Then you can continue managing your domain. But what happened is that they target the end user that they want to target, they

compromised the email account, they went to the MYNIC site and basically request a password reset.

Once they had done that, they got their email validation, and then the rest is history, because basically they basically circumvent the loop, where basically if an user forgot their email, or forgot their password, in the managing of the domains, they then basically then have the capability. But that is not enough. So what happens is, when we look at it, it's basically when they do all of this stuff, things got – their angle is to modify the genesis of the information.

The October was done by a Pakistani team, while the 2nd of July one was done by the Bangladesh team. So, even though we improved our security on our web systems, the problem is basically there is an issue with the end user. If the end user does not maintain a certain level of security, the system also will be perceived as not secure. So the hacker will basically go and find the weakest link.

This happens to a lot of banks in Malaysia. As last time what happens is basically once the bank start to improve their security, they start to target the weaker banks. So they don't go to the stronger banks, they go to the weaker banks. So on our side, we look at it that, just looking at our own security, it's not enough.

And because the expectation of the end user, and the perception of [?] is different. End of October, in [?] when we look at it in retrospect, a lot of the perception is that MYNIC itself got hacked, actually it didn't. All the records, all the changes are actually valid changes, but perception is very important because therefore this key in the market, and we want

to be portrayed that we are trying to improve and trying to maintain a secure registry system.

So, due to that basically what we do is, what we are implementing is a multi-factor authentication system, which similar to what is implemented in the bank systems, where we go and find an [?] validation capabilities. So, I mean, this is just very basic. What is basically multi-factor authentication. So for this who are very experienced, you know that basically the first factor is basically what you know, your username and password.

The second factor is what you have. It could be a token, your phone, a smartcard, your certificates. And a better factor is what is unique about you. Your fingerprint, DNA, retina, or any other matter that is unique. Multi-factor is basically leveraging on multiple types of factors to create a more secure authentication of authentication that – to minimize the risk of compromise.

So with multi-factor, if one is compromised, there is obvious way for you to protect the account or the system. So, in our case, we actually implemented two types of authentication. One is via SMS, this one is basically using a computer authentication, so we can validate basically the computer that you use. It's basically authenticated.

If it's a new computer that is used, you actually ask for a new SMS code. So what will happen is basically an user, let's say they didn't know someone got their account and request some changes, or some authentication, still get a SMS saying that we are trying to login in this so-and-so machine and so forth. And basically they then know someone actually got their maybe email account or something like that.

The second one is basically smartphone based. It's using IOS or Android. So what we are looking at is points of protection, authenticate authentication process, the password recovery process, and also when there are requests to change sensitive information, domain information, technical information, or any other information that basically effects the operations of the end user.

This is basically just a simple screenshot of how it looks like. So in the app, basically, when you load the app, it asks you for a pin. And then basically, you just turn it on. If you see from, see like, I don't know if I can do this. If you say from here, it's just basically if you turn it on, then it connects back to the hub system and tells it actually, okay, this user is online, and this user is authenticated.

For the second one, it's more on the transaction. So if you can, I don't know if it's very clear from that screen. Let me see. If you see here, these are the information on the changes that they wanted. So what happens is when you change something, you'll get – the app will tell you, okay there is a request to make these changes. So these changes can be the same percent or could be an authority that is supposed to be approving these changes.

So this allows basically another level of security so that basically, if people were to actually get your passwords, username, they won't simply be able to just change it. So we're trying to basically improve a bit on the current level of security that we have in MYNIC, to offer these types of capabilities.

Other matters that are important, because we never keep mobile numbers for the users, so that's pre-registration process that we need

to do. The delegation capability in Malaysia basically many of the domains are actually delegated to a domain administrator. Not many of them are actually managing it themselves.

So the delegation capability for end users, for outsource is important to be incorporated into the systems. The awareness of why this is needed, and also the additional point of, from a help desk point of view, because this is a new way to, at least to our current users, a new way of doing things. And the integration, because it depends on what you want to protect. If you just want to protect from authentication point of view, or you want to create multi-level authority.

So that will increase in complexity. So actually, basically, that's, that's basically my presentation. Thank you.

EBERHARD LISSE: Thank you very much. [Applause] What I'm, if I understand this correctly, so if a registrar makes changes to a client, the client gets informed? Or the registrar gets informed?

HASNUL HASAN: Okay. In Malaysia, basically we are the register and registry…

EBERHARD LISSE: Okay. So it's between registrar and the end user?

HASNUL HASAN: Yes. So in Malaysia, we have resellers or domain name administrators, which is acting on behalf of the end user. So we have those two types of end user groups.

EBERHARD LISSE: On the registry, and I of course want to have nothing to do with end users, if I can avoid it. We're automatically integrated registrar that has about 12 clients list, and every year we lose one, we're very happy about it. Yeah, because of exactly this, yeah? But I think it's a good thing. The source code, is it available?

HASNUL HASAN: Actually we didn't develop ourselves. It was a local partner who actually developed it.

EBERHARD LISSE: The point is, yeah, it's an app for the iPhone, it has to go on the thing, but if it was available for other, tool was available for others, for example, [?] ccTLDs, for other registrars. Would you make that available?

HASNUL HASAN: Since it's on the commercial arrangement, that's what's happening. So we work with a local partner who developed this locally. It's deployed in one of the banks of Japan, also. Basically they say it's about 20 cents per user, that's the cost.

EBERHARD LISSE: Okay. So it's commercial. Warren?

WARREN MURRAY: So Warren Murray. Three questions. First off, this presentation isn't posted online, is that just an oversight?

EBERHARD LISSE: Ask Christina. It will be posted online.

WARREN MURRAY: Cool.

HASNUL HASAN: That's my apologies, because I submitted it at 4 AM.

WARREN MURRAY: Not a problem. Then also the October incident, you said that the user lost access to their – their email got compromised. Can you clarify who the user was in this case? Was it like a large registrar? Or was it end users? Because a bunch of domains changed at the same time.

HASNUL HASAN: In Malaysia, many of it are reseller who actually manage on behalf of the end users. So their email system actually got [?]. So it was a targeted attack through [Google?], so they basically studied the background, who was managing and so forth, and they targeted that user first.

WARREN MURRAY: Okay. So the reseller of the thing.

HASNUL HASAN: Yeah.

WARREN MURRAY: And then the last one is, do you guys have a 24 hour knock, or emergency contact thing? Oh, sorry, I forgot to mention, I'm with Google so I have some knowledge with the incident. It took a long time for people to be able to reach you guys, and I don't know if you have a 24 hour emergency contact or something that registrars have.

HASNUL HASAN: Okay. We do have. The only thing maybe, we need to improve the collaborations, because usually what happens in Malaysia is we look at contact first, then the resellers. We contact our commission, and also the cyber security. I think we need to put in there basically other measures so that it would be easier for us to collaborate.

WARREN MURRAY: Cool. Thank you.

JAY DALEY: Hi. I'm Jay Daley from dot nz. Do you use external security auditors to test systems?

HANSUL HASAN: Yes we do. So we basically, we also [?] certified. And we basically employ [?] for the security assessment, [software?] review. Actually the second incident that happens, was actually [?] so hackers welcome for me, that happen. But during the July [?] there is a lot of things that we saw that basically, big gap. Which specifically they then took steps to further improve. So we had basically we called in the experts to even review our source codes, to do basically a [strategic?] posture, policy [?] network reviews on the lots.

JAY DALEY: Okay. Before the first incident, you didn't have those, and then you implemented it after the incident?

HANSUL HASAN: They were having it, they were [?] into ISO 2700 [one?]. So during that mid-stream, that's what happened. So, that's why when the second time they said it happened, it wasn't more on our side, but then we realized that that was another chain of events that happened before that.

JAY DALEY: And would you recommend to other registries that they should all have third party security assessments conducted?

HANSUL HASAN: I believe from experience, we need to have internal team and also external team, because it's like [?], so sometimes we seek – because we

are so used to our own systems, and we tend to overlook certain things. So we need sometimes, another team's perspective.

JAY DALEY: Okay. Thank you.

EBERHARD LISSE: Any other questions? All right. Thank you very much. Next will be Mehmet Akcin, who we all remember, fondly from implementing the DNSSEC into the root. He's now working for the enemy, Microsoft.

MEHMET AKCIN: So hi guys. My name is Mehmet. I work for Microsoft right now. And so forget everything you know about Microsoft. Next slide. Please. So why am I here today? Microsoft has been engaging in several technical communities, as you are aware, Microsoft has a big event called Tech Net. It's a global event, it takes in several places around the world.

And also, it has been engaging in ICANN community in technical community. Why I am here today, is to actually introduce Microsoft's – I'm actually going to read this, introduce Microsoft's experience in operational capabilities to the top level domain operators. I don't want to limit this to the top level domain operators, I want to expand this to anyone else that actually runs DNS.

In order to explore possibilities to collaborate [?] security and resiliency of global domain name system. So next please. So the most important thing is, we don't really know what you guys think. Today, my biggest

reason of being here to listen. Listening is the most important thing that I want to achieve in this meeting.

I don't know the day to day operational problems that you face. I did run a TLD in the past, and a root name server for 10 years when I was at ICANN, but I don't know the day to day problems that might be related to registries, registrars, some sort of relationships, databases. I don't know. But I'm here to listen. I'm here to listen if you have problems in the name server platform, within Microsoft products. I'm here to listen and learn.

And learn your experience, and see what kind of expertise we have that we can help you succeed what was missing, and help. So I'm going to talk about three different things, but the most important thing here is the blue one, your idea. I want you to help us help you back, because again, I want to repeat it, I'm here to learn. I'm here to learn from you to be able to deliver you what is the best product out there.

It's easy to reach at Microsoft. For DNS related things, it can be operational need, it can be the DNS software itself. DNS at Microsoft, that's it. If it is DNSSEC related, we are actively working in several platforms to make DNSSEC available. If you have any questions, we have recently published a step by step guide. I don't know if you follow the DNSSEC deployment guide, mailing list. Let us know by just contacting DNSSEC at Microsoft dot com.

And the vulnerability scanning. So, we have this service that we announced two ICANN meetings ago in Shanghai. I was actually working at ICANN during that time. And what this service does, is basically, we sign some sort of agreement, and it basically scans your

infrastructure and we have some scanning abilities that can scan for bogs that are not necessarily released yet, but there are known bugs for us that we are working on fixing.

So this service actually can point out lots of things. I'm going to talk about that more. But again, the most important thing in today's, my presentation is going to be your idea, not our idea, because again, we don't know, we're here to learn. Next please.

So, the DNS team, the team that writes the DNS software in Microsoft is very enthusiastic. I'm going to give you guys one example. So, I started and I'm an IETF guy, I look at the RFCs, and I was like, "Okay, we don't have NSID." The guys were like, "Okay. So what is it? Which RFC is this?" This RFC. And one week later, here, update. Now we support… It's just they are not necessarily aware of everything that the community needs.

Again, this involves your input. Please provide us any input you want about the protocols, about the abilities that this software needs to support, and we convey that to the right team. Just email DNS at Microsoft dot com. I'm not going to go through and make this like a sales presentation. All I want to do is point out, think about diversity in your infrastructure.

How many people here runs TLD? I'm going to make… I'm sure you have thought about diversity in the past, and I'm pretty sure you've never thought about Windows. Think about it. Give it a try. I am not necessarily saying it's the best, but my goal is to make it the best. It's pretty good. That's all I can say.

How many people actually knew that Windows 2012 R2 Server comes with a BGP speaker? So it does. You know, you can have a single box that can do any cast. It can be just one that is out there that's doing any cast out of the box. DNSSEC obviously enabled, IETF RFC friendly. Anybody, in this room, this is a personal, not Microsoft, thing so just want to make sure that. Anybody in the room that comes and tells me, and direct violation of RFC, I will get them Xbox 360, myself.

I went through that list, by the way, just so you guys know.

EBERHARD LISSE: Can you make that an Xbox one please?

MEHMET AKCIN: If it is a direction violation plus a bug, Xbox one, just because you said it. There isn't any, that's why I have gone through this my first three months, but if you find one, please let me know. And it's a personal commitment, it's not Microsoft commitment, but you guys know me in these last years. What else?

Documentation. So we have right now, step by step guides. We're working on some kind of videos that people can watch and people can learn from. And then the next steps are actually going to be making these available in several languages. It might be spam, you name it. We are a global company. So wherever the customer demand is, again, feel free to send an email and suggest that, hey it would be great to have this guide in this language.

Czech is one of them that we are working on, because Czech Republic is one of the most DNS adoption in the world. Anything else? Capacity wise, performance wise, I'm not here to sell you guys a product, it's pretty good. Give it a try. If you are were a ccTLD, if you are interested in… The next slide please. I'm going to tell this on the next slide. So DNSSEC is well – I mentioned several things over and over. Are IETF RTC friendly. If you guys running a ccTLD, if you guys want to give it a try to DNSSEC platform off Windows, find me.

I will give you guys some licenses that you can demo the product, and we can work on – you have to be non-profit, that's the only requirement. But we can work on some things, and I can help you set it up, or I can point you out the right online guides where you can see a lot of information. We support, obviously, Active Directory and Dynamic DNS, automatic queue rollovers and SEC 3.

One of the key things that I have spent about a month when I joined Microsoft, was to make the GUI, you know Microsoft has to have beautiful GUIs, right? It has to be easy. Once we made that GUI, is that it's right now, if you have the latest, greatest update that was released on February, next, next, next, just works.

Meaning, next, next, next will just get your domain automatically signed and at the end of the day, you will have a fully signed zone that automatically rolls the key without you checking any box. At the end of the day, all you really need to do is sign this domain, and I think no matter – I have worked with open DNSSEC before, I work with everybody, it should be just easy. Happy to vouch for every other

software out there that makes it easy, and suggest that diversity, in terms of diversity, to use multiple of them.

