DUBLIN – Tech Day
Monday, October 19, 2015 – 10:00 to 18:00 IST
ICANN54 | Dublin, Ireland

EBERHARD LISSE: Good morning everybody. Can we settle down please?

CHRISTINA: Hello. I'm just going to announce the session for the recording.

EBERHARD LISSE: I didn't want to start it… There is some echo here somehow. Before we start, I want to just make sure whether everybody is here. We are supposed to start at half past 10, so if all of the presenters are here, we can actually start a little bit earlier and see what happens. I see [Mary?] from Venezuela. I see Patrik Wallstrom sitting over there. I haven't seen [Urzua?] From Chile. There you are.

Steven Farrell, I haven't seen him yet. But if we can, if necessary we can start this, then [Wullink?], there you are, and then Korcynski? Not yet. We're missing also presentation, and Conrad is also here. So what I think we do is just start a little bit early and make it up as it goes along.

Yeah. All right. Then Christina can start the proceedings.

CHRISTINA: So this is the 19th of October, and it's the Tech Day in room Liffey H2.

EBERHARD LISSE: Okay, now for the recording. God morning everybody, my name is [inaudible] Lisse. As you know or not, I'm the ccTLD manager of dot NA. And I chair the technical working group, which used to be the ccNSO technical working group, but we have been starting to branch out a little bit, and we are trying to not focus on exclusively or predominantly on ccTLDs.

We have a relatively packed session, in many topics are focused on data. And I think I must go to the other laptop because that's the one representing from.

I just want to go quickly through all of the presentations as we are having them. The first one will be from dot VE, they did the data clean up. I myself for dot NA, I am also busy with this. Many of us are busy with this. So it's obvious thing to hear what we're doing.

Then Patrik Wallstrom and a colleague, I forgot the name of, it was a French speaking colleague, have written a replacement to the French zone testing tools which I like very much when I played with it. And they were willing to present it both on the

front and on the web, and back and in [inaudible], so it can be twisted a little bit.

Then we will hear from the new dot CL. We wanted to have the same presentation last, in Buenos Aires, but it couldn't happen for [compassionate?] reasons, so this time we do it. Steven Farrell will speak a little bit about some joint project he has done with dot IE. Then we will have two presentations from basically from SIDN.

They have back on a project analyzed, they are data and also on reputation metrics, how you look at, computation not only as… For example, if you have got 10 [foreign?] domains in a domain TTLD, there is 500 names that has a different impact, then you have got 10 in one that has 10 million.

So the normalization aspect of this, how that it gets into, affected into the assessment, we had about a month ago, a really stupid commercial, things running, running over the Internet, how bad TLDs are, then when we looked at the methodology, it was really silly. So the more thought we put into development, how we assess this, when you accuse some TLD of being [inaudible] it's better to have one's stacks lined up really well and compare this, compare only what's comparable compared to the font sizes.

And then, there have been some changes to the technical team of ICANN. David Conrad will give us an overview, as I've said joined them. so we felt that it was a good thing to do. Somehow I must go on to the next page here.

Well Jan will then give us a historic view into Internet domain names, early Internet domains. So it's more like historic view, and since he was late and it's a light topic, he gets the after lunch. Steve Hollenbeck is going to speak about [inaudible] is he? Scott, sorry, Scott. Oh yeah, there you are. There he is.

Andrea Philips will give us a second reading of [inaudible] project. We have heard about this a year or two ago, so it's a good time to receive an update. Shane [inaudible] will talk about the [inaudible] DNS service. I haven't really heard about it, but it's sort of simulation thing where you can look at key turnovers and key holdovers, and from the whole service.

At last, then we will have a [inaudible], Janelle McAlister and David Piscitello. Janelle McAlister is from [inaudible], and Dav Piscitello is from ICANN. They will talk about an incident and their response. And we will see whether we can use this for future incidents. Fair warning, [inaudible] is a registrar that deals with large commercial names, so that may or may not be transportable to other systems, but they usually make nice presentations so we felt that we would have that.

[Inaudible] is giving us the regular IETF spot. He will talk about several projects that they have going on with DNS privacy. And a late edition is from [Crypt?] data, also I must give full disclosure say that, [inaudible] in Germany and they provide a backend for several gTLDs and they also provide DNS service for deserving ccTLDs. So I'm just in the spirit of full transparency, I wanted to make this before.

And they have mentioned something that dot Swiss has to publish the domain names for two to three weeks before they get registered. And they have chosen a way to go through an RSS or [atom?] feed, and I found that particularly interesting because one may be able to use this for other things such as [inaudible] and things like this.

Then Jay will, as usual, give us some interesting thoughts about what happened during the day, but he will also get 10 minutes to show what they have done with their data. They have, a while back, got [inaudible] and collected the data, and they're making it available for review. So he will get a few minutes to make a short presentation about that, which ties in what dot NL has been doing.

So, sorry. I would then call in Mary, [inaudible].

For all of you who haven't been here before, we present from the Adobe Connect system, both for the local participants and for

the remote participants. So you have to bear with us a little bit until we get the presentations up. We want both the remote and the local participants to see the same pictures and the same presentations.

RUMARY RICAUTE: Good morning. My name is Rumary Ricaute. And I'm from Venezuela. And today I will present the running process that is VE compliant for cleaning up the database.

Okay. Let's see the topics. First, I present the VE timeline, to know the event that make history on the VE, on the Venezuela ccTLD. Then, let's see the VE structure. And the main process…

And then migration and update from the VE. Let's see the statistics, and suddenly the results of effort. Okay, the timeline. In 1991, the ICANN delegates the VE for Venezuela. In 1998, creating the Network Information Center. In 2000, implement the Openreg, with the model registrar and register. In 2005, the process of the registration of the domain was automatic.

In 2009, starting the transfer to the ccTLD to the regulator of telecommunications in Venezuela. In 2009 there was a virus, we had a virus attack. Even a record of more than 40,000 domains. And the same year, the NIC VE finally becomes a part of the CONATEL, but without the domain release process.

In 2012, VE down for a few minutes, but, by a deficiencies of hardware. In 2013, implementation of replica with Anycast. 2014, start a new sys admin on the new work team for VE. We are improved the new process for NIC VE.

In this year, reopening the WHOIS service through the 43 port, and implement of new replica in July, approved by IANA. And it started a test for any new release process.