Next slide please. How am I doing with time? Stop me. All right. Next slide please. So security. Why are the TLDs are target? Why? Well first of all, you can take down the whole company at once. You can take down the whole service at once. You don't have to worry about really, "Oh, did I attack their web server? Or do I attack their mail server?" No. You can take down by changing just name server record. You can change everything.

So next slide please. When we did our scan, I'm actually thinking just go to next slide real quick. Next, we will get that. So we actually are in touch right now with 21 TLDs today about these TLD security scanning servers that I was talking about. Seven of them we did full scanning. And we found 138 bugs. These are cross eyed scripting, [?] injection, open ports, serious stuff.

I'm not talking here about low level bugs that will actually not cause any problem, but rather things that might cause big issues. We can back just one slide. Yeah, this one. So how does this work? How do you get on? First of all, it's free. No money, nothing like that involved. The only thing is that we need to sign a NDA basically that we're promising that whatever data you give us, we are not going to be using for anything else but just for this scanning.

And then it lists all the other options. Why we are doing this, what's our goal. Our goal is, let's understand, ccTLDs, you guys run ccTLDs, but our business depends on it too. We have customers that are on these TLDs. And the other reality is, not only we have customers, but we have our

own domains. So we're trying to help you, indirectly it helps us. It helps everyone, if possible.

So, you can go to this website, AKA dot MS slash ccTLD scan, and you will be able to find more information. Or simply send email to that email address, if you can be on, you know, constant email forward back and forth. Feel free to talk to me, prior to sending emails, if you have some concerns.

Other than that, it should just work. Next slide please. Next slide. So we have not been present in many technical communities actively, playing a role where we sit down, listen, learn and develop, and deliver. We want to change that. These are the paths that we're taking right now. We will be in LACNIC 21, we will be in DNS [?] workshop.

By the way, if you have not submit, or if you're interested in submitting some presentations for DNS work, sorry, can I make an advertisement? Okay. DNS work, we just extended the time, so you can actually submit papers. Please do so, it's an amazing opportunity to talk about DNS. ICANN 50, ITF 90, and much more. Next slide please, and the last slide.

The most important thing guys, we are here now. I'm only here to listen, learn, and deliver. Most of you, may or may not be my customer today. My ultimate goal is to deliver the product that is better than anyone else. If you can afford it, we will find a way because Microsoft is very generous delivering these kinds of services for the good of the community. We will work together, for the better and for more stable Internet together. Thank you. [Applause]

EBERHARD LISSE: Thank you very much. I gave him leeway beforehand to advertise a little bit because he promised he would buy us lunch next time. So I will, Christina you can book him for the packed lunch for London already. Cocker Tools, the one that is run by a number of ccTLDs, has been participating, dot cx has been fully scanned.

I run the latest and greatest version two weeks behind, and I think all parts are close. But what is really important about, if anyone of us uses credit cards, or accepts credit card payment in any way, make no mistake, it's coming. They're going to close you down, they're going to close down your merchant facility, unless you comply with this credit card industry standard.

We have, I have been through this recently. My processor in England is called Global Payments, and they have got something called Global Fortress, which is another method to squeeze some money out of me, but basically now we are in compliance. They do probably the same scanning as they do, it's fairly simple. We don't, Crocker Tools doesn't keep any data online, so it's not a big drama, but you have to look at this.

If you keep, if you accept merchant, credit card data, they will do vulnerability scanning. If they detect it, you're done. If your business model depends on taking credit card payments, and you don't comply with the standards, you're done. Never mind that is a good idea to close all the security loopholes so that they can't mess with your registrations and your data.

And then there was a certain advertisement here that was making, that I asked him not to put in the presentation, Microsoft is hiring. So if any

DNS engineers are interested in taking up a job with the enemy, there here is. Any questions?

MEHMET AKCIN: We are actually building a DNS team that has experience of running real, authoritative name server fleet. As you guys are probably aware, we have our own top level domains, the new gTLDs that are coming on, but also we have very critical domains, like Skype. Who uses Skype? Can I see hands? Come on.

Who doesn't use Skype? All right. So, good. So all of these domains, all of these DNS are behind this team provides, we are hiring managers, engineers. Feel free to talk to me, or drop me an email. My email address is Mehmet at Microsoft dot com, so just feel free to contact with me, or to see me here. Questions?

EBERHARD LISSE: Now that we're doing with the advertising, any questions? No questions. All right then, I'll close this for lunch 15 minutes early. Let's be on time. From my previous experience here, I can recommend that one can go downstairs into the local food mall or something, and buy $10 worth of stuff. Last time I enjoyed this very much.

UNIDENTIFIED: We are going to wait a few minutes before we start. Feel free to sit down, turn on your laptops.

EBERHARD LISSE: Okay. Can we settle down, or sit down, or…? All right. Ron, can we sit down please? Okay. The next thing is what we call IPv6 round table organized by Leo Vegoda, who doesn't happen to be here, which puts me a little bit in an uncomfortable position. And I'm going to have my merry way with him when I get hold of him next time.

Mehmet has some experience in this regard, so he has been delegated by Leo to chair the round table. Sunny Chendi from APNIC, and David Summer from the [Bear?] Group. I'm going to present and talk a little bit. IPv6 seems to be a bigger issue in Asia than it is where I come from in Africa, where nobody cares. Yeah, but probably because the Internet penetration here in Asia is much, much, much bigger.

The limitation for IPv4 is an extra issue. In [?], nobody cares. Yeah, because as long as the button works, yeah, no one really cares about this. Anyway, without further ado, I'll hand this over to Mehmet. And the format is going to be that each of them is going to present, and then we're going to try to get some questions. And if there are no questions, than Jack [?] will have to ask the first one.

MEHMET AKCIN: Thank you. Sunny?

SUNNY CHENDI: Thank you very much Mehmet. My name is Sunny Chendi, I'm from APNIC. So we were invited to present in this session. We thank Leo and ICANN, and ccNSO as well for giving this opportunity. As you know, APNIC is one of the RIRs, one of the five RIRs, and we serve the APNIC region.

So my focus on this presentation is mainly looking at the deployments of in this region [?], and how we're doing it. In terms of not only making allocations in this region, we have to also measure these allocations and see how much the deployment is actually happening in the APNIC region, compared to the world.

So we do quite a bit of analysis of the data. Our chief scientist [Jeff Houston?] he puts a lot of time in analyzing, to looking at various statistics and producing some for us to share with the community. So I'm going to look at various aspects of the IPv6. But it's quite important that we look at IPv4, just to see where we are currently on the areas.

So this is the status of IPv4 exhaustion. You probably aware already that APNIC exhausted – IPv4 exhaustion occurred first in this region, and then followed by [?], but in terms of other RIRs, we believe, just a prediction, we believe it will be LACNIC that is going to exhaust soon, and then ARIN. In terms of where they're standing, just a prediction, we're thinking it's the end of this year, or at least in November that LACNIC will hit their last slash eight.

But we here from LACNIC that the rate of allocations is going quite rapidly now, and I think everyone is trying to get what is left there, a very quickly so they're crediting actually, it may happen in May or April, sooner than what Jeffrey is predicting. But we have to just wait and see.

As for ARIN, it will probably be early in 2015, next year. So from this, we do see that it is real, the exhaustion is really for exhaustion is real, and we have to continue to work on IPv6 deployments. So let's look at some data on IPv6 and see where we stand in terms of a multi-

dimensional aspect of this, including the transit providers, the content, and websites.

So if you look at this chart, the green one is the transit, the IPv6 transit AS numbers, which is transiting IPv6 and the orange one is IPv6 enabled AS. So, looking at this, we believe. The tier ISPs are 100% ready. And as we go down, we see that those who are, the content is still not up on the IPv6, they still need a strong push on making the content on IPv6.

And also, down the chain, the food chain, there is still not much IPv6 is in the process. Or is it in the deployed state? There is still a lot of capacity building has to be done in that tier. And we actually are constantly engaging with the community to provide capacity building and some workshops to make them understand the importance of IPv6 and the deployment of IPv6.

Looking at some of the statistics from other entities who does a lot of statistics out there. This one just shows the top 10 economies that their websites have already looked at. But the green color and orange color, the green color is [?] that exists for the primary domain. And the orange one is the color that exists for the second domain, so the subdomains.

So from this region, we see Singapore and India. And number one, they know what they are doing [?] they have a very steady growth in terms of making their websites IPv6 ready. Looking at the Alexa top 50, again, I think we need to still see a lot of content to be pushed to IPv6 as well. But the bottom line, you see the subdomain [the orange one?] is the subdomain and the green one is the color for the primary domains, which is reasonable.

So the progress is good. It's steady, and it's grown, and it's picking up, at least with the primary domains. But when it comes to the secondary ones, there is still a lot of work to be done. In terms of who is leading the IPv6 in the world, this is among the commercial operators, not including academics and the research center [?] networks. The development of LTs picking up, moving broadband to mobile, which is good to see, and most of this are actually happening on IPv6.

So [?] zone that deployed IPv6 when they implemented the AT in 2010, that's where they're updating, in the USA. And the Japanese company, Chubu Telecommunications, they deployed dual transition. That was a very good move. They are [?] one person now, preferred rate over the network is IPv6. Within ICANN, we believe there is a registrar activation agreement in June 2013.

This might come in force in January 2014. I'm not speaking for ICANN, but someone can clarify this one later. But as for this one, all the registrars accredited to ICANN, they have a mandate to move to IPv6 along with the DNSSEC and IDNs as well. We believe this is a good initiative from ICANN. We are hoping to see much more deployment happening in this area.

Some observations among the regions, economies and individual operators, it defers from local policies. The deployment does not happening all at once, it's a slow process. We understand that. We all understand that. But looking at the previous slides, you know, we are quite satisfied the way it's moving on with the IPv6 deployments. But the awareness is still important, we believe, at least, it's still important.

We still have the goal that end user level to create the awareness to make that business proposition, coming from the bottom up, moving like we want to IPv6. So looking at, in this region, we have a colleague from China. You probably know more about the deployments in China. But we understand that the China Telecom, China Unicom, and China Mobile, three major operators in China, they have clear mandates to move to IPv6, and it's very visible even though we don't see much from outside.

But the publish and growth and the broadband users, the mobile users over there, the amount of growth that's happening, you know, they don't have any other alternative than moving to IPv6. [C-NIC?] [?] from China, they also have discussed this, and they have created awareness, and they have organized summaries along with the BII.

To encourage the previous deployments in China, and they have shared – the awareness programs that they've been doing in China in other conferences as well. So it looks that we are all working for the same objective here. We don't get much measurements out of China, due to various reasons, but it looks like the end user [?] is what the global Internet is as difficult as we can, in China as well.

But one thing we can tell from this one, this chart is definitely some activities are happening and we believe it's the reflection of the IPv6 mandate by the Chinese government. Now the office of the government chief information officer put out via mandate in Hong Kong, they're actually doing very well in terms of moving to IPv6.

A lot of government initiatives there, you know with the Internet gateway systems, and… We work very closely with ISOC Hong Kong,

ISOC chapter in Hong Kong in creating, providing awareness and capacity building. And they have a very clear IPv6 in action for the 2012 event, but it's moving steadily in the 2014 and beyond.

So Hong Kong constantly working on IPv6, where the ISOC Hong Kong, the ISOC chapter in Hong Kong is doing a very good job, along with the government and the operators in that region. It's not up to the world average level, but even though it's quite small, but the industry is speaking up and they will see much more, better statistics, probably in the years to come.

Japan no doubt is the forefront in this one, IPv6 deployments. They have indicated very early in the stages, apologies. And the government have very good initiatives in Japan. I don't have to really speak much about this in Japan. You can see the statistics, they are very good partnership between public and private sectors. And same goes to Singapore as well. The deployment initiative iDA. You probably all heard the speaker from iDA this morning in the [planning?].

They have clear mandate, and most of the deployment that is happening in Singapore that is the mobile broadband deployment. You can see from the graph, the speaking publically, well the startup and Mobile One, they have toned down IPv6 in May 2013. And since then, you can see the graph is going up. That's the preferred rate by month, IPv4 rate. So very strong initiative in Singapore.

In Taiwan as well, we have NAR there [?], along with Ministry of Transportation and Communications. They've been organizing the well-educated programs, v6 forums, and the industry is going along with them in deploying IPv6. Otherwise, the statistics are not that great. But

moving along, I think with all of these slow initiatives, we believe Taiwan will do very well as well in this region.

So the good part of the Internet, if you see mobile cellular subscriptions, it's about 100 subscriptions, about 100 inhabitants in major economies in this region. The mobile deployments, the LT deployments, is the trigger for IPv6, moving to IPv6. To me, the demand of the globe patterns in increasing the business and the business strategies, these are the things which will drive v6 deployments to the mobile.

These are just a few examples from this economies. How many are the TLD deployments or the subscription rate among various economies. So, we believe that the LTE will pick up significantly, as we move along to mobile broadband. And this one shows the global LTE subscriptions will surpass one billion by 2018. So, that's the significant deployment will happen by then in IPv6.

LTE user devices. Going with smartphones, and tablets, and all of the Internet on the go, is the driving force, as we all know, and this is what is going to drive in the future. So in mobile network, the business competency of operators is very important. It's a technical and a business decision, it's not just a technical decision. It's very difficult to hear operators saying that the technical community understands the importance of deploying v6, but not the business – the management doesn't.

But it's very clear, you know, if you want to be in business, you want to have more customers, you want the [?] use this mobile, there is no other solutions but to implement the new technologies that we have, and go with the demand that's coming from the customers. There is

one case study that I have included, involving T-Mobile USA IPv6 LTE Story.

Very well planned, very well processed, very well agreed strategy among the management and the technical groups. And this deployment happen. And you can see, it's a successful deployment, so if anyone is thinking about LTE deployments in the region, I believe you can look at this as a case study and learn how to move your customers to mobile.

Conclusions. I think we still have to work together, the public and the private sector, for the good of the Internet, among all of the regions. And we are still at the critical point, as I showed statistics from other areas. All the areas, at one stage, will be mute down to one slash eight, and then based on the local policies, there will be a lot of restrictions in how much you can get before from them.

We still have to work and the deployment, when we're thinking of deployment to v6, I think the strategy here is, you know, what is the best technique that's out there? What is the best approach that other organizations have taken? Learning lessons from them and moving forward with that. We have collected a lot of information, and made it public.

It's on IPv6 website on our site. We have actually focused more on governments, on [?] makers and [?]. We will be passing on the information to those stakeholders who are interested in just knowing what they have to know, not putting everything out there that might confuse them.

And we also have a flow chart, but just answering a few [SN] notes, you know, you can see which deployment strategy, you can buffer it and how soon it should be moving your networks to IPv6. So that's about it from me. Any questions now or later would be fine. Thank you.

MEHMET AKCIN: Any questions for Sunny? So, I'm going to ask and actually talk about the T-Mobile case. T-Mobile is based in the same city as us, in Seattle. So anybody from T-Mobile here? No? Okay. These guys actually we talk like over the beer all of the time. They really went through and reviewed a lot of documentation that's out there, including the best current operational practices, everything you can think of.

And I believe, if I'm not mistaken, there will be [?] event you can remotely participate and see their presentations. Either it was already given, in that case it's available at www.[?] dot org. It explains almost like step by step what did we do kind of guide. It's really amazing. It's impressive work from there.

SUNNY CHENDI: Yeah. They have given this talk in the recent [?] meeting conference in Malaysia and we have the slides up on the website, on the conference website. If anyone wants to have a look at them.

MEHMET AKCIN: [Apricot?] dot net, right? Yes. Any other questions? All right. We will move on to the next presentation.

DAVEY SONG: I want to stand up and give a face to face communication with you all. Because IPv6 is an old question. I don't want the topic [?]. Okay. Firstly, I want to thank the ICANN [?] center to engage with us to here to discuss this problem, transition problems. And Patrick, and [?]… And in this slide, I don't want to go too deeply to the technical detail about the transition tools, the technologies.

I just want to present some statistics, some reference, and my observations about how IPv6 translates, [?] such a problem why its [lust?] for so much time and so long. Okay. I want to begin with a little picture. That first thing, implement to IPv6 to the PSD, and when users want to use the IPv6, he can verify his connection with IPv6 when he saw the swimming turtle.

I once had a chance to speak to the main contributor of the KAME project. Why you design a turtle as a sign for IPv6? This implied [?] a long way to go. Yeah, it's just joking. Next slide. We changed, if we change this as a rabbit, maybe something different will happen. But yeah, it's my first time to talk about the KAME project, the contributor of the KAME project. Next slide.

Okay. I will give you a brief description about the situation about the global developments and efforts just in one slide. Many people may notice that IPv6 implementation has a problem, maybe start from the announcement by IANA in 2011. And followed up by the IPv6 Day, and IPv6 launch.

We can see a jump and increasing traffic from the [?] or from [?] statistics, and we can also notice that the major ISPs and the ICP content providers provide their services, the IPv6. We are glad to see

that. And we also notice that some policy making organizations, such as ICANN, and other Internet community organizations have, give out [?] the video statement to call for a transition to [?]. It's a priority in their task list. I think, next slide.

I think it's good aspect, the bad aspect or challenges here still exists. The lack of global consensus. I think the lack of global consensus on the importance of IPv6 is still exists. The major ISPs and ICPs prefer that all trust other IPv4 workarounds to continue their services. The IPv6 here is optional and too complex IPv6 solution. Later I will introduce how complex it is. The dual stack deployments prove to be the best practice in many best practice stories.

But it still has some problems or some concerns. There are too many transition scenarios and tools that end users, or the decision makers, cannot decide which way they should go. I think, and of course the global, as the hosts said, the development is not equal. In APNIC area, they're exhausted in the pool, but the Internet [grew?] with much [?] demand. They need address space to develop more innovative applications.

So the development is unequal. Next slide please. Okay. IPv6 transition process. I want to show that when I – it's maybe 10 years ago when I first learned IPv6, when I was still a college student. Later I know, it's more than 20 years the technical discussion and the experiments have begun. But from throughout 20 years, there were many discussions, and working groups, and many [?] engineers, and other experiments.

We do a lot of things, but there is still less than expected, the traffic is less expected. We notice that Google allows 3% of the traffic now were IPv6. It's a little small and it's not widely accepted as a mainstream technology so far. We are happy to know to say that more parties joined the discussion, and recent five years, I think, is the [?] of that IPv6 traffic. And so we can predict that…

One prediction said that maybe six yearly, half of Google's traffic will be IPv6. Many people will say yes, it's a matter of time, but I still have a concern here. Next slide. Yeah. There are many transition tools already being talked about. There are translation technology, [?] technology, or to stack ways, and there are many, many [?] surveys about how to translate from four to six.

And I provide a comprehensive introduction you get. Next, yeah. I also observe one discussion of the ITF at universities, that we can change the wording in certain ways, if we put our effort into other part. I guess. The picture shows that just in four over six scenario, there are 14 proposals for ITF standardization. Yeah, I had a chance to join some discussion, the technical discussion of transition mechanism.

I found that there is not a consensus, there is not a way to get the IPv6 deployed. You know, there are more different implementations, different scenarios that are discussed. I think it's too much. Next slide. Okay. According to my experience, when I was a college student in [?] University, we – in [?], to use it the student have come part of, let's say two, 20 gigabytes per month free for each student. And they also have IPv6 configuration. The IPv6 configuration was free.

And there are four, or five, or six application in the campus, locally, and all between the campus. So the student is smart. They can choose IPv6 for what they want. We also notice for our monitoring, the IPv6 export network bandwidth almost fill up between the campus in Beijing. [?] has just introduced that from outside of traffic, the traffic rose not so high as Google indicated.

But I want to say the IPv6 travel, in the channel itself, it's much very expected, I think. Next slide. So I had observed the observation that I think the new technology may be, it's not perfect. It's impressed, people attract new users only when it's shows its difference and benefits. So the IPv6 transition is probably much in the industry, promotion problem more than a technical one.

That's my one observation. So we do a lot of work to, not only on technical specs to promote IPv6, but also from the industry, for the policy makers, as well as ISPs, and also government policy maker to understand that the whole process needs us to join. Okay. Next slide.

Another observation is that the roadmap of the IPv6 transition from the initial problem, is addressed shortly to [?]. The plan and the reality, I think it's – you can already know, the picture shows. There are three ways. One is to implement v6 purely. It's proved to be difficult in the beginning. And another is to share the IPv4 address with multiple users like the IPv6, IPv4 workarounds. So certain ways prove to be effective and now a lot of people invest their energies into that. It's a transition.

The idea is to translate the Internet from v4 to v6 mostly. The idea is good, but the reality is not so exciting. Yes. Well, from the little deeper discussion about the transition roadmap, there are some, most of the

transition is true, I think, are to firstly rely the IPv4, the contents. The users came from the IPv6, use the transition or tunnel technology to get IPv4. You would count it then, the ISPs and the users can update firstly to IPv6 when the content providers [?] increasingly invest their content to IPv6.

Then the traffic will go to the IPv6 world. That's the basic roadmap of what they are doing and what they are expecting. Next slide. Yeah. This is a map. You can, yeah. We [fall?] with a little with six, to [?] six with the [little] four, that's the roadmap, the roadmap. We notice that there are private v6 CGN, NAT44 take much – in this stage, take much note of this transition. So people want to finally, keeping point here. Yeah. The [?] point you can see the two, the three lines interact with each other. The [?] point that IPv6 will, well with no promotion, it will grow automatically with the market beat.

But there are potential trap. The potential trap is that if the dual stack with the CGN and v6, probably IPv4 go on without any direction. Just to free the demand of the Internet nowadays. But IPv6 may stop all just to remain in the experiment stage. That's maybe one potential trap. The APP where one expert, one professor, told me that one day all APP will use only some certain ports, you know, the DNS and the webpage ports.

That's happens now. Next slide please. Okay. Next slide. One observation of my [?] that IPv4 workarounds like the NAT and the CGN, with NAT light technologies, still proceed like the – we all know the Client/Server mode communication, the STUN/TURN ICE stack firmware to support the services of the Internet. And also the PCP protocol, they are all NAT friendly technology.

And also some virtualization and server sharing for the services. And happy eyeballs implementation is also one. You may know that IPv6 once [?], IPv6 has priority over that the IPv4, not now but happy eyeballs changes that balance. The run the two, has the two horse, which one go fast with less RTP will gather winner. So the happy eyeballs implementation drives the user to decide which road should he go. So it's, yeah. The following is the SPDY or we can call It the HTTP 2.

There are some experiments with this technology, with the CGN, and the, yeah. There are some picture from the study. We can say that with the study, NAT session is nearly the [?]. So it's – get reduced the nodes of this [?] so that the quality of CGN, the quality of [?] maybe enhanced in that technology. So this are all the cases that may enhance the IPv6 workarounds. That's my concern about this.

Okay. Next slide. And there are also some pictures which mark 15 points with the latest technology study. Yeah, we can get the full services that are expected. Okay. Next slide. So the development of IPv4 and IPv6 now are in [?], which affect the speed of v6 adoption worldwide. Yes. So what do we want? Do we want a way forward v6, do that forever? I don't think so.

So how we define the success of IPv6, one question is to use, please, yeah. So we can ask when can we pull the plug of IPv4? Are your ready for that to be a reality? Okay. Next slide. Yeah. I give a picture of what happened recently in Beijing, have – has in Beijing. I think the environment will affect everyone. I think it's happened also, the Internet itself is affected by the – not to say that old technology, but some with some patching, the technology. So next slide.

So do we need more sophisticated anti-haze mask? Or do we need another way to have a bright future? So I think in ICANN platform, on this occasion, I want to share with you my concern, my thoughts about why do we change our mind, or to do some more promotion work. Not only technology, but with more awareness of IPv6, the problem, it's reality, and how we can come up with more concerted effort.

I think we want not only one Internet, but also one bright future in the future Internet. Also, I want to introduce some efforts doing IETF, IPv6 only efforts. Some are doing years ago, some are doing recently. I just give you a connection of it. You can download and analyze what they are doing. I think it's a good direction, and have a – to make loud sounds about how we can use a different view, maybe in a pure v6 application or experiment. Maybe a proper view.

Slides. Okay, there it happens recently in the IETF. And I fully agree with [Lee ?] about what IETF should do. It's happening, sounds like that, working group, but not so many people pay attention to that discussion. I pulled these slides here, and want to use this platform to make a loud voice to encourage people to think about the IPv6 only. Okay. Next slide.

There are also discussion, maybe not everyone noticed that IPv6 only root in ITI discussion. There is contribution from Paul Vixie, and he mentioned that we should have… If we have to change, why not consider that v6 only root name service? And separate the domain from different, the connection. I think in the early stage of the DNS ways of IPv6, the may discuss how to support IPv4 and v6

simultaneously, and – but I don't know why they just combine them as just a certification.

I think more IPv6 only work and discussion done in this field. Okay, next slide. Okay. Simply introduce where I am from. I'm from the BII group, located in Beijing. And we're doing research, certification, training and networking mainly focused on IPv6. And every year we host a global IPv6 summit and we invite experts, and anybody who is interested in this field to come.

And this time, we bring some introduction and some unique invitation letters to you as a bonus. My colleague [?], please raise your hand. Okay. We have some material that you can also get access. And what I want to do is, awake the [?] stuff, the IPv6 development especially for some general problem to conquer some problem that we all faced. Next slide.

Okay. Thank you and that concludes my presentation. [Applause]

MEHMET AKCIN: Do we have any questions?

EBERHARD LISSE: I said already earlier, Jack is first.

JACK: Hello, hello. Jack [?]. I'm dot ca. And in Canada, I've been trying to push IPv6 for a couple of years now, and I've been doing presentation on how to migrate, and how to adopt v6, and all of that. And the only

thing we see is a big wall of resistance. Human nature, I think, is we're going to try to adopt, make v4 last as long as we can, we'll do CNG NETT, we'll do every ugly thing out there to make it last.

So the question is, is IPv4 going to break at one time? Like is it going to break, or are we going to be able to sustain that for long? And the other question is, the date is April 4, 2014, so 4/4/2024, that's when we turn off v4.

MEHMET AKCIN: Question for who?

JACK: Is it going to break?

SUNNY CHENDI: I think it's a good question. Personally, I don't think, I believe it's going to break. We still have a lot of those small island states and small economies where growth is rapid, significant. They'll probably still continue to provide their services on v4. We have to look at the deployments that are happening across the region, and see who is actually moving to IPv6 and what is driving them to move to IPv6.

I see it as a business proposition, as I mentioned before. That's actually, in case of T-Mobile, it's the growth of the business. The kinds of standard v4. And other alternatives to move to v6. So that's the business proposition that moves, but it probably will happen in small states as well, where you only see one operator or two operators that

are serving the region, or serving the economy. Very slowly. So I don't know where you heard that date, but personally I don't think…

JACK: [CROSSTALK] 4/4/2024.

EBERHARD LISSE: We must just repeat it often enough, then it will be believed.

JACK: We need to do that. Somebody has got a plan, to state, this is when we're going to turn off IPv4. Another option, observation about IPv6 is the IP eyeballs feature is working too well. Now, in the beginning, it was to avoid bad IPv6 sites from being connected. Now it's avoiding good IPv6 from being used. So I'm trying to compile data on this, but a browser, because of microseconds is going to pick IPv4 instead of v6, even though it's a valid site.