Okay. About the VE structure. Distribution of the DNS replicas of VE, you can see in the world wide distribution. We have seven physical replicas, five in our country for our management, but one of them is out of capital of Venezuela, Caracas. Another in Chile, another in California, that I mentioned previously with anycast type. About the domain extensions, we manage the commercial extension and the special extensions for the government entities.

And the commercial extensions, we have dot com, dot net, dot co, dot org, dot info, dot web. And of the special extensions, we have gov, the edu, dot mil, dot tec. We have other extensions. Use by domains names, active and inactive. Okay.

Well in this slide, I will present the main services for NIC VE. We have billing database server, the webpage or web server, the mail server and the [share] registration system with the domain database.

This server, have an obsolete hardware. Work in the Debian three dot one. It's a few below debian version. In the shared registration system, we have the zone file, and others zone files. We've managed the open rec with the EPP on the WHOIS and the domain name server.

Doman name database. And we expect to separate this service, services sorry, for different server to improve the operability and security on the platform. And actually, running the WHOIS in other server and the domain database master server of DNS, is under stress testing.

Okay. A part of the process, we expect the cleaning domains database because the, or thanks to the virus attack, we managed several information that will should be cleaned. At the same time, we start the project for a national project, a running project for national plan for the domain names.

We manage the, sorry. We manage the terms of use for the registration of domains. About the domain release [inaudible], we manage the extension. We have other extensions that dot arts, dot store, dot int, dot nom, dot rec, dot gov, like government, dot firm, dot COP, dot E12, and that use by the many, in many domain names. With the old domains, the banned domains, and the inactive domains. They are part of the domain which should be released.

Okay the number of the domains to release is nearly through 230, sorry, 230,000 domains. And the process is actually running, it started a few days ago. Implemented five or six seconds, to release one domain. The final times, I will present the next slide. Okay.

About the statistics, since 1998 to the third [inaudible] of the present year, in red, we can see the domain names registered active. In the blue, you can see the sum of those domain. You can see in the 1998, with, when the domain names released was lost, the sum of the total of domains increased suddenly.

For the, until the date of September, the number of total of domain, reduced 13 VE, is the 319, I'm sorry. Okay. Actually, for one, two, three, four, five, six weeks, the domain name release will be complete.

You can see the number of, the total of domains will, sorry. I'm so nervous.

The process team, the process they used, six seconds, what I said, and we implement with 69 weekly hours for the 213,000 domains, for a total five or six weeks. And per week, we released 41,000 for hundred domains. And a pain of the result of effort, improve the VE platform. We have the more availability for the, we can see that is 59% to the 98%.

We expect we reached, we did a main release, 100% availability per web portal. And about the reliability, we increased more of the 20% with the contact all of data release the register domain. Of course, increase the efficiency and increase the security. And about the planning improvement in the VE platform, well finally we expect the veracity, the total veracity information.

Reducing response time of our process. Well sorry for this, but thank you very much for attention. For information, we can contact that email account. And if you have a question.

EBERHARD LISSE: Thank you very much. And all give a big hand please.

In particular, it's not very easy to present in English when Spanish is the main language. And but you did quite well, I must say. What I'm saying is, all presentations will be posted to email addresses on there, so if direct email is required, please feel free to do it.

Are there any questions in the room? Jacques has a question. There is one question from remote so we don't forget that when we're done with the floor.

JACQUES: So I've got a quick question. Maybe I missed it, but why are you releasing the domains?

RUMARY RICAUTE: Sorry?

EBERHARD LISSE: Why are you releasing the domains? The reason.

RUMARY RICAUTE: Because we have many, many problems with… We have many problems with the time of process, because technical persons, we have delay because number of domains. Because the information about the registrar was not, was not real or ever the payments to handle [names?] or statements that have no match with the information in database.

UNKNOWN SPEAKER: So basically you went and sorted out all domains that you felt were not complying with your rules. That means that you must be able to identify with the domain holder. And people who didn't come back to you, they got their domain kicked into the wild. And then if they wanted, they had to come and register it again with proper details in your system.

Basically what I take from that. Fairly aggressive method to do it, but I think it seems to be effective. Another question from the floor. Please identify everybody, identify themselves for the record.

UNKNOWN SPEAKER: [Inaudible] from Fellowship, newcomer. I have a question from my friend from dot VE. Do you have disaster plan for dot VE, domain name?

EBERHARD LISSE: The question is always what disaster you mean? Software engineering and hardware engineering is obviously created to be redundant, but what disaster do you mean?

UNKNOWN SPEAKER: I am technical. I already know that. Technical [inaudible], so no problem for me. Disaster plan means, if you have any failure of hardware of the registry system, of dot VE, how do you manage that? In case you have virus and your zone file is corrupted, how can you manage that?

If your DNS server is down, how you manage that?

RUMARY RICAUTE: Well, we have managed with the database backups. If the DNS server have, if 15 days to [business] days to recover, to wait to recover the information in this, the same day we, the [inaudible] me and [inaudible] me, repair or try to recover the DNS service.

[Inaudible] they haven't pass out an hour, I think.

EBERHARD LISSE: Any questions from the floor? Then there was a remote question.

UNKNOWN SPEAKER: Edwardo Menendez, he would like to know if you have any plans to deploy DNSSEC and WHOIS RDAP on dot VE?

RUMARY RICAUTE: Well yes, I planned that, I'm here to the DNSSEC because I'm interested to apply the DNSSEC. Maybe the next year, the second [semester?] of the next year, we will start with DNSSEC work.

EBERHARD LISSE: Do you find that there are many people in the room that will be willing to help with DNSSEC, and we can also give you some contacts with Spanish speakers, why is one of them, at least his wife is, we have some who will even be able to contact you in

your native languages, which makes it much easier, but there is many people here in this, on the floor here, who are quite willing to help.

UNKNOWN SPEAKER: [Inaudible] for ICANN. Additionally, Wednesday there is a DNSSEC workshop, and I'd like to invite you.

RUMARY RICAUTE: Okay, thank you.

EBERHARD LISSE: Okay. Thank you very much.

RUMARY RICAUTE: You're welcome.

EBERHARD LISSE: And the next presentation would be… [Applause]

The next presentation would be Patrik Wallstrom. And what was it again? Vincent [inaudible]. My French is probably better than my Spanish, though.

UNKNOWN SPEAKER: But your French is better than my English.

EBERHARD LISSE: Don't worry. The point is dot AF, dot FR had a zone check or two, which I have used in the past, which was quite cool because it automatically goes through and gives you a lot of information you can sort of even click on how [inaudible] you want it to be, how much you want to be informed or just whatever. Quick check, and recently I heard, I don't know how we claimed about it, we heard about it that there was a new tool.