So something… Somebody should look at that. There should be a percent. If it's like 5% slower v6, start using v6 sites.

DAVEY SONG: There is some [?] APNIC, to [?] …about different [?] and different [?]. I think you can… I can give you some information, but you can Google the research. They use a different way of happy eyeballs. Maybe three or four different way to implement it. As far as I know, Google [prune?] used very short came out, just came out to just for [?]… Not simultaneous the file will do that.

SUNNY CHENDI: I just come up the suggestion, and then you will laugh in the back of… So in Canada, the way to adopt IPv6, I would go to the Hockey Federation and tell them that you will only stream the game over IPv6, that's going to be the fastest adaptation ever in the history of IPv6. Just an idea. Nothing is impossible. And for Canada, it's going to work.

DAVEY SONG: First [?] takes on, what Jeffrey is great work. He can see all of those stats on labs, dot APNIC dot net website.

UNIDENTIFIED: [?], Costa Rica. This is not the question, more of an observation. Two observations. About the transition from IPv4 to IPv6, NAT 64 breaks, than DNSSEC. I don't know if you're aware about that. There is an [?] if that breaks, DNSSEC. And the other is more about the promotional of the IPv6. All of the RIRs are promoting the IPv6, basically, [?] IPv4. But I think there is a lot of more benefit of IPv6.

For example, manager of [?] where basically on the DHCP, minus that [?] is great improvement about the management of the scopes and that kind of thing.

Then the [?] in all of this is better in many ways, then it's a good thing to promote IPv6 like a better technology, not like an [?] of IPv4. And thinking about good [?], to say something like the real IT guys use these things. It's a little bit difficult to start, but to prove that you're good with the technology using the IPv6, something like that.

SUNNY CHENDI: That's a good point. The innovation is driving that IPv6. There is a lot of innovation that has been behind it. The end to end connectivity which is driving that IPv6 as well. Yeah, thank you.

MEHMET AKCIN: IPv6 and IPv4, I want to use a small analogy if I may. Think about it like, how many people have a home telephone? How many people have a cell phone? Can you raise your hand if you own a cell phone? Can you please keep your hand up if you also have a home phone? Really? Okay.

I just think it's the same thing, right? It's just like phasing out. IPv4, I don't necessarily see that the translation is going to be the solution longer, but it's just going to phase out eventually than people can reach everything in IPv6, it's just going to disappear itself, I think. Maybe in like 50 years.

UNIDENTIFIED: Thank you. [?] from [?]. My question is for Davey Song. I do appreciate all the work that BII have done in China for IPv6 promotion. IPv6 actually returning the Internet and the industry around. And IPv6 translation is not only for the next few years, but grips the Internet for a long term. My [?] is also recognized from a domain name system, Beijing Engineering Center. The BEC provides gTLD hosting service, [?] service, and high performance DNS device.

Well, there are some recent links, application and addressing infrastructure, [?] from network to user. Also, new gTLD are required to support IPv6, which offers an extra drive for IPv6 promotion. A coordinate, I wonder, are there any IPv6 transition well designed for DNS by BII? Or something on the way in the years to come? Thank you.

DAVEY SONG: Yes. You know, I bring in these slides to share the operation and concern about IPv6 transition. I'm not particular some proposal, technology. I think the DNS industry is a key part of the Internet, the support for IPv6 is the key, I think. The [?] panel, already introduced some discussion to fully use that IPv6 as the mask feature.

What I want to emphasize here is that the Internet is consist with the multi-part, multi-player just as the multistakeholder, and they should proceed accordingly finding a home. So not a lot of network technology or the service technology, or other expect can bring the IPv6 to the real world.

So, that's my really concern. If you want to know more about the DNS IPv6, we can chat later.

MEHMET AKCIN: That is going to be the last question. If you have any more questions, feel free to reach the speakers later on.

DAN YORK: Okay. I'm Dan York with the Internet Society and with the deployed 360 program. First of all, thank you for the presentations. APNIC, thank

you, you're always doing great work in that. I was not familiar with the BII work, so thank you for presenting that. Just listening to one thing that we certainly found that resonates with people as we talked about, you know, getting more IPv6 out there has been the fact that most statistics that we have right now, or most estimates, show that there is maybe two and half billion people online, and that leaves about another five billion to come online.

And that's one of the biggest drivers we've certainly seen for people. They start to look at that and say, oh yeah, and all of those devices. The other point, too, the discussion on case studies. Just last week at the v6 congress in Paris, there were a number of presentations that were given around some very interesting case studies, including one out of Facebook talking about how they were moving their entire network now to being all v6 only internally.

And some amazing statistics and info about what they're doing. So when you talk about innovation, there is companies like that who are already looking at how do they put themselves in a position where they're v6 only, and they're going with that, and they have gateways on the edge to go back to the legacy IPv4 networks.

So some very interesting work that would be there. And I just, to throw it out to people too, we are, myself and I'll point, Chris [Grund-i-men?] in the back there, the two of us are here from the Internet Society, the [?] 360 program, and we have – our goal is to help provide resources to people for making the transition. So if you… I guess I would encourage people to come find us and talk to us about any issues that you're having getting the IPv6, because we're trying to understand what our

people's pain points and how we can help put up information to help take those away.

And I love the idea of the hockey. You're right, in Canada, that's the winning thing. Associate anything with hockey and you're good. Thank you.

MEHMET AKCIN: Thank you. Thanks for the comment, very good comment. And five billion people online and 10 billion [?] online.

EBERHARD LISSE: There is baseball, there is basketball, there is American football, there is soccer, there is cricket, not to forget. And so we'll probably get a few.

MEHMET AKCIN: Well, I think we are about to finish and wrap up here. Just one thing I want to point out. The ISOC guys have this Twitter feed, they announce a lot of activities there. If you're active on Twitter, you should follow them, and they really do some really good stuff. Thank you for both of our speakers to enlighten us, both the China and the APNIC region. And look forward to seeing you next ICANN meeting.

EBERHARD LISSE: Thank you. [Applause] I would also like to think the speakers, but also Mehmet for chairing on such notice. Okay, next presenter is Mike… No, for him it was really short notice. They only came up with this today.

Mike O'Connell from… What are you called these days? Dot Africa these days, okay. Is going to speak about Marks, their policy engine…

MIKE O'CONNELL: Good afternoon everyone. I will try to, I have a habit of… My name is Mike O'Connell. I'm from dot Africa DN services, and we've got a product called Mark Validation System. Next slide. The Mark Validation system is such an entry to the TMCH to cover local industry. It fills a market where the TMCH doesn't, which is unregistered marks, company names, and other custom Marks such as public benefit organizations, charities, and any other companies that may not be registered in the Mark system.

We've based the MVS API on the TMCH version two API. It's word for word, the only alterations we've made are user name, links, and we've added additional Mark types to [?] or [?], Mark spec. We've also added an additional Mark document type, official declaration, which was missing, and that's for your actual Trademark certificate.

The company names, unregistered Marks, XML is all available. I'll give you a website at the end of the presentation. We've got no UDRP support yet for case numbers. We will be adding that in good time. It's built on the policy framework that I've documented in a number of precious ICANNs. The entire API spec took me about two weeks to implement on the policy system, so it was very rapid in getting it out there.

One of the complaints I had was the ICANN TMCH was not built on APP. It was modeled closely, but it was not the same. So it required a

number of alternations on the registry side to actually get it up and running, which was a bit of a nuisance. We got around that. The API documentation is available on the website. Just you go there, you'll get some more information. You can also feel free to login if you're going to be creating Marks and viewing the API from there.

Additionally to the LOUDN, it's a list of registered domain names. We're adding additional functionality called the LOUDN, which is a list of updated domain names. That includes in any chance to the registrant's ID on the system which affects the Mark holder's rights. So if a registrar changes a registrant ID, that notice will go out to NVS and NVS will send it out to the Mark holder. It's just a supplementary run for indefinitely and perpetuity for registries life cycle.

Launch phase specification version 12, which I worked with Will [Ten] and James Gordon [?]. It offers support for all of these functions, but our local industry had issues where they wouldn't have time to implement such a specification on such short notice. They were like tortoises through peanut butter when it comes to implementing things, unfortunately. So we've had to investigate alternate rights verification mechanisms.

And that basically involves going directly from the registry to the registrant or the Mark holder to verify their rights. Bypassing the registrar so the registrar doesn't have to implement the launch phase spec, and that's, you know, last minute stuff, unfortunately. But we need to get 300 registrars onboard, and to do so this is one of the best ways to get through that.

So what's going to happen is, the registrant, the Mark holder will go through the agent to MVS to registrar the Mark. The registry will download the domain name label list from MVS, or TMCH in this case, then a registrant will then go through a registrar that does not support launch phase. Issue an APP create, which will go into a create state.

On the registry side, we'll create an application for them, defaulting to land rush phase. We'll also send an email out to the registrant asking for any additional documentation to support this application, such as SMD. The registrant will then go to the agent to get the SMD or the validation token from the MVS system, and upload that via URL to the registry directly.

At that point, we'll escalate the land rush to a sunrise application, and they'll go through standard validation and eventually, hopefully no contention to registration. The same with claims period. We might or might not run this indefinitely, we're uncertain at the moment. But the objective here is again, to not – or just allow a direct access to our system without the support of launch phase during the first few months of GA.

Most of it is fairly soft [?]. We'll get the Marks into Mark Validation System. The email is then sent to the registrant with the Mark holder information. And the registrant will then say yes or no [?], but I don't. And very straightforward, we bypass the registrars in this model again, just because it's a hassle on their side.

Currently on NVS, we're relaunching web za in just under a weeks' time. Web za is going to be supporting NVS and this additional functionality. We do have about 140 registrars on the za network that will support

launch phase, but we still have about 300 and something that don't. So with that in mind, the alternative rights verification is quite important. Just to getting them onboard.

Because we want to use this as a test bench before we launch Dot Africa, just to get the systems up and running and make sure that the bugs are ironed out before launch. Org za is also re-launching probably fourth quarter 2014, depending on the bureaucracy.

The gTLD support for NVS and the API and all of the functionality is listed there. We're going to see how it goes on the local industry. We again, we can't dis-include Cape Town, Joburg, and Durban from our local registrars, and then requiring an ICANN registrar to register. So they all have to find a local ICANN registrar to come through, which is not necessarily at launch phase, so the alternative is validation. That is very important.

We're going to be adding additional Mark types for public benefit organizations, charities. And effectively, MVS is available for any rights moderation. I think dot horse, we're a perfect example for this. You know, you have a registered horse and you have a system up and running that you can moderate the domain, NVS validation tokens is a good option.

I'm otherwise looking forward to launch. That's the end of the presentation. I've spoken too fast again, unfortunately. Five minutes short. Yeah, but big thanks to James [?] and [?] for modifications to the launch phase spec. We added additional attributes for validating to IDs, just to support MVS as well as TMCH on our gTLD level. So thanks guys. I really appreciate that.

And we are having a public Dot Africa signing ceremony on Wednesday in the Canning Room at 5 or 6, around there. Yeah. So please be there to witness a historical moment. It's been a decade of hard work, and we finally got there. Thanks.

EBERHARD LISSE: Okay. Thank you very much. One thing that I didn't really understand, so this is not registered Marks. How do you solve contention?

MIKE O'CONNEL: Each Mark type has a priority level. The information is available on African one space dot org. The Trademarks, local Trademarks, registered Trademarks will have the highest priority, alongside court Marks, and treaty, and statute Marks. Then international Trademarks, and then we'll go into unregistered international, unregistered local Trademarks, or company Marks, deprioritizing as we go.

The policy engine will take care of the prioritizing and sorting of that into contention sets on the highest priority only.

EBERHARD LISSE: Any questions? All right. Thank you very much. Next presenter will be Stephen Deerhake.

STEVEN DEERHAKE: Hi. My name is Steven Deerhake from AS Domain Registry. This is a version two talk. I was able to do some expansion on this. I originally presented it at the APTLD meeting in Kuala Lumpur last month. And it's

based on some interesting things that I discovered looking at daily emails that I get from NeuStar Ultra DNS, who provide our secondary DNS servers.

A couple of acknowledgements. One for Don Hollander for suggesting the title of this talk, and for permitting me to present it at the APTLD meeting. And another acknowledgement for the gentleman to my left who looked at some work I did early on in this, and suggested some improvements, and then took my work forward into his own registry and found some equally interesting things, which he will talk about shortly himself.

Shameless plug for American Samoa at this point. It's unique among the US populated territorial possessions because it's both unincorporated and unorganized, and as a result of that legal distinction, which is unique to the US territories, it has a higher degree of autonomy than that enjoyed by the other US populated territories such as Puerto Rico and Guam.

It is also the only inhabited territory in the United States below the equator. There is one other, it's unpopulated, and there is extra credit if you can name it. The registry was established back in 1997, as such, it predates ICANN. It was delegated directly by John [Pastel]. Currently has an excess of 17,000 domains. It has a high entry price point of USD 100 for a minimum two year registration. And this tends to keep miscreants and short term registrations out of the thing.

We offer few registrations, it's a policy standpoint for anyone on island, in American Samoa. And we also offer free renewals. We do not charge for transfers between registrars, nor do we charge for modifications.

Most of the registrations can be broken down, it's pretty much local on island businesses, although some individuals have been adopting them of late. It has got a fair amount of brand protection, which is fairly typical for a cc of this size. And our niche play is in Scandinavia, particularly in Norway and Denmark, where the designation a slash s stands for joint stock company, and hence dot as looks nice on corporate CEO's business cards.

So we're a pretty dull registry, all things considered. I never thought I would consider myself, we'd be a candidate for short term registration simply because of the price point. We've been asked repeatedly to adjust that price point, we're very reluctant to do so because we feel that it's part of our brand at this point to reflect and maintain the quality of the registrations.

And as I said before, it's mostly brand protection, on island businesses, and Norwegian/Danish corporations. Our DNS setup has evolved over time. At one point, we were like most other small registries, we had using e-net at one point, and a couple of universities at one point, etc., etc. You know, 12 plus years ago. The current situation is we have seven distinct entries in the root. One entry is under the direct control of the registry. This is a non-NE cast server.

The other seven entries are ultra DNS, NeuStar entries, and an endpoint to their various NE cast networks. This graph shows a typical month of DNS queries, and as you can see, day by day, it doesn't vary a whole lot. And the actual spike we had in the month of December, when this was taken, is around 15 plus million queries.

So this is fairly typical. It doesn't vary a lot. Every day I get an email from NeuStar. This is from the NeuStar system, this is not from our single DNS that we control. I get an email from them listing how many queries per hour and so on, with a total at the bottom and it varies around this and so I don't pay much attention to it. But then I started noticing a spike in January, so I threw this stuff into Excel and that's the stuff I got.

And at the peak of that spike in queries, it went from a daily average of 116 million to about 140 million, so it's quite an increase in query of volume. This is dated against the Ultra DNS servers. I didn't really pay much attention to this until a few days later, and I got interested about it, and said, "Well, where is all of this coming from? And what impact on the registry did we have?"

And the answer is obviously, the bulk of the traffic was handled by Ultra DNS, we did not have any operational impact on the registry. We continue to serve DNS just like we're supposed to. But we, I did not have any data from Ultra DNS. I do not get log file information from them, I just get this summary email mentioned earlier. So what I ended up doing was pulling the raw log files from the locally controlled name server, and then I dropped them and then the SQL is using post Postgres rather low end, ancient, ancient 386 boxes of some sort, X86, with [Sentos?] and Postgres eight, something on it.

Not very fast hardware at all. And then I started applying a little SQL against it, and since this is a tech talk, I thought I should show a little SQL. So we have a little code up here. What I was doing here was going, well, this is probably coming from one or several but not many IP

addresses, and the intent of this code was to try to drill down and find that out.

And as you can see from this, there is one pretty clear source of all of these queries. Again, this is just against the one machine I had logged data for, but 36 million versus the next guy at 610 is pretty clear, this guy is the guy to pay attention to. Did a little further work on him going through WHOIS at RIPE, etc., etc. Traced back some address in Dresden, it's controlled by a German ISP.

Initially I thought, well, is this guy a victim? Was he just getting hammered by somebody attacking him? Or was he an attacker attacking us? And I quickly concluded he was attacking us. Eberhard will go into detail on this particular miscreant because he has some further information on that in his talk coming up. I started drilling down further into my log files, which is fairly easy to do once they're in a SQL table. There is no magic SQL involved, it was just increasing the granularity and poking around, limiting it to the address, the IP address, etc., etc.

And there were three distinct phases that were evident from the entries in the log file. The first was a code development phase. I never heard from this guy until like around the 12th of January, never saw back in through December any queries from this IP address to us. And yet, in looking at the log beginning where we started seeing these addresses, you can see that there would be several hundred entries and then it would go quiet for five or six minutes and then it looks like another code test was run, and they ask us 500 more queries quickly, and then go away again for a while and so on and so forth.

And the code development began on the sixth and it took them a few days to get it sorted out. And then they went into a code testing phase. And when they went into the code testing phase, they switched from asking from A records to NS records, which I thought was a little interesting. I didn't really notice this at the time until I did the whole summary of the attack. And they tested the code. They would run 5,000 and go away from a while, and then run another five or 10,000 and go away for a while.

Then they went quiet for a day, and then they did the full on attack, which started on the 14th, lasted almost two days. They were just pounding away at our name server, and presumably against NeuStar Ultra DNS as well, on their [?] network. And it just stopped, and I have not seen a single query from this IP address since. From an examining the log files, it's pretty apparent it was a dictionary attack. And it had English and non-English strings. Had numeric and alphanumeric strings as well.

The usual patterns of zero dash zero dash and variance like that. Could not establish any firm evidence of many repeated queries. There are some, but not very many all things considered, the volume. How successful was the attack? Well, from a NS record harvesting perspective, it was reasonably successful. They managed to take from this one non [?] name server, more than 60% of the zone file, based on matches I did between the query names in the log file, and our actual zone file also imported in the Postgres for the purposes of a comparison.

It was a very handy tool to use that. From an efficiency standpoint, in other words, how many queries did they have to run before they found something that actually was in the zone, it wasn't very inefficient. They ran over 36 million queries over the height of the attack, and came away with about 11,000 entries from the zone file. Consequences. As I mentioned earlier, we lost over 60% of the zone file, but in reality, one has to assume that based on their volume of query against the [alter] DNS NeuStar servers, and the fact that there seems to be very little repeat in their dictionary, that their dictionary was probably of sufficient size to essentially make off with 90 plus if not indeed, 99% of the registry zone file as of mid-January.

On the good news front, there was no WHOIS data taken. This was not an attack against the WHOIS. It was simply an attack against the DNS, exploiting what we do as registries. We get asked a question, does this string exist? Or what are its name servers? We answer that question, and we just happened to answer it a lot of times over a very short period of time.

We did not have any operational issues affecting the registry. The DNS continued to operate just fine. The website and registration services operated just fine. And the WHOIS operated just fine. Ongoing. I'm looking at impact on other registries. One of the things I did when I started doing this analysis is I sent this information down to Eberhard, and he came back with some very helpful suggestions for Perl scripts improvement and so on and so forth.

And he launched his own little episode, which he will tell you about. Yes, in his usual charming manner, we must not forget that. When this

slide was prepared, I was debating whether or not to actually discuss this with law enforcement, since it was prepared I have opened that dialogue with the FBI, and they're pretty interested in pursuing this further, so we're going to have some developments down there at some point.

In terms of limiting this going forward, obviously it makes some sense to try to implement response rate limiting or something similar, just so we can identify addresses that are suddenly sending us a lot of stuff and see if we can't slow them down a bit on that. I would also like to look at the data held by NeuStar Ultra DNS, open that dialogue with them. The last time I spoke with them, not to bash them, but it was well over a year ago, and the sales rep talked to me about selling me my log data, was younger than the age of my contract with Ultra Star DNS, and that didn't go very far.

And that's, then looking at, you know, doing some real time monitoring, I'm discussing with Eberhard and some other people some possible design ideas for building tools that will spot spikes and queries, and perhaps allow us to do mitigation in near real time. And that's it. I'm working on documenting on a lot of this in terms of the scripts, the table layouts, and some of the analysis in a paper. If anybody is interested when it's done, just drop me an email and I'll shoot it off to you guys.

EBERHARD LISSE: I initially wanted to first do mine, but you can of course, since you asked your question.

UNIDENTIFIED: I just wanted to give some extra credit here, Jarvis Island.

STEVEN DEERHAKE: You have won the grand prize.

UNIDENTIFIED: Thank you. I actually do have a set of questions, but I'll wait until Eberhard is done with his presentation, because I think your presentation will be about a similar subject, is that correct?

EBERHARD LISSE: Yes.

UNIDENTIFIED: Okay. Then I'll wait with my questions. Thank you.

EBERHARD LISSE: So, I will talk about the same thing that happened on dot NA, but more sort of, not so much on what did they do about it, but how did they do it. Okay? As I'm not able of any original thought, I basically took Steven's software and had a look at it. And it's extremely well written and documented so even I could understand it, it's written in Perl.

So I basically brushed it up a little bit to make it so it's readable and it's working very nicely now. I [?] and Steven told me he wanted to present this. I, of course, thought, wait a minute, let's have a look because I was bored. As you know, I had a day job and until recently, also Nigel, but because of the high cost of mud practice insurance, and maybe I've

decided to stop doing deliveries, and I get now – now that I'm not called every night anymore, so I usually stay up late, not playing with myself, but playing with my computer.

And I was thinking, what can I do about this? So, we looked, we didn't run any log files. So I first of all turned the log files on, and because I could turn on any log I could, and he said, "Jesus, what are you looking?" I look at everything now. Every log channel that bind this is not turned on.

I took his Perl script, and basically we plugged it into the MY SQL database, and we come to how that came about just now. And then we looked, I looked up with WHOIS, what CIDR class. It's not just an IP address, but what net block it belongs to, and then I'll try to look, are we having queries from the same net block? Is he distributing more than one. And like in Steven's case, I'm quite sure it was a mistake.

They didn't figure out what was going on and they turned it on too high. And I come to the reason why I say this just now. We then wanted to contact a secondary, so we then wanted to do some analysis on the secondaries to find out whether they're seeing this too. PCHC said yes, but we don't bother, we don't even query by net group, we only query by countries. It's just too much effort, our pipes are too big.

The Germans said, yes we see two and we know the host, we've had occasional issues with them. If you can't [?] mentioned a word access control list to them. [?] noticed this, dynamic noticed this, dynamic DNS has some new stuff, very competent people but they obviously haven't heard of me, so they said this is confidential, we can't give you your information.

Then I had to politely inquire whether it was, or too much trouble to stay in the root as the dot NA server, and then I got very polite messages that they have now been fully briefed. It's a matter of growth, probably a good thing. And then I wanted to contact the perpetrator and he's very, my usual charming bedside manner comes into place.

What Steven has forgotten to mention is that under the… Let's put it like this. I recently read about somebody being locked up, some hacker for four and a half years under the Computer Abuse and Fraud Act, a federal pen, four and a half years. So I looked up what it says, and it says that a government computer, a banking computer, or any computer that is used for interstate commerce. There you go Steven. That is being attacked or being accessed with an unauthorized purpose, is a Federal offence. In the same act, it also says you can be liable for all of your – you can hold the perpetrator liable for all of your expenses you incurred to investigate and to rectify this matter.

Steven fortunately didn't have a volume based chart, but he mentioned, if your prices goes within one day by 10 times, it can throw you out of business. This is an example for all of us that you people know. In any case, when I mentioned this politely to Steven, he probably started, he must figure this out whether we shouldn't open an investigation. I think we should not discuss [door down?], and we'll come to that later.

What software did we use? This is my secondary is an old Ubuntu 8.04 box. I know, I know, I know, it's not supported anymore, but it's been broken so why must I go and fix it? We have resurrected the next box in the wreck, and we figured out that for some other reason, it wasn't

running for three years because nobody went in there and pushed the button after the power had failed.

So we went and pushed the button and it came right back to life, so we put 8.04 on. We were experimenting with it and then we brought that name server up to 12.4 as well. We run My SQL 5.5 version 16. 5.6 is not that latest, but 16 is the latest version of 5.6. You need to run 5.6 for some certain queries, field types, that you note. Perl is the latest version of both Ubuntu and on the Mac.

As far as the analysis software, I used 3.2 but 10 days ago 3.3 came, so we loaded that, we are now current. I like to write my reports and my presentation on this one, LyX, which is a frontend [?] and because I like to toy with things, I'm using a beta version. It only crashes about once every three days. I've learned to use the backup button. Beamer is the module that allows you to write presentation.

The cool thing about it, you can, if you're really play with it, I haven't been that far with it, you can put your presentation and the report, or lack of dissertation, into one document and when you push a button it comes out as a presentation and when you push the other button, it comes out as a report. And if you then plug this other module, Knitr, it's a play on something called [?] which is the old version, you can put tables, or rather the source code for a table, and graphics, or rather the grammar that describes these graphics, in your report, push a button and the current data comes up.

This presentation was written on the 19th because I was being – I wanted to lead the charge to sending the presentations to Christina early, and I don't think I wanted to change my presentation after I send

it in, because it would be two different presentations on the Adobe Connect, and on the one that I do here.

But this morning, when we had coffee, I ran my confidential report that I wrote for my colleagues at my company about this, and we just loaded the data, ran it through, and it's now 16 billion record entries. Off the shelf hardware. An old Pentium, and old Pentium dual core. My iMAC is sort of loaded, it's one and a half years old, but fully loaded. I'm probably going to invest, next year, into this new power MAC just because I can.

But it also makes analysis a little bit faster. Okay, now some results. This is from the 19th. We had 12 million 900,000 queries. With a mean of 390,000 a day. And a maximum of 344 per second. If you were really to analyze the mean, really per second, by dividing the number of queries by days and then seconds, you would come up to 5.5. I analyzed the queries not on microsecond, which is what bind does, but by second so I can easily group per second. And then at the most we had 344 and the mean was 4.7.

Now as far as the tool is concerned, in the presentation, in the original presentation that I typed in on my screen, it doesn't have these numbers, it has place holders. Whenever I save it as a PDF it runs the query through to the database and puts the mean in there. And I find that quite cool, especially if you have recurring reports, or if you want to keep a presentation up to date.

Now, this is a smoothed scatterplot, which shows the black dots which you don't see very well, and we don't really want you to see them very well, because that's not the emphasis that I want to make. If you

smooth there, you see the average approach in the light green, lime, and you see around, before the 24th there is a drop, which is paralleled in the red line which is the spam queries.

This is a very – I find this a cool looking graph. It took about a week to figure out how to do this in the language, I have since bought the book, with a little bit of help from an Usenet group, we figured out how to do this. And the cool thing about it is whenever I want to update it from 14 to 16 or to 20 million queries, I just need to push a button and I don't need to go to Excel and do things.

Okay. Now, post SQL has a very cool method of grouping into CIDR blocks. You can just use the quick, smaller sign times two and then you can compare where an IP address fits into a CIDR block. SQL Lite has a similar message. You must write a library, or get a library, compile it, load it, and then it can do it too. My SQL can't. So I had to develop a view and manually put this into a view, but now I can group by this.