And I had a look at it and it was both a web thing and a PERL script, some modules that you can work offline. And I found this so cool that I sort of passed that Patrik, I couldn't come here, and so Vincent and Patrik then, Vincent approached us independently so we then felt that this is something that we should do. Go ahead.

VINCENT LEVIGNERON: Good morning. So my name is Vincent Levigneron. So I know that it's not very easy to pronounce. And but just give me five minutes for short interventions, so I am sorry if I speak very fast, because my English is not very easy to understand. So sorry for that.

So I will take 20 minutes, just for the intervention. So I'm here to present you [Zonemaster], she's a new DNS delegation tool. And

I am from AFNIC and Patrik is from IIS. Why Zonemaster? AFNIC and ISS have their own DNS checking tool, and IIS use DNSCheck, and AFNIC use Zonecheck, but those tools have their drawbacks.

For instance, DNSCheck doesn't provide deterministic results, and Zonecheck is written in legacy code, in Ruby. Ruby is not bad, but in fact, nobody at AFNIC use it every day, so it's always difficult when we have about [inaudible], if we have request for a new feature, to add it.

So we decided, IIS and AFNIC, to create new better tool for check DNS delegations. And we decided to collaborate because we have the same goal, and we decided to collaborate to create a new reference tool with a join requirements and specifications. The collaboration started in October 2013, and it took us one year of work to organize a project and task between us, and to discuss a variety, common requirements and specifications because a new tool should be able to fit with what DNSCheck was about to do.

So it [inaudible] sometimes to write a common specification. And when you are to develop a new tool from scratch, and it has been written in PERL. The first release was in December 2014, but it was obsolete later because it was not completely stable.

So the publically announced in February 2015. It has done one year. What is Zonemaster? Zonemaster is an open source project. As I told you, it is written mainly in PERL with some Javascript, mainly for the going. It aims to be a state of the art checking tool for the Internet domain names. And it can be used to check delegated and non-delegated zones.

It's very useful because, for instance, if you have an issue with your DNS provider, you can use our tool and check domain names that already exist. And if you want to modify your domain name DNS information, if you want to create a new domain name, you can test a zone even if it's non-delegated.

Process is hundreds of tests. And Zonemaster also provides three kinds of interfaces. The ones that can be used for non-technicians, on our website, it's a web interface. If you able to use a common line, you can use the [inaudible] interface. And if you can program, you can use and integrate in your own system with API.

It also provide different kinds of level of output from very low output to very high output, if you active a bit more. Than you have many, many more information screen. And you have three kinds of root put on web, text, or JSON.

Output can be in English, in French, or in Swedish. We can add new languages, but I speak French, Patrik speaks Swedish, and

we need someone who speaks other languages. So if you can help, you are welcome. And Zonemaster can be tailored for your needs. And Patrik make you a demonstration of that just after the introduction.

The title is Zonemaster. What is the master in Zonemaster? Is composed of four different models. There is the engine, the CLIs, the backend, and the GUI. The engine is a piece that must be installed all of the time because it implements all the test cases. There are 10 different categories of tests, and there are 56 different test cases.

The CLI is used to interact with the engine in the command line user interface way. And you can add different kinds of logs as texts or in JSON. The backend is a module that is used to interact between the engine and the GUI. And it offers a very interesting feature. You can store results in a database.

And the GUI is a user interface that runs tests and presents the results. And you, so an interfact that gives you access to the history. I show you just after. And we have many quality considerations when we developed Zonemaster. We used Perl Critic and Dever Cover. The Cover is pretty good because it is about 90%. And we have hundreds of non-regression tests, which are used when we accept full request in GitHub and regressions tests are processed by Travis.

This is Zonemaster web interfaces. You can test it if you go on the website, Zonemaster dot net. Yes, it's visible. So what is very interesting, it's very easy to use. You just have to type your domain name in the, and the [inaudible] runs the test. What is very interesting in the web interface is that each test has a unique identifier, as we can see, this one is 399128. And it can be referenced in an URL.

So you know that DNS is not static, so you can make a test, find a mistake, and it can be fixed 10 minutes later. So you already mention what happens when you do the test. So you can use that URL with a unique identifier to refer to the issue unit. And as you can see, there are different categories.

I told you there are 10 categories, and they are in green. So if it's green it's because everything is [inaudible]. And you have very few details, but if you click on the checkbox, that's what I did for the last category, which is zone, you have all the results of all different tests that are done on those different categories.

This is an output you can obtain with CLI interface. Of course, it's not colored in reality, just for the presentation. So you have different reverse [inaudible], etc. And you have an [explanation] each time you meet something for [inaudible] which is a warning.

It's very interesting to have that kind of tool because of course, you can make your own testimony, but some tests are not very easy to find when you do them manually. For instance, two nodes, DNS key with a target 7533 user [inaudible] number is not something that you will see very easily using [inaudible].

Or even if you see that URL is not correct, it's not easy issue. [Inaudible] So that kind of tool is very useful with automated tests. What you can do with Zonemaster is that you can tailor it very easy. You can add your own language. As I told you, there are only three languages at the moment. But you can add your own.

It's only one file create, there is no need to understand PERL. It's very easy. It's completely documented and [in the repository?]. Of course, if you do that, please create a request that would be to include in the project, because we need to have more languages. You also can adapt Zonemaster policy to yours, it's just JSON file to modify. You can choose which test you want to execute. You can modify the severity levels.

Of course, all registry don't have the same policy. So if you decide to test, it's more important for you then for another registry. You can modify the [inaudible]. And if you are a developer, you can develop your own tool with the library provided. You just have to use Zonemaster on your [inaudible]

and write your [inaudible]. And if you want to contribute, I encourage you to use the tool, of course, through the web or other interfaces.

And if you find bugs, please report them on Github. If you needs some feature, some enhancements or improvements, please ask for them on Github also. If you are a developer, just clone the Github repository of Zonemaster on your server and use it. It's very easy to copy and to use it.

If you develop your own tools based on the API, please share with the community. And if you need output during this meeting, me and Patrik are here to help you. So if you have any questions, don't hesitate. So it's for Patrik.

PATRIK WALLSTROM: I have used Zonemaster in my work since it was stable. So I'll talk a little bit about the applications. We have used DNSCheck for [inaudible]. We have since five years [inaudible] to publish health check reports, which says a lot of the status of the DNS and the Internet is with them.

And it was mostly based on our work with DNSCheck. We have regular [inaudible] to the registrars about their DNS quality as an opt-in service. Tried to do some zone cleaning, where we looked at the quality of DNS. We tried to find the easy things for DNS

operators to send, ask them to do it. And we also run the service, DNSCheck dot IS dot SE. As the web. Now we're switching all of these services to your [inaudible]. So, as part of my work on cleaning up the SE zone, I started to use the command interface to run tests on massive amounts of domains.

But I started quickly to, the zone [inaudible] directly [inaudible] easy to use. As an example below here, it's just use and you get a log when you run a test zone command. It is much more flexible than that if you need it as well.

So I rather Zonemaster collector to collect massive amount of domains, or test of domains, and I quickly found out that Zonemaster wasn't threat safe, so now we fixed that. It's run multi-threaded, so it's fast as well.

I store the JSON result directly from Zonemaster on the MongoDB database. A very convenient, since it's a JSON document store. So it can just store more output from Zonemaster documents, I think, and query. This is how I run the collect command, so an example application, I show you how I run all of the TLDs through this system.

So, what I do here is that I say which database we run with, and where to store the documents, and how many threads use. Use

a lot of threads, it creates a very load on… A very small firewall or something, that will…

And I chose the level of debug for Zonemaster. This is the most verbose. For that I want to have, and I also have a file with TLDs that contains all of the TLD domains. But how to analyze this? Selection of JSON documents is very hard to browse. MongoDB queries are not the easiest to do either.

So the first command here just gives me a list of TLDs which contains errors. The second command is searching for TLDs with the name server, as example of [inaudible], and giving the results among the common things.

This is a more specific question for the MongoDB database. Search for, it's a [inaudible] tag to find which TLDs have name servers that access [other servers?]. I list those [inaudible]. So this is not a very user friendly way to write your creation.

So I created something easier for this. I created a larger set [inaudible] defined queries. And I created the web interface. And as an example, I took all the domains and I published this on a public web server. So what you see here is all of the TLDs listed, and you can sort this, and you can click on the TLD and look at the Zonemaster result of that.

So as a demo, you can sort it on the error level, and just look at all of the TLDs that have an errors according to Zonemaster, of course. The errors can be a result of that Zonemaster trying to be very hard to look at a delegation problems.

All problems are due to the fact that there is a problem with a TLD, but there can be policy stuff or problems with the, best case in Zonemaster itself. But you can actually look at this and find problems with a zone, by the way. And you can see the whole log. You can click on any of these and find more common problems with a name server, IP address, or a tag, find things.

You can sort it according to the log level and quickly find the problem. This is the top level domain [inaudible], dot AX. That has different set of problems. It is not signed here, there is no DS. As a policy the Zonemaster will say that not being signed is something that you should take note of. So it's not…

You see also that there is only ASN number for IPv4, there should be several ASN numbers. Different queries like this, you can have top list of the most popular ASNs for the TLDs. This top 100 on the website. And also see the ASNs for IPv4 and IPv6 separate as well.

You can also see all of the domains. If you click on a DS number here, you will have the domains that are connected to the DS number, which makes it very easy to find whatever you're

looking for. Some cleaning in the SE zones, and find other types of problems. But very interested in TLD stuff. There is another set of things that you would like to look at.

The most popular name servers, you can see these as well. Both [inaudible] numbers and [inaudible], and how many TLDs that are used. This is the top level errors and warnings in TLDs. So an activity problem is the most common, that I cannot connect to a certain name server. And my Internet connection, I do not test remote locations on the [inaudible] next, obviously step for some of us to be able to do.

The hard part there is to actually present the results in an understandable way. There is lot of common problems among most TLDs. The popular problem is this [inaudible] series, which is not really a problem because many TLDs do very fast updates, and for all of their name service update, exactly the same time is very hard.

But for a smaller zone, it might be a bigger problems, since the zones must be upward referring, something we look at as well. This is where name service referred to the root, when it does not answer, so it can be… So there is a lot of interesting here. Also notice in [inaudible], I do not display the debug, or the tests we do with ZoneMaster, as the test specification, coming directly from the [inaudible].

And understand which test case we use for each tag's logs. We have a page, that being each log message to what test specification we have. So it looks like this, it's a really long list because there are more log outs then test cases. So there are many. We do support for different test profiles. We have yet created what we call the IANA test profile, where we only test a set of requirements that IANA have on that delegation.

I will probably do the same test again with the IANA, when I have one that I want to try. I'll run it and publish it on this website. This is the current set of IANA test requirements, but so they test a lot less than we do in [inaudible], they have like 12 requirements and we have…

There is also ongoing work within Center. We have RTF working group, or the test requirements test course, where we look at the requirements on the DNS delegation. That is the outcome of the Zonemaster product as well. You want to write down the requirements on the DNS, we should have a very good document that we can improve later on.

So the current status is to write an Internet draft where we actually try to detail the requirements now, so this is work that I'm currently doing. So, all the code is here. You have the Zonemaster on Github, and now we can also play around with the web interface, [inaudible] dot com, which is my… And the

code if you want to play around with the web interface and the collector tools by Github.

EBERHARD LISSE: Thank you very much. You're not forbidden to give a hand, if you want. [Applause]

I have downloaded this on my Mac before I upgraded to 10.11, so now it's gone again. And it installs on the Mac. Patrik told me he doesn't develop on the Mac. I just loaded it, it took, it loaded all of the modules. One module give me a little bit of problems. I had to manually basically install this module and then it installed flawlessly, and then it worked.

I never really bothered because it went so easy. So it installs on the Mac, it installs online obviously, so anybody wanting to play with it, it's fairly simple.

JAY DALEY: Jay Daley from dot NZ. So, thank you. That's very interesting. …one is, can you rake limit… …one is, do you publish the data or aggregate it?

PATRIK WALLSTROM: No. The only data that I published right now is the thing I did on the TLDs.

EBERHARD LISSE: Don Hollender?

DON HOLLENDER: Don Hollender from [Book Haven] in New Zealand. So the web interface, were you able to put in the domain name and it goes and check. Can that support IDNs?

PATRIK WALLSTROM: Yes. [CROSSTALK] Yes.

DON HOLLENDER: Thank you very much.

VICKY [RISK]: Hi, Vicky [Risk] from IOC. We'd like to encourage you to include a few tests for EMS compatibility, specifically checking to see what the behavior is when presented with an option, because we found some bad failure modes with that. And I would be happy to share with you some test case. EMS comp dot IC dot org that…

PATRIK WALLSTROM: We don't look at specific, it's a DNS testing tool but…

VICKY [RISK]: You have one test case already for…

UNKNOWN SPEAKER: [Inaudible] from [inaudible] of Egypt. Newcomer Fellowship. I have a comment that [inaudible] focus for [inaudible]…

PATRIK WALLSTROM: Of course, yes.

EBERHARD LISSE: Okay. Any more questions? Christina? Thank you very much. I will personally approach you for one or two small things that I want to play with, but thank you very much for this presentation. [Applause]

And I just looked at dot NA, and it only has got two warnings about some name servers not our control, having not [inaudible] PTR. Okay next one would be dot Chile.

JOSE URZUA: Good morning to everyone. I am Jose Urzua from NIC Chile. And I want to present the new dot CL, and that was a big project that we [inaudible] in dot CL to replace the system. At the beginning, I want to show this timeline. And with main goals that we reach

our time. Dot CL is 28 years old, and in 1997, we start with policies in dot CL registration payment.

In the middle of 2000, we started with any casting dot CL DNS servers, and IDN support about 10 years ago. And IPv6 eight years ago, and DNSSEC four years ago, and in 2015 we start with the new dot CL. This chart shows how many domain names we felt in dot CL.

Currently we are near 500,000 domain names. Well, the whole system was working between 1997 and 2012. And it was a monolithic type of recovery system. This time, have a big update in 2002, when we move from [data five] to [inaudible] database. And that system was based on demand. That means or tied to the users demand new functionalities, and we add the functionalities to the system.

And that in turn had only one registrar without a user account. And it was built with local tools. And the local tools include naturally the number for contact data, and the [inaudible] or [inaudible] for others. That means when you want to register domain names, and decided what's necessary to include the [inaudible] number. You have to include the [inaudible] number and the district…

And the [inaudible] tools, the old system was made, was [inaudible] in programming language. And my support for

database and Apache for web server. This time around, [inaudible], chose the main [component] in the system, in the same space, in the same architecture, and we have some generation to dispute [inaudible] system. And all domain names operations, payment, and invoicing.

Well, new dot CL has these main goals. [Inaudible] this account, that is necessary we will reach the modern 400,000 domain names and the [inaudible] was completely [inaudible]. Other main goals is to implement the registrar registry model with EBB of course, and to have a new architecture.

We want to [inaudible] and we want to manage three million domain names, and we want to check that new architecture with that amount of domain names. And other main goals was to have a [inaudible] dispute resolution system to replace the old fiscal dispute resolution system. And of course we want to move the old domain, so the domain names from the old system to the new system.

And was necessary to implement a domain name [transfer]. When the new system start, we went to stop the registration of the old system, and the last main goal will have a new website for NIC Chile. But this chart, [inaudible] and the protocols, we define to use in the new architecture. The first component, part,

is the registry. The registry has to register the new system and the old system.

We update the old system, the register the new system with just accounts [inaudible], and we use EPP to communicate the registrar and register. And the other two main components are the dispute system, and in this part, and that component, we're using web service to communicate with the registry.

And with the arbitration system, the online arbitration system. The same chart, with the client and the operations, we have the client on the bottom side. The client can create, new update, delete and transfer request domain names in the user account in the new system. And in the old system, they only can update and delete.

And transfer concerning, if it is necessary. In the dispute resolution system, we are receiving the domain name compliant. And when the conditions are okay, the system create new arbitration into arbitration system. And in that system to arbitrate the domain holder and the complainant start, and when that process finishes. The arbitration inform using web service, to our dispute resolution system to [inaudible].

…system update in the registry. [Inaudible], in this new dot CL we use Java, like our main programming language for verifications. And we use spring and a start frameworks, other

technologies related to web applications, like [inaudible] CSS… For a skilled tasks, we use Java, PERL, and Python, trying to use the right technology for the right tasks.

If we want to do small tasks, for example to collect information for statistics, probably scripting language… But if we try to implement a big task, we get different system using different models, using Java the framework would be the best option. To have higher availability and performance, we use [inaudible] server [inaudible], and [inaudible]…

…is an important component in the system, we use to communicate the old system with the new system. And Apache still do the web server work, and [inaudible] is our Java web application server. The servers architecture is in this chart, Apache is the first component [inaudible] user request and the Apache sends the request to our three [inaudible] Java web application server, and [Tonka] is used in a proxy to send the transactions, two, three, database work.

These systems in 2015, in November 2015. And the green area chose how many domain names we have entered in the system. And the blue area, shows how many domain names we have… In the green area is increasing our time, and the system that we expect that is decreasing.

In the new system, currently we are, have more than 300,000 domain names. And the old system, we have 108,000 domain names. Why did we do it on our own? It's a very [inaudible] question, but we have to… There is number one, that we checked external tools like [inaudible]… That check was probably six years ago, and we knew that the external tools need a lot of customization.

At this time, we have more than 400,000 domain names with local policies, like [Brute], [inaudible], the local objection systems, the payment in both systems, we have different payment systems. And we are part of a university, and we have internal [inaudible] for different process.

Other reason is the requirement evolve. [Inaudible] we have more than 200,000 lines of code in the new systems. And the most system start also with this line, but our time it was necessary to add the new functionalities, to solve some [inaudible], to and to answer the different [inaudible]. And we have more than 350,000 [inaudible] domain systems. And the last reason is that we have an engineers with experience.

We started with engineers, and all of the time we have an engineer area. And we have developers administrator, a network administrator. And [inaudible] the main reason is to coordinate this project with organizations [inaudible]. And it's

completely necessary to have the tie to customer service and in other areas, like the legal area, administrative payments.

We start with new policies and new procedures, and what was necessary to coordinate the different areas. Use the right technology for the right tasks, it's when you, when it's necessary to implement like a big system, like a web application with different user roles, of mean tasks, scaled tasks, probably use a framework, it's technology. And we did a pen test, a penetration test, and that test was okay after first time.

We were very proud about it. But after hours, went in the system start, we receive an alert informed security. We were very worried about it. We solved [inaudible] but the pen test failed, and from that it [inaudible] is a big lesson, try to surprise companies and abilities in different system, and probably different technologies. And is very helpful to carry out the stress test and disability test.

With the stress test, you can find how many requests your architectural, geographic applications can manage. And with disability test, you can improve your user interface and time in the system.

EBERHARD LISSE: Thank you very much. No applause? [Applause]

It's an old, something works on the first try, something is wrong. How many attacks do you see on your system? You said bout penetration testing. How many attacks do you see on your systems? Penetration testing.