And we find that we have a few blocks that are interesting. The 74 125 zero zero is not interesting because that's a legitimate block. That's Google. Google is not really sending us too many spams, but they are being – as an open name server, of course, they get their fair share of queries. We contacted them to figure one particular one out, but most of these queries are legitimate.

144 76, 176.9, these are all illegitimate ones. There is one, a block of queries coming from Colorado. These people have not responded to us. I'm quite sure I will convince them to respond to us. I will now, now that I have ready computer fraud act, I will convince them to talk to us

or they will talk to somebody from the FBI, because it is fraud in what they're doing.

Then there is somebody from Romania. I think I will probably waste my time talking to them. The software that I have put here in links. So if you download your presentation, you click on each link, you get to the website. So now what did we learn? The guy that attacked us is the same guy that attacked him. It's one Mr. [?] from [Daupe?] DE. Is he present perhaps? I invited him, if he was in Singapore, to join us here because I told him I would [?] this.

Somebody from [heads] is here? I invited the host [?] if they are also in Singapore to be present, because then we could discuss it and they could give us their side of the story. They have a very well-functioning ticket system, and very [insolent?] staff members. You can have 10 IP groups. You cannot write one question, you must write 10, even if they belong to the same person, they want you to write 10.

So eventually, I figured out how to automate, I think, 10 complaints to them. And probably this is the idea to not make it too easy to complain, so that people start basically, I'm not doing it. Now, German [?] just don't like the words [German]. So that means copyright violation. That's a criminal offence. [German] means public prosecutor. It is a peculiar type of the German law that there is two types of criminal offences, a public one and a private one.

If I insult Steven, it's a private offence, and he would have to lay a criminal charge against me at the police, and if he was [?] it was done. If I was insulting a police officer, that's a public offence. If a public prosecutor was standing there listening to it, he would have, by law,

press prosecution against me. I think [?] is a public crime. So I'm going, I told him I would do this, and I will now, on Monday, when I am in Germany, go and familiarize the public prosecutor of the county of [?] with the terms of this situation, and see what's going to happen.

Whether he would like to knock on this guy's door. What German [?] don't like to hear the word [German]. If you are a host under German law, and you're guy does crooked things, you're in it too. It's called something like, [German] means aiding and abiding. The American Computer Fraud and Abuse Act has clear cut provisions for doing this for commercial purposes, for doing it together. It's not really RICO, but it's – they will be held responsible if they [?].

And then they really don't like access control lists. If I were to ask all of my primaries or my secondaries, big [German]… If I were to ask them all to put them into access control lists, [?] would probably not be able to resolve half of the ccTLDs anymore. So, when I use that word, they actually responded. And here you can see, you can actually spot three lessons, but I only listed two, I show two. One is that on the 19th, on the 20th of February, at 17:00, I received an email that he didn't do anything wrong.

No, no, no, no. But I noticed the queries stopped. Okay? I don't know whether you can see my, here it is. This is the 20th, this is when the queries stopped. Here is the pure, the identified spam request. There may be some in here that we haven't identified yet. Also you can see it stopped.

You can see that it's a basic amount of background queries that we all have to encounter, botnets and others, but there is two other things

that is important. One is that I can see every week, every weekend rather, we have a drop in spikes. Because we are such a small TLD, we see this here better on a gazillion queries, it probably doesn't matter so much, but it means that 10 to 15% of my DNS queries are purely business related. Hence the digital economy of my country, which is not very big, depends on some extent on the proper functioning of my name servers of my company.

So I have a bit of a responsibility to go after whoever comes after me, never mind that it pisses me off to no belief, that somebody is using my work, and my effort, to make financial benefit. Namely using it as an intelligence tool to go and tell people, oh, in dot AS you can register this and this name. In dot NA you can register this and this. Oh, and dot NZ, you can still register this and this.

If somebody wants to use my intellectual property, or Steven's, or Jay's, or whoever's, the should come to us, enter into a contract, and pay us a price, a license fee. That's what the law says in every country. [?] so that somebody decides, no, just because I do use two, three credits a second, I can do that. Now they can't. And whether Germany, whatever, extradite a German to America for something like this, is highly unlikely because the Constitution forbids it, but it will restrict this guy's travel a little bit.

And, the Germans prosecute Germans for criminal offenses for which they don't extradite. So he will actually have to go and figure out whether this was legal. At the very least, he will have to pay for his defense attorney. And this is not something, malicious prosecution, this is something where I will go to the prosecutor, he will go to his

prosecutor, and we will see whether this guy will not stop it. And if he can stop one, he can figure out methods to stop others.

And if you look what he did to two, he did it 240 others. The Germans, in New Zealand as we say, they have big pipes, they don't really bother, but they're still doing this. It's wrong. If we can develop methods to figure this out, we come a little bit further, we can actually also look at botnets. Identify them.

The third thing, which is very, it's probably not so easy to do, is at least one spammer who works business hours, who reduces his spams over a weekend. You can see this here on the red line, that's also regular little down take here. I found that quite interesting. What we see is, queries like this. It's [?] and [?] here because that's on the name server, zero minus zero minus five, zero minus zero minus eight, zero minus…

So it's a systematic five character string, they go from zero to z, and permutated this. And it comes, 74 is Google, the other is disguised from [?], and I'm going to have a word with, and the other ones are from Romania, who probably I'm wasting my breath even thinking about them, because they will not respond. This seems to be not a botnet from them. I've asked them, they just don't answer. So I don't know whether it's them, whether they attack, or whether they are…

So now, the interesting thing is what I did with this when Steven sent me this. Postgres can read about 1,000 insert statements a second. I was thinking, what happens if a name server gives you more than 1,000 queries per second, and you can't load them into Postgres anymore? So I figured out that Postgres has a copy statement, where you can bulk load, so which means I could increase the loading to 5,000 per second.

And then I was thinking, My SQL can do about 4,000 insert statements, but it can do 25,000 by copy statements per second. SQL Lite can do about 30,000 bulk load, but then the analysis is very slow. And My SQL gives you the fastest loading time, and also the analysis time is reasonable.

My thing would be to sort of get BIND to write not to a text file, but what is called a Named Pipe on Unix. It's a special file in which one program can write and other programs can read from. So we must ask developer [?] whenever I stand by this program that is being started, and then it will, in real time, read all the queries. BIND has got a hook for My SQL, or for SQL, but not for log file writing, only for configuration.

I don't know whether anybody has thought about it. This is where, shortly spoken with Roy about it. Now, we must sort of look at what queries… How do we identify, not by volume, but by spikes in the queries, if they get under the threshold, we won't even see it, or we only became party to this because they made a mistake and increased their queries by 10 fold. I would like to find out a way of identifying queries that are not part of my data maze.

If I take my plain three and a half thousand names, and look which one of those are not represented by queries, and leave some other stuff out, that thing runs about five to six hours. So, it's a pure SQL statement, it's probably not going to work, especially when you not when you have a bigger database, even – for me it's not time critical. If I let it run at night, it doesn't matter.

We probably need to find some learning software that looks at queries and find that these queries do not belong to us, but this is something, I've seen this before, and then developed sort of an interest. And the other thing which is really irritating to me, that if I find an IP address, I have to manually go and look at the WHOIS, and then there is different WHOIS versions around. You can't even automate that, you actually have to look at this.

I would like to have a database where I can query in real time, send the IP address, and get back, this belongs to the lowest network, not the slash eight, but the slash 23, slash 22. So I can then look all addresses up that fall into this network and see what I'm having there. It's not too difficult to develop a view. My SQL [?] addresses in [?], but it's not really that much problem when you have got the network, when you know it's a slash 24, to calculate which, the [?] of these two addresses.

And you can do this in automated manner, so you don't hear, I just have a function so you don't have to do this manually. That's the thing that I have learned. Any questions? Now you can come more.

RAY ADAMS: My name is Ray Adams. I work for [?]. I've got a few questions. You guys, both of you, keep on talking about the illegality of this. Yes, [?] this is an ethical issue to do this. But the sole purpose of having a name server is to serve queries, right? They obviously innumerate the entire zone file, 60%. What people are doing in general, what we're seeing for instance, take the zone file, which is free available.

You can get it, they [strip com] [?] dot UK, and boom, they have our entire zone file. There is nothing we can do. Another example…

EBERHARD LISSE: Let me answer this. I'm not an US… I think also under the German Computer Fraud Act, but I haven't really looked that up yet, use of a computer for a purpose that is not intended is a criminal offence. And if we do not wish to hand our zone file over, but we wish to make this available so people can access websites [CROSSTALK]…

RAY ADAMS: …not lawyers, but the thing is the sole purpose of that name server, the intent of that name service is to answer queries.

EBERHARD LISSE: It's not to answer queries. It's to resolve names so you can access websites. This is not to make the intellectual work that you can I and we have put in there, freely available to every camel so that he can then sell it to other peoples. That's the illegality about it. I'm not a lawyer, but I'm going to take it up with the local prosecutor and see whether I can push him into doing something. I will report back on the outcome.

RAY ADAMS: Thank you. I would love to see the result of that. Just a hint, and I understand that this is being recorded and people might be listening, etc., etc. Can I just urge both of you, at least one of you, not to deploy a zone file with DNSSEC? Use NSEC 3? If you sign your zone with NSEC…

EBERHARD LISSE: I've got NSEC 3.

RAY ADAMS: You have?

EBERHARD LISSE: Yes.

RAY ADAMS: Okay.

STEVEN DEERHAKE: When ours is signed, it will be NSEC 3.

EBERHARD LISSE: Or rather the zone that is commercially valuable, the dot com, that is signed with NSEC 3. The root is only a few number names, it's not a commercial issue.

RAY ADAMS: Okay. That's it. Thank you.

EBERHARD LISSE: We looked at that. We only started signing… We have got the split model, the second level has got a few hundred names and com dot NA has got two and a half thousand names, and that one is [?] the

important zone is that one. And that is done with NSEC 3. It's not as easy as I thought it would be to change running zone from NSEC to NSEC 3.

So, I have deferred to Jay, when he wants that, I should leave my fingers off even looking at this. But it's commercially not an issue because it is small – the way our model is. But we looked at that, and it's very important that if you have a large zone, you must use NSEC 3 if your – because then you make it even easier for somebody like this, you can just work it.

RAY ADAMS: If there are no other questions, I can go on, for instance, about the data detection stuff. The guys who do this very, very cleverly. They are – they do this in such a way that it is very hard to detect. They use various different sets of IP addresses. They know it exists. The moment they know it exists, that won't query for that anymore. Or, they basically keep on querying for those that do exist, so it basically intermingles with the stuff that doesn't exist, etc., etc.

They can do this in such a slow rate that it's hardly detectable. They can harvest other zone files that are freely available. They can check through Google, for instance, to see – they do that almost automatically. You can use Google search, use Google API to basically do sites called NA, and see – or com dot NA, or com dot UK, and that way it gets the entire zone file.

Then there is a lot of companies that do passive DNS. It not only uses security tools, also uses data mine tool. So it's not just the queries that

you see, that release harvesting, or leaking your entire zone, or your entire zone being stolen, to use your terminology. There are other techniques as well that doesn't include name servers. Just, that's all.

STEVEN DEERHAKE: This was pretty evident though. I mean, the dictionary was alphabetized they used against us, and it all came from a single IP address, not even a slash 24.

RAY ADAMS: Yes, just because they're malicious doesn't make them smart.

EBERHARD LISSE: It also doesn't make them stupid, yeah? While I get away with prosecuting them or not, but I'm not going to lie back and enjoy it. Yeah? This is not so much what I'm going to do, this is sort of something that I wanted to bring to everybody's attention, even smaller zones, smaller ccTLDs, even in small developing countries. There is some form of digital economy, whether it's like on American Samoa. And maybe it's small, but if your DNS doesn't work, banks don't work, several car manufacturers or car sales cannot order spares anymore.

We had an incident with Volvo once, where we figured that one out. It's important that one realizes that this intellectual exercise that we're doing has a real commercial value, and real problems can occur, which have real financial implications. And of course, people who lose money on us, they will go look after, will go and find out who is responsible,

and they will, of course, try to get to someone who is located close to them, which means in country.

And really, I'm going to try whatever I can to sort of see whether I can make an example out of this guy because if I can do it once, we can all do it again.

STEVEN DEERHAKE: I'd like to follow on what Eberhard said with regards to the economic costs. There are, in fact, very real economic costs potentially involved in an attack like this. I'm sure most of you budget your DNS connectivity cost in some coherent, consistent manner. I certainly do. My contract is actually zone size based, not bandwidth used based.

Consequently, in this particular instance, with regards to NeuStar Ultra DNS, there is no economic impact on the registry, however, we maintain a private zone file on island and had they discovered this zone file and attacked that zone file, this would have cost me probably easily mid five figures. Had it happen several years ago before the advent of the cable on island, the satellite usage costs would have been well into six figures.

So, there are real economic costs. It is a conundrum in the sense that yes, this is what we do. We run servers that answer the question, does the name exist? Does this name record exist? Does this address exist, etc. We are obligated to answer those, but I don't think we're obligated to take what, in my instance here, was rather simplistic, but systematic abuse of the obligation that we have as registries to answer these questions.

And had my contractual arrangements been structured differently, or had it been earlier in time against a different name server, that they may have discovered by hook or by crook, it certainly would have bankrupted the registry overnight. So there are real economic costs to this stuff. You guys should be keeping an eye out for it, and you should be trying to mitigate it and take what action you can against it.

EBERHARD LISSE: All right. Thank you very much Steven for bringing this to our attention. And the next speaker is Paul Ebersman. He will go through systematic workup on DNS and DNS attacks.