I know that. I know. They have had a penetration test, which first worked, and then later picked something up. Have you noticed any attacks on your system since then? Any attacks?

JOSE URZUA: Attacks, no. Only one user report the back. One. And that [inaudible] was in the user account where the user can update, and was only affected others and the name of the user. And that was the [inaudible] ID for the subject in the form. [Inaudible] and when the user submit the update, they can change [projects]. I don't know that was the problem.

EBERHARD LISSE: Thank you. Jaques?

JACQUES: Jacques with dot [inaudible]… question. So you're running two… How come your plan didn't include migrating all the old…?

EBERHARD LISSE: Say again.

JACQUES: So why, how come you didn't migrate all of your data from the [inaudible] and just operate one system?

JOSE URZULA: When we, at this time, to move from the old system to the new system, is not mandatory. And we want to add new rules, when the client are coming to renew the domain names in the old system, we want to force to move to the new system. But we want to implement rules in the next time, the next month.

EBERHARD LISSE: There is an easy way, just increase the price, double the price for the old system, they will all switch.

That's correct. We notice this in NA, we have got five or six clients who are refusing to switch no matter how much we increase the price. We are not complaining. Okay.

UNKNOWN SPEAKER: …means you have two zone files?

JOSE URZUA: Our system is here like a registrar, and the new system is a registrar too. File is in the registry. We have only one registry.

EBERHARD LISSE: So only one zone file. So basically the new system is a new front, so users can use two different systems, now it just depends on what strategy is being used to entice them to go to new system.

UNKNOWN SPEAKER: End users, if they have to make a lot of changes…

JOSE URZUA: Change in the…?

EBERHARD LISSE: What's the feedback from the users? What's the feedback? Do they like it? Do they have, do they think…? Do they have to implement many changes?

UNKNOWN SPEAKER: Yeah, technical changes.

EBERHARD LISSE: Technical changes.

JOSE URZUA: Technical change. No, dot CL has a web interface from probably 15 years, and we changed that. And of course, at the beginning, we changed the interface, the policy, the user [inaudible]… Over time, we have a new, some new functionalities and we have update in the interface change.

EBERHARD LISSE: I mean, this is the method that you will use if you want to get new… You have a [inaudible] registrar structure, then if you want to get more registrars in, you set up a system that talks to the registry with EPP, and then you can publish that specification, and then more registrars can…

But if you do it for internal purposes, to make registration safer and so on, and you must force the users from one to the other. And many of us have done this, as I said, the end users, they go, this is basically, at the moment, a monolithic… They had one registrar and now they have two registrants, and the registrars have got resellers.

Eventually, we'll probably start saying okay, this should become registrants, and then they can go straight in with EPP, that's all the future I would foresee from here. That's a policy decision.

JOSE URZUA: The next step to have in this model is to have an external registrar, using and we are working on that.

EBERHARD LISSE: The point, what I'm saying is you must have the technical basis there. [Inaudible] dot [inaudible], we did the same thing. We first got a structure working, we set up administratively so that it was easy when we had the technology working, that we could then just create the technology and now get the administration, so that eventually, the registrars, like in every [inaudible] also will reduce the work load, especially the help desk functions and so on, can be off loaded to the registrar.

So in order to do this, I would think thing here shows that you have to lay the ground work technically before you can do the administrative operational changes. You have been, very often you can go one to one, any time, anyway, and thank you very much. Quite interesting presentation. And the next one will be Stephen Farrell. [Applause]

Usually we have what is called a host presentation. We usually invite the host who hosts the ICANN meeting. Since…

Since we don't really know who is the host this time…

Yeah I'm just making fun of this. Usually, in the past, it used to be the ccTLD manager who was the host. I think this time it's [I

next], the Internet Exchange, and we asked them, they really didn't have a presentation ready. So we asked dot IE, and they had something going. So not only a tradition, but we also have heard [inaudible]… whenever they start.

UNKNOWN SPEAKER: So I guess we engage with Stephen and the IE domain registry. We help work with him in his IETF work as well. And in one of our sessions, we kind of put our heads together and we tried to develop [something] that would increase the usage of DNSSEC and try and just make things a bit easier for people to use.

And so, I guess I'll hand it over to Stephen in terms of the grunt work that he did on that.

STEPHEN FARRELL: Thanks. Okay, so the idea is basically from using various kinds of domains and knowing kind of virtual hosts in various places, it's always too hard to get everything working and get the DNSSEC work, get HTTPS working. So we thought it might be useful to see if we could do it at, initially a proof of concept to just show that it can actually be easy and try to talk to some registrars and other folks who could actually deploy it.

And [inaudible].

So this is the working with dot IE with this, nothing to do with the IETF really or Trinity College. So that's the [inaudible]. So the problem basically is, you know, many years on, not everybody uses HTTPS for the web, and… Getting better, but it's not getting better very quickly, especially maybe for the kind of long [term]. That's kind of the people we're thinking about here, more or less the new long [term] registry. And the downsides are, I think, probably are very well known. That if you have [inaudible] text, you have a larger attack surface, and you, you're just trying to distract me?

Okay. Thank you. So you've got [clear text] you've got a larger attack surface. [Inaudible]… More attacks is more support costs, more trouble, so it's all bad. And getting certificates for domains, and websites, and DNSSEC, set up is too much trouble for the average person creating a domain, or for somebody administrating stuff, or they just don't care.

So one way to try to tackle this is to make it, even if they don't care, it still can [inaudible], which would be nice. So the initial goals were to try and produce a, you know, demonstration that you could have a little checkbox that somebody is registering a domain, particularly if they're going to register a domain with somebody who also is hosting and will host a big forum, that they could just check a checkbox, or even no checkbox, and always have them, and that DNSSEC would be set up and HTTPS

would be set up, so that from the very first DNS query, or from the very first HTTP request, it's all using the [inaudible] protocols that we have, and just have it.

The goal is not, you know, very high assurance, because probably most of these websites are somebody's blog or somebody's initial stab at site for the company or something. So, you know, we're aiming for low insurance. And that's kind of consistent, I guess, with this idea of kind of, idea of opportunistic security were you try to make things a bit more secure, in a way that's easy and reliable, and then later on, it will be more simple to get [better] insurance, compared to starting with nothing.