PAUL EBERSMAN: All right. So, as you've probably all seen over the last couple of years, DNS has become much more popular, not only service but vector of attack, as well as one of the more popular targets. I'm going to sort of go through all of the slides, both the authoritative and the recursive. I realize most of you are probably more on the authoritative side, at least externally, but probably run recursive internally, and certainly we all have customers who are concerned about these issues.

So there are two kinds of DNS servers that exist. There are the authoritative servers and the recursive servers. The interesting thing for attacking a recursive server is that if I can as a bad guy get data into your cache, that has the information I want, you as the zone owner are no longer control your information. So, I won't go through in detail how recursion works. The key here is that you start off and you step down.

From the root, you go down through each of the iterations until you finally get to the point of an answer. And what is usually used by most people is a recursive cache, much in the same way that your computer has a stub cache that goes off and talks to your ISP, the IPS's server actually does all of the work, then it stores that information so if somebody else asks the same question in a short period of time, it doesn't have to redo all of that work.

Usually that's considered a feature in its time savings, however if what I can do, as the bad guy, get you to put the wrong information into that cache, I can create a whole series of problems for you. Unfortunately, DNS was designed back in the 80's, and at that point in time there were no ACLs on anything, routers, posts. They still, in most cases, we still ran our finger server that let you see who was logged into your machine, for anyone on the Internet with no restrictions at the period of time this was going on.

So, the whole concept of bad guys was just not something that was designed at all. So there are some flaws in name server implementations, which have, for the most part, been fixed, but there are fundamental flaws in the actual DNS protocol design that we did not foresee at that point, that have made for some interesting times for us of late. The other issue, which I'll talk about later on, is the fact that there are a great number of recursive servers that are open, or will do recursive look ups for anyone on the Internet with no restrictions.

So, if I manage to somehow poison your cache, what does that really mean? Well, it means that essentially, I can do all sorts of nasty things. I don't have to break into your web server, I can simply send you to my

web server that I happen to have stolen the images off of your web server, make it look the same as your bank, your business. I can totally clop all email, send it on to you, but look at it first.

Pretty much anything that a man in the middle attack, or just sheer denial of service can be done by pointing you to a different machine, or by breaking the correct DNS. Now the first instance we saw of this was seen by Eugene Kashpureff, and it actually required you to have a DNS server that broke the rules. What you did was, you deliberately went to a cache and made a query that you knew that you had crafted a special answer for.

And then when you resolve or gave that answer back, you actually give back additional information, sort of quick aside. Like glue records are additional data that I as a name server might give you when I give you a referral. If I am the com server and you want to go to example dot com, and my name servers are NS one dot example dot com and NS two dot example dot com, you can't get to me because you don't know what the A records are because they are in the zone you have to ask the name server for.

So in order to get past that chicken and egg problem of how do you have name servers that are within the zone that you're querying for, what we did was we came up with glue records, which are essentially useful extra hints, additional extra information that you might find useful. So when you come to the com server and say, you know, example dot com, or dub, dub, dub example dot com, I would actually give you A records for NS one dot example dot com and NS two dot example dot com.

Pretty much necessary for that problem. Unfortunately at the time, we also didn't do checks on glue records, so if I came to you and gave you the name server that you asked, the authoritative answer that you were looking for plus additional records, you would just happily stick that in your cache. There was no validation. So I could put in stuff that was unrelated to my zone completely, as in the example here, where even though you are asking for alternate dot net, I would put in Internet dot net stuff.

And you would happily take that information, and you would stick it into your records, and you are now poisoned because you now have a record for dub, dub, dub dot Internet dot net. You won't actually ask the authoritative server for it. One of the other problems is that message IDs are a way that you know that this is the answer to the question that you asked.

So when I as a recursive server want to go off and find out the information for dub, dub, dub dot Google dot com, I will eventually get to Google's server and I will send out a query with the message ID, Google will send back the reply with that same message ID, and that's how I know that's the answer to that particular question. It's only 16 bits, which isn't a whole lot of data on today's market.

And we discovered that there were implementations that had flaws where it wasn't even that random. In particular, at one point in time, BIND random number generator was such that, if the message ID was even, you had only 10 possible next values, which made it trivial to guess what that message ID was. The problem there is if I know the message ID that you're expecting for the answer coming back, and I

want to forge a reply, the only validation there is whoever gets there first with the answer to the question you asked, tagged with that message ID, that's it.

There is no password, there is no validation. That was considered sufficient at the time. Here some math if you want to look what's in the slides. I didn't do all that while in my statistics class, so I will not attempt to explain that piece of it. But what it boils down to is, birthday problem says, what is the odds for a given number of people that two people will have the same birthday?

And if you go through all of the math what you discover is that very quickly, you get to this point where you are pretty much are guaranteed that two people will have it, and it's far less than 365 people. This room probably has enough people that there are probably two people in here who have a birthday. Using that same math, what it means is that instead of having 16 bits, or 64,000 possible random message IDs, I don't really have to do that because of the way the birthday attack works. By the time I get up to five or 600 message IDs, I have an overwhelming chance that I'm going to match it.

So what that means is, if I send you 500 packets with the answer, and different message IDs, I have a 90% chance that if all 500 of those packets get there before the authoritative servers [re-lancer] gets to you, I win and you will believe. And that's where we get to the Kaminsky Vulnerability. In the summer of 2008, anybody who was supporting a DNS server had a very exciting summer and fall.

As did pretty much all of the TLDs and anybody who was notified early. Because what happened was, Kaminsky was a researcher who sort of

did some math and suddenly discovered, gee the birthday attack actually applies to message IDs, and that's the only validation we have. So the attack works as follows. I query the authoritative… Or I query my recursive resolver that I wish to poison with a string that I know it can't have possibly cached.

I take a deliberately random string, so if I'm attacking example dot com, I will have completely random string dot example dot com, and I will ask the recursive server for it. He's going to go off to the authoritative server for example dot com, but in that period of time before example dot com comes back with no such domain, if I can send five to 600 packets with what I want to shove into your cache, and it gets there first, I win and my record goes in.

And because it sits in the cache with whatever the TCL is for the server that's in that record that you get back, I as the bad guy am going to put an inordinately long TTL in. So once that cache is poisoned, it stays poisoned for a really, really long time. So, the goal is this. Essentially what I do is I put stuff into the glue as one way of doing it, that basically says that is the A record.

If I want to be nasty, I can poison the NS records for that zone. And now my name server gets all of the subsequent requests. But either way, if I can get stuff into your cache, and it's a very large cache. You know, my employer Comcast, we have millions of customers. If I somehow got Bank of America dot com poisoned in Comcast's recursive caches, it's a very bad day for the bank.

So, one of the things that we did when we were trying to fix this quickly was, we said, well, first of all, most of us actually did everything on port

53. So, let's try and make this more random because you do have to match the IP address and the port of whoever is sending the query, as well as the message ID.

So, in order to increase the impossible entropy that they have to match in the number of replies, if we randomize that source port for the query, we have increased the amount of packets they have to shove at us to give them a chance to poison it. Doesn't solve the problem, but it makes it a little taller, a little harder to jump over. And so, we increase the source port depending on the implementation and what you configured, eight or 16K, so this actually increased the time that it took to poison the cache.

It doesn't really prevent it, but at least made it more reasonable. Unfortunately, even at the time where 10 meg pipe was fairly fat, it showed that it went from minutes to hours, but it was still, you know, effectively only hours before you can go through. So, what do we really need to do? Well, obviously increasing the randomness of message IDs, source port all of it certainly helps.

One of the other things that happened was we started becoming much more picky about what went into glue. One of the things that happens is we take glue, because you have to in order to get that problem of what happens if the name servers are in the zone you're querying. But the first thing that most implementations now do is as soon as they actually do that, they will do a dig at that name server for DNS records.

They will take that data, and they will overwrite the glue with the authoritative data from the authoritative server. And we started keeping track of the fact that I got this as glue, not everything in my

cache is equally valuable and equally trusted. Only the stuff that came from authoritative servers was trusted fully.

The real solution, of course, was DNSSEC. Where if it doesn't validate, I know that it wasn't from the authoritative zone that has the keys that sign that record. I do have to validate, and they have to sign there zones. If I'm running a cache that is not validating, I can still be poisoned even though you signed your zones, but at least, if everybody starts doing it, we have some reasonable defense.

The other piece of fun, of course, is that a lot of folks have started overwhelming authoritative servers. Sometimes it's an inadvertent overwhelming. They have a script that sort of gets a little out of control, sometimes they are deliberating trying to take your authoritative server out for some reason.

Denial of service, possibly overwhelming your server so that it's easier to poison you with Kaminsky hack because you can't get your answers out as fast. Most of the stuff we're seeing, I mean, right now is just sher brute force. Because of the fact that most DNS is done with UDP, which is a single packet one way and a single response back, it's very easy to just spoof source addresses in your UDP packets.

And unfortunately, because of malware, and whole security things I'll talk about a little later, as to why most end hosts are easily corruptible, there are botnets running around, multiple botnets that have tens of thousands of bots available to them. Open resolvers, resolvers that will make that DNS query for you with no restrictions, and again, there is a group that is trying to keep track of this problem and help people fix it.

It's the open resolver project dot org. Check out their webpage, they have a link that will also let you check if any of your servers are open.

But at this point, when Jared [?] did the initial sweeps, started, it was on order of 36 to 38 million open resolvers on the Internet. Fairly quickly, within a number of months, through a whole lot of effort, we have dropped that down by 10 million, which is nice. The problem is that we have been stuck at 26 million-ish for quite a long while now. And a lot of those are CPE device, customer routers that have truly broken, truly antique software that acts as open resolvers.

They are not remotely upgradable. We've had talks here about that as well. And so they are likely to be that way for a long while. And that combination of 10 thousands of bots, each of which can generate spoofed queries, each of which can then send those queries to an open resolver, means that you can generate huge amounts of DNS queries.

Because there are large records now out there, specifically DNSSEC signed records, you can send a 40 to 80 byte packet and get two to four K of output to your intended victim. So when you do tens of thousands of bots, millions of open resolvers, we are seeing not just gigabytes, we are seeing hundreds of gigabytes on some of the biggest attacks.

The latest one tipped a little over 400 gig, and that has been increasing even in the last year. These statistics are actually already quite dated, but pretty much everybody sees some kind of DDOS, at least once a month, once a week. Many of them are having literally hundreds. Verisign, who is the one that has publically admitted this, is continuously under attack at all times.

They basically build all of their nodes for all of their root servers and all of their TLD servers, with the assumption that they have to survive hundreds of gigabytes of attack, because they are. So, part of the problem is that UDP in-balance, this is probably also why you hear about this with the NTP implications as well. If you can send tens of bytes and get thousands of bytes to your intended victim, you can generate huge, huge attacks.

Something else to keep in mind, NSEC 3 was mentioned here. It is obviously a way to prevent people from enumerating and walking your zone. The way it works is, that instead of having what the next record is, and NSEC record, their purpose is to cryptographically prove the nonexistence of a record. So what happens is, if I have AA dot example dot com, and BB dot example dot com as two records, and you say, give me the records for AB, what happens with NSEC you actually have a record that says, hi this is the NSEC record for AA and the next record in my zone is BB.

Therefore you can look yourself and you can go, oh, AB is between those two, both of those, the AA NSEC is signed so I can validate. BB is signed, I can validate it. Therefore I know that there is nothing in between those, therefore AB doesn't exist and I know that with DNSSEC. But you can walk the zone.

NSEC 3 does the same thing, but what happens is you actually have a hash of AA and a hash of AB, and it puts those in a zone file but when you ask for AB, you essentially have to come up with the hash. And what gets back to the resolver is, those three hashes to prove that it doesn't exist. The problem is that you then have to do prime number

math on your authoritative server every time someone asks you for that nonexistent record.

So if I know that you are using NSEC 3, and that you are DNSSEC signing that zone, and I want to make your name server's life more miserable, I just sit there and I just do random string dot your zone, plus DNSSEC, and you get to wheeze. Something else to keep in mind, one of the other problems that's out there for authoritative servers, is making sure that it is actually your authoritative server.

There have been two talks here about how to not have people doing some form of phishing or hijacking attack, to change the registry entries for your name servers. If I can get the com server to use my name service instead of yours in the com zone, for example dot com, I don't have to poison caches. I win. I already have the name servers. So, definitely make sure that you are getting those security checks. That you're using decent software.

That you're using multiple validation methods. That you're training your folks on how to deal with phishing attacks, which are disturbingly effective still. And think about what TTLs you put in your name server records. You want them long enough that it's not abusive, but short enough that if it does get into cache with either a hijack or a poisoning, that you're not looking at two weeks of fun to get it back.

So, what are we doing to defend our servers? Well, there are two sort of ways of looking at name servers. If I am a recursive server, I know who is asking me, because in theory I know who my customers are. If I'm not an open resolver, I should know which IP address ranges are going to ask me questions.

I don't know what question will be asked because I have to resolve anything on the Internet. But I at least know who my users are. Authoritative servers are exactly the opposite. I know exactly what questions I will be asked, because I will be asked about the zones for which I'm authoritative. But I don't know who is going to ask it and I can't restrict via IP address, because in theory my job is to answer questions from anyone on the Internet who needs to know about those zones.

So, ACLs are more usually effective on recursive servers. Not absolutely correct, but more generally the case. Certainly one of the things you can do is you can have servers that can take more abuse. If your servers can deal with a million queries a second, then somebody firing 500 queries a second at you, isn't really going to make it work too hard. You may not even notice the noise level.

You can try doing things with clustering and load balancing, but that is really mostly effective on recursive servers, not authoritative, because you just don't know. Response rate limiting is something that has gotten a lot of interest. This is where, essentially as an authoritative server, or at least potentially as a recursive server, the less useful.