So it's kind of consistent with a broader approach with what's going on. Those are kind of goals. One would be the benefits, well hopefully, you know, you'd have things like cookie theft and so on, a little bit less often. If you look at the way the browsers are developing and the web is developing, there is a whole bunch of features that are maybe more likely to be more accessible if you're running over HTTPS.

So it's things like having permissions for the microphone, and camera, and web RTC, and so on and so on. I think the browser makers are more and more assuming that they want to add in new features, or sometimes powerful features, and only have those available if you're running over HTTPS.

And as a result, you would like people, as soon as they were setting up their website, to get there immediately, and not have to get a website and then realize web RTC is crap for them. So that's the kind of, you know, the general thrust. We also, you know, get rid of some of these stupid browser warnings, and as a theory, I don't actually know if this is true, but it would be fantastic to find out is, if you just do this all out of the box for nothing, without even, maybe without even telling them [inaudible], would actually that reduce the support cost on people like hosts and registrars?

My guess is that they must get a lot of calls from people who are struggling with the open SSL command line, or other things, or you know, following up some blog somewhere that tells them how to set up HTTPS, and it's now working and they'll call somebody. So I don't know if this is true, it would be interesting to find out if it is. And I think there is a common good in all of this, and even if you only just, if you're only interested in making money, you're about, this is probably not the right room, maybe it's another one somewhere here.

But all of these kind of [inaudible] things, I think the likes of Google and so on say that there are going to rank sights higher if they're running over TLS and so on. So again, if you have more, better security kind of posture, even it's kind of opportunistically done for a low assurance environment, chances are that will

actually score better in whatever kinds of ranking or evaluations people are doing, whether that's directly for marketing reasons or others.

So. What we want to do, we tried to figure out something that looks like it's traceable, so if working with doe IE, they have, the dot [inaudible] is unassigned. So basically, if somebody is coming around to a registrar to create a new domain name in dot IE, and that registrar is going to offer hosting, quite a few of them do. Then you can basically just have this little option that's turned on, or just have this. And in theory, you could have like a more advanced one where you can place in public keys, or do the key generation yourself, command line.

But we haven't really done much there because it's really unclear what would actually be unusable, what actually become less usable than current things, depending on how you approach it. And technically, I suppose we want to try to do is the, you know, have the registrar do all of the DNSSEC stuff, talk to [inaudible] in the right place, and then [inaudible] as an instance of the CA, probably using the IP protocol that's been standardized in the IETF, but that's encrypted. They're at the point of having a service life [inaudible], and so that, you know, again with the goal from the very first DNS query, the very first HTTP request, that it's using whatever protocols we have.

And I think the benefit there, so it kind of seems to work as well, which is a good thing, but there is benefit there in investigating, and there is a bit more work to be done, but investigating how we can leverage DNSSEC for the initial domain registration, in order to get better domain validated certificates in a service like [inaudible] encrypt?

So do people know what [inaudible] encrypt is in this room? [Inaudible] I see nodding heads but not that many of them. [Inaudible] script is a fantastic idea, that a whole bunch of people are doing, offer essentially a free CA service, spitting out the certificates that you need initially for websites, but also later on for other things. And so we're kind of trying to cooperate with them on this and there, they have open source tools and the client just seem to work after some wrangling.

I don't think there is anybody, there is an [EFF] guy here, but I don't think there is anybody from [inaudible] encrypt directly. But if you need, if you want to talk to them, you can ask me and I can point you the right way. So the DNSSEC setup is basically the registrar does all of the key generation, [inaudible] the DS to a registry and you know, in the case of dot IE, they have an API that they're using with their current registrars, that exists but is not turned on in the wilds, not widely used, I think is fair to say.

And the new zone will get signed, and so on. And DNSSEC key rollovers in the newly registered domain are kind of important, so we think we've got that working, but it's one of those things where time would tell. The tools actually seem to do it all correctly, as far as we can see because we're not actually writing any tools, we're just using…

And then for the web server, the assumption here is that the registrar is also the host and then will spin up some kind of website, maybe with some editing tools. In our demo, we haven't actually done anything like that, it just creates a site with one page. But we'll generate also the red server keeper, and then use the [inaudible] encrypt client to talk to encrypt and get a certificate for that.

And again, it's all happens out of the box for free. One of the nice things about the [inaudible] encrypt CA service that they're setting up, is that they also, when you use their client, it will automatically renew the certificate. So one of the other things, if you had a system like this working in the real world, hopefully it will also get rid of that stupid thing where you have forgotten to renew your certificates for the website.

And their model is they're going to have, I think, like a 90 days certificate lifetime, and when you install their tools it will also do the key rollover and… And you just, so the goal is basically, like I

said, you know, you get DNS kind of list, looks at happy at the beginning. SSL labs look [inaudible] as well.

So I think reliability is a kind of critical goal. So the idea is don't make things worse, and I think, you know, key rolling for DNSSEC is probably the most likely thing that you would suspect you would follow on, or the certificate. So I think that this [inaudible] seem to work like this, but you'd have to try it a bit more scale, the [inaudible] of it.

So this is our kind of secret plan, it's no longer secret I guess. But just to work with, we're working with the dot IE folks to develop a proof of concept, and you can have a look at that now. I can't do that live because it's not my computer, but [inaudible]. And then we're going to basically try and talk to registrars and hosts to see if we can get some folks who are interested enough to actually try this.

And one of the interesting things would be, you know, if they offered it as an option, how many people would select it? I think that would be a quite interesting question. You know, it it's a really stupid simple option like, would you like it secure or insecure? [Inaudible]

And actually, they're really, I can't see a reason why there should be a major to change more, [inaudible] non-profit money. So we'd like to basically discuss the details with some folks who

liked try it and then work with them, and try to help them to get this in place. So I think that's part of a broader thing that I am doing with…

EBERHARD LISSE: Exactly, yeah…

STEPHEN FARRELL: And then sort of hopefully it would be kind of all working happily and then other people could pull down the code and have full and [inaudible] status. So we have that thing that I did, because I did it, it's kind of pretty badly done. But you know, so it's chewing gum and string and other things.

But if you go to HTTPS testbed dot IE, and if you don't, if you go to HTTP, you'll land somewhere entirely different. But if you, just because [inaudible]. But HTTPS testbed dot IE has some very highly graphical content there, that pretends to be the dot IE ccTLDs. So you can basically go there and [inaudible] testbed dot IE, [inaudible] testbed dot IE is gone, sorry.

But you can create a subdomain there. And it will go through the process, and about five minutes later, the subdomain will be created and will be DNSSEC signed, and let's encrypt [inaudible] verified with the website, and the website will be there.

So it just …happens. But I don't know if all of you do it at once, I don't know, maybe it will work. So you basically, you can kind of create a kind of [inaudible]… The idea would be to try and talk to a registrar, and then have this as a bigger kind of thing we can do with real dot IE domains.

This implementation available, which is mostly scripting. There is some details in the back up slides, it's all kind of tools that are existing, so that's encrypt Python [inaudible] are the DNSSEC tools, so zone signer and [inaudible]. Soon as I get a chance, I might put it into Github or something that is more socially acceptable these days.

But if you want, you can go there. What does it do inside? It's kind of basically, like I say, it's chewing gum and string. There is a website, it's testbed dot IE, you can fill in the thing you want. There is the request staging, hands off to another machine that will do the key generation for [inaudible]. So that will ping back and forward on most machines for a few minutes, and then when that is done, and then when that's done, it will go back onto the first machine where it will create the web server, [inaudible] and talk to [inaudible] encrypt, and get the search, and put up the website.

And then eventually, at the end, there is a timer that will count down. When it gets to the end of that, it will look like it's broken,

and you can come back later, and it's working again. And the reason is, because one of the authentication modes that's encrypt support is to actually talk to the web server itself, and I have to restart the server.

So it's looks like it's broken for a minute and then it comes back. But it kind of works. Although it's kind of pretty hacky, I think it has the right kind of, different sets of functions split up correctly so you could actually include those in a real deployment. Now you might have to re-implement one or all of them, but it should be split into the right kind of chunks that you can put them in the right part of your work flow.

And finally, I found [inaudible] if it's useful would be something nice to find out when we get to [a more] realistic world. There is a picture, there is another picture, and there is a conclusion. So it looks like, [inaudible] with not too much effort, if somebody wanted to do this, it seems entirely possible to do it. That you can actually just have, if you're setting up a new domain and the hosting for that, you could just do all of this, with not much effort.

I mean, I think it took, you know, [a couple of] evenings, you know, maybe a dozen evenings, I don't know. It took a bit of hacking around. The tools are all there. They're all open source. That doesn't mean they'll fit in every environment. It will be

interesting to find out. I think, you know, invisible or only very slightly visible security should become the norm.

We've really got, if you know, if we could come back in a year or whatever, and say that the number of new dot IE second level domains that are using this kind of stuff and not doing anything bad, not doing lots of clear text instead, it was such a percentage, that would be really good.

And once you kind of get this basic stuff going, I think you could extend it to other things like, you know, more interesting domains who are doing mail, [inaudible] and so on. It looks like registrars who are hosts, in particular, and working with a ccTLD, maybe particularly these smaller ones, are really well positioned to do this kind of stuff. [They're] not much cost and not much pain, if you could do it well.

[Thank you.] That's it. [Applause]

EBERHARD LISSE: Thank you very much. I was, I noticed something about let's [inaudible] dot org, and I just had a look, I didn't know about that. That's… Can you talk a little bit about that?

STEPHEN FARRELL: Sure. I'm not associated with it, other than liking it. Let's encrypt dot org, maybe, is there anybody else in the room who is better to talk? No? Okay, I'll do it. It's a bunch of people who are operating essentially a service that wants to spit out free [inaudible] certificate for the world for that scale. So there is quite a bit of heft behind it.

They have quite a bit of funding. They're rolling out, you know, an Internet scale, whether it be, that will do stuff for free. And one instance of a… So there is activity in the IETF called [Acme], which is an acronym for something. And the idea there is that this is a protocol you can talk to a service like [inaudible] that could be operated by one of the existing CAs.

But now focusing on getting the automation of the certificate lifecycle correct, because in the past when we did this, we screwed it up. The idea that you can, if you're talking the [Acme] protocol to such a service, like let's encrypt, you can get your certificates rolled for free, you can get your, you know, or they can charge you I suppose.

Key rollovers and so on could become not a problem, when you do a, you know, [inaudible] install, Apache, it might be just…

EBERHARD LISSE: Thank you. Warren first and then…

WARREN: …I'm planning on doing HTTPS and automatic redirects…

STEPHEN FARRELL: Yes. So currently what I do is I set up the subdomain on, it's actually another version of host on the same thing, so you will get that, but yeah. I'm trying to follow the better [inaudible] dot org, better config for the browser stuff.

WARREN: …browsers…

STEPHEN FARRELL: It's on the way. So for, to the best of my knowledge, they have the [inaudible]…

I think they have a [inaudible]. Their root will be going into browsers. If you use their client today, [inaudible] CA, but I think that's in the process of happening.

UNKNOWN SPEAKER: [Inaudible]

STEPHEN FARRELL: Which one? I've made so made so many of them.

UNKNOWN SPEAKER: When you said you wanted to improve people using secure… DNSSEC is hard in the beginning, but easy to roll on. [Inaudible] the case for all of your clients when you're a big host, it's just maintaining them that's hard. The DNS industry is a little bit [inaudible] than you eluded to in here. There are parties in it that are not the registrars or registries or registrants. Some of us [inaudible], yeah, all of us is [inaudible] and it would be wonderful for the registries, to tell people how good every provider is, because when you see a bulletin board as a user, whether there is anything behind it [inaudible]…

[What kind of] technical competency is there?

STEPHEN FARRELL: [Inaudible] a statement I didn't make.

EBERHARD LISSE: Personally I think for registry to tell which registrants are good, or should be used or not, is a little bit complicated. I think users may vote with their feet or with their dollars.

UNKNOWN SPEAKER: [Inaudible]. So we just kind of picked one slice, which is as I say, you know, not too big ccTLD, and a registrars, and new domains,

and hosting. I think there is lots of other ways we can cut down the space. And for many of those, there may be like a nice, easy way to get this, you know, better security stuff, rather than to try to do it for everybody at once, if you can kind of cut out the bits and address the easy bits and then, you know, move on from there.

UNKNOWN SPEAKER: …protocol in place that allows third party operators to [inaudible] the registries. I [inaudible] DS records.

UNKNOWN SPEAKER: So neither let's encrypt dot org or…

STEVEN FARRELL: Yeah, we're talking to them about how to do that best. There are probably some subtleties. So at one level, it should be a no brainer. You should get better domain validated certs, if the DNSSEC is in place. But, you know, that just has to work logically. Now the details of how to make it work well with a service like [let's encrypt], so some of those details we need to figure…

They don't currently support it, that's true, we're talking to them about how they could integrate it as one of their validation