If I notice that somebody is hitting a particular record really hard, above a certain threshold, or I notice somebody is doing it by query type, any records are a very popular resource record type to be abused, because they don't have a lot of uses but they do get a really nasty payload. I can start throttling back, and there is a paper mentioned here by Dave Knight, that talks about some things.

Basically what the usual method is one of two things, and it has to do with what's called slip value or of two. A slip equals one. Every single time I am in throttle mode, I will set the trunk eight bit and force you to reconnect to me with TCP. Because TCP can't be spoofed. Therefore, I know who the real query is, so in most cases, these botnets will go away or stop bothering you about that. And slip equals two is basically alternating between dropping packets, every other packet is dropped, the other ones are set TCP bit, which in theory cuts down the load on your routers.

Dave had a presentation which basically talked about how effective this was on some of the root servers. It seems to be in pretty good use with TLDs, and with root servers. Gets a little more dangerous potentially with other name servers. The second paper here is the one that talks about cache poisoning, because remember that cache poisoning is me getting enough packets to your recursive cache that I'm trying to poison, before the authoritative server gets the legitimate NX domain back.

Now, if I am a bad person, and I know that the authoritative server is using response rate limiting, then all I have to do is deliberately abuse that authoritative server to kick it into that mode, where it is now throttling back massively what its responses are and forcing TCP. That widens that window in which I as the attacker get to send my forged packets, and get there before that authoritative response.

So, not saying ROL is bad, but it is definitely one of those, make sure you understand all the tradeoffs before you make the decision may be doing you more harm than good. Other things you can do, fatter Internet

pipes will certainly help. Again, if your problem is that you're shoveling out, you know, 10 gigabytes, or 10 megabytes of traffic over a pipe, and you have a 10 meg pipe, life really sucks.

If you have a gigabyte pipe out and you're travelling 10 megabytes, okay, it will probably show up in your logs, but you are still going to be getting out the legitimate traffic you want. That's the good thing. Unfortunately, on the other side, it also means that, for instance, if you are a recursive cache and you're shoveling out lots of responses because you have a nice fat pipe, it does mean that you are a much more effective hammer to be used against someone in an attack.

You can increase the number of authoritative servers you have and spread them around, up to a point. There are several issues, one of them is client stacks don't always do the things you want with authoritative servers in a list, or multiple A records in a list anyway. So, it may not solve the problem for all clients.

And there is a practical limit in that, in many cases, if you go over 512 bytes in your response, there are still enough places that don't do the right thing with DDNS zero, or with [trunk-cation?] bit, and they want get a usable response. Any casting is probably one of the more effective responses.

In the same way that the reason that it's very hard to kill botnets, because they come from so many different address ranges that are trying to do the [hackles?], to stop them is operationally painful because they continually change addresses. So you have a moving target at all times. In the same way, if my servers are any casted across multiple

different addresses, the botnets will only be able to tap the one that's topologically closest to them.

So I'm actually reducing my attack surface to the entire botnet army by being any casted. High availability is very popular with an awful lot of services, like web servers and other things. Unfortunately for DNS, it's not really a terribly effective offense. One of the other things that I will talk about here is, stressing how important it is that we all be good citizens.

You know, you may not have somebody attacking your TLD, or your recursive cache. But the reality is that at some point you may. And if everybody stopped worrying about it until they were the one attacked, we're going to have a very long period of time in which this is ugly.

Basically, I stress this whenever I give talks on IPv6, or DNS, or DNS security to everybody. If you're letting spoof packets out, you are part of the problem. BCP 38 and BCP 84 have been around for years. It's not that hard. There are various other forms of source validation for routing that are in production routers that work at line speeds from the major vendors.

It's just a matter of configuring it and turning it on. And it will be very sad if we had to wait until there was a very expensive lawsuit from somebody who had been attacked, and claiming somebody else was negligent for not having done this before we actually move.

One of the other things that I've mentioned, the open resolver project, make sure that you're not running open resolvers. The last thing that's not on my slides here, but I'll mention. The SSAC group, the security

and stability group for ICANN, has a document that came out recently. SAC 65 which talks about the DNS amplification, has an excellent set of bibliographies and references on where you can learn more about how to not be part of that problem.

So, what do we do? Well, we probably need to revise some DNS standards, but the ITF is not a speedy or nimble organization. Realistically speaking, if we wanted to change existing RFCs and do business of them, or updates and [?] of current RFCs, we're looking at years.

Source validation is a much more likely thing to be done right now. At least it's something that we have the technology for. There has been a lot of debate in the DNS op working group and on various mailing lists about, should we be using TCP more? Since UDP is fundamentally insecure in that it is very easy to forge.

The reality is that TCP is also much more expensive on servers. By default, most servers have very few TCP slots, even if we up them practically speaking, we can't keep up with the volume. There is a reason why UDP was picked for the volume of traffic and quantities of queries per second that most people are seeing.

One of the scary things… When I turn on my iPhone, and it has not been on the network, if I'm using a fairly normal set of apps including mail, Facebook, Twitter, and a couple of other things, my phone instantly makes, I think it's 74 DNS queries, the instant it comes up.

And this is just going up. Everything is doing DNS queries, webpages do dozens and dozens of queries. UDP is the only thing that can keep up

with the current volume. There is an interesting draft that was done a while back that is being resurrected, which is sort of a halfway between, which is DNS cookies. Which is, kind of, sort of state-ful UDP. Don't know if it's going to go, but you should read the graph.

You should probably get on the DNS OP mailing list, if you're not already on it and talk about it. So, just very quickly some of the things that I've seen and heard talking to various customers at various shops I've been at. DNS is not just an attack point or an attack vector. It's a critical piece of their infrastructure. Most of the botnets are done with command and control, whether a small number of servers that actually tell them what their job is at any different point of time, what the new code is, all the rest of it.

And I like to joke about this, but it's not really a joke. I would kill to have programmers that had the kind of good habits that these botnet guys use. You know, I have gotten tired of explaining to coders why hard putting IP addresses in source code is bad. Why using partial DNS names and counting on search lists and suffixes is a hideous problem. Why you should be checking result codes and all of the rest of it.

And sadly, the guys that are doing all the right things, are the botnet guys with their command and control. They are using DNS names, they are changing those DNS names with these little algorithms that cryptographically go to the next one in a sequence. They are doing what's called fast flux, single and double flux, where basically they'll rotate through and constantly change the actual A record, or a group of A records that respond to a label.

They will change a label completely on pre-determined intervals very quickly. They will actually the clients wake up and do a call home when they come online to let command and control know that they are back and active. All of that stuff is being done with DNS and DNS queries, and state being put into labels in DNS.

And all of the big ones that you hear about in the news, storm configure, all of those, DNS is a critical piece of that infrastructure. And basically there are three pieces to cleaning this up. Stop spoofing addresses, stop open resolvers, and stop botnets. And that third piece essentially, how do we stop having clients as botnets? We'll talk about all three of these.

So, antiviruses, we need them, people are still running old code. You know, the reason why a six year old rootkit still works is that they're people using eight year old operating systems without patches. Unfortunately, there are so many munitions that there are literally thousands a year. I think the last time I did the math, if you don't get new virus definitions at least every three minutes, there is a new virus that is out.

That's how often they're coming out. Virus detectors are pretty easy to disable, and they don't have a very good success rate at catching a lot of the new viruses that are out there. Perimeter defenses are necessary, but they won't catch everything. Part of the reason why botnets are using DNS instead of IP addresses is because most of our defenses have to do with [hackles], which is IP address, port number, protocol number.

If you continually shuffle between all of those, most perimeter devices can't catch what you're essentially left with, they're the signature based

ones, which do work, but are much harder. Or you have IP based reputation lists. Some of the things like spam, but again, someone has to notice the infection, somewhere it has to hit it, some honeypot has to be updated.

They don't get updated all that quickly. One of the interesting trends is RPZ response policy zones, where you do have a reputation feed, but it essentially has a wild card match, and you can match based on DNS names, or IP addresses of name servers that are common to those.

It uses all of the things that DNS is good at. It's very easy to replicate because [?] fast, it does protect your clients. It does gives you options for essentially lying in DNS and caching it. Another interesting tool not completely, but pretty much, we need to do everything. We need to be able to do source validation, BCP 38. We need to clean up open resolvers, and we need to do all of these things with our clients.

We need to keep up with antivirus. We need to keep up with perimeter defenses. And we need to do things like RPZ and continue to develop new stuff. So, that's the basics. Questions about anything related to this?

EBERHARD LISSE: No questions? End of the day. Overwhelmed. Thank you very much. For me, this is also very interesting. As I said, the easy way out is to just get big pipes, yeah, like PCH wasn't really bothered by our attack because the volume is too small. But if it can get one or two people who do this into jail, others are going to start thinking about maybe this

is not such a cool prospect, and the energy required becomes much, much bigger.

Yeah? And hence I am really a fan of using law enforcement. In any case, thank you very much. [Applause] That leaves us with what I usually call the host presentation. Since we have been here two years ago, we're not going to dive too much into the local setup, but Mon-Loi is going to give us a little bit, a few slides about that. And he's basically talking about the abuse system here in SG.

MON-LOI PEREZ: Okay. Good afternoon. I am Mon-Loi Perez, associate consultant at SGNIC. So I'm here to talk about abuse management system that we have developed. So first, a little bit about my sides. We'll talk about SGNIC, the types of abuse, and the measures that were addressed by using this abuse using our AMS.

Some statistic and experiences that were gathered from the AMS, and some conclusions. So about SGNIC. SGNIC is a national domain registry for dot sg names in Singapore. So it's a wholly-owned subsidiary of IDA, which is the Infocomm Development Authority of Singapore. We interact with external organizations such as ICANN, IANA, APTLD, ccTLDs, APNIC, SingCert which is our local Cert.

And some statistics, we have about 155,000 dot sg domain names as of December 2013. And about 70% annual increase. So the chart below is a breakdown of our domains by categories. So as you can see, 58% is from dot com dot sg which are for business entities. And next is dot sg, which are for usually personal registrants.

More into our resources. So we are a small team. We have a total of 10 employees, three are technical. We have a shared registry-registrar system. It is an outsource source. So basically it's also, we have an HA setup, so typically two servers for critical applications, RAID 1, redundant power, load balancer, etc. We use a mix of Solaris and Linux operating system, and also a mix of SPARC architecture and Intel.

For the application side, we use Java Web Application and Oracle Database. So all in all, the servers, we have about 10 servers. But soon, for this year, we have a project to virtualize it into three. So basically, two for the production and one to be a DR. And we also have a monitoring system plus a NOC vendor and in-house. So the NOC monitors uptime, performance, and incident correlation.

If I have time, I can present a view of our NOC. I didn't put it in the slide, because it may be a bit sensitive. So, last week we had four secondary name servers. So the types of abuse. So, all of these boxes are abuses that SGNIC is concerned about. But the green ones are the abuses which we feel more effective measures can be done.

So, one of the… What are the effective measures that we did? So that's where the AMS came in. So it was developed in house. It's typically built on a LAMP stack, Linus, Apache, My SQL, PHP. It's completely open source. It was operational in 2012 after three to four months of development.

It basically detects and tracks some domain name abuses such as malware, phishing, incorrect/suspicious registration, suspicious registration from bulk registrations, and some DNS wildcard usage. It

also provides the statistics and tracking to better understand the nature of these abuses.

So, this is the automated… One of the features of the AMS, so it does automated scanning of domain name against third party website scanners, or reputation databases for malware distribution and phishing activities. We do manual verification of flagged domains to confirm abuses. And then if abuses are confirmed, we send notification and continuous reminders to the community such as the registrant, admin, hosting provider.

This is an operation view of the AMS. So on the left hand side, you can see the domain name, the domain name and then the IP address. Next you can see the status of the cases, meaning if the status is open or is the status – or if it was resolved. And then you can see on the next column, is the next application or system that was used by the domain, meaning is it running on an ISS, is it running on a LAMP stack?

And then next is the type of attack that was used to let's say inject malware into the domains. And then the last few columns are the third-party results. Whether it is flagged as malware, or is it like a phishing. So for the third party results, we started with two, sorry, two services, two free services.

But then eventually, we realized that we weren't able to detect more, so we added a few more feeds. So, detection and tracking. It also provides inaccurate registration… Providing inaccurate registration information is often a precursor to domain name abuse. So in AMS, we built an early warning by checking the accuracy and completeness of new registration information.

So how do we do it? So this is an example of simple domain name registration. So first for the owner, we check the company number versus the company name. So if the company name, if the company name doesn't… So we basically do that to check if the company name is in fact the company which was registered with that number. So same case for the postal code. So with a given postal code, we counter check if the given address really matches the postal code.

So we also, we think that some suspicious registrations come from bulk registration. So we built a bulk registration detection in AMS. So one of the rules for the bulk registration is like, for example, if registrant registers 10 domain names, all with different registrant name but same email address in a day, we think that something is suspicious. So our manual verifiers than go and check those names.

And the rule is 50 domain names for the same registrant name using the same email address in 30 days. So that's another rule we have in for the bulk registration. So, AMS continuously monitors all domain names. All new names are scanned weekly for three months, there after monthly scans follow.

So, these are some of the statistics that were gathered over time. So of these statistics, as of January 2014, we were able to scan about 156,000 domain names. About 30 average confirmed abuses per month. So the graph below shows the number of confirmed cases from 2011, since it was – in the year that it was operational, up to 2013.

So as you can see, in the first quarter of 2011, we didn't have much detection that's because like I said, we started with two free feeds from the Internet. So in mid-2011 we realized that we weren't detecting

more, so we purchase a service for detecting more malware and phishing cases. As you can see, the number of confirmed abuses grew from then on.

So this is additional information that we have gathered through AMS. So we were able to find…

[END OF TRANSCRIPTION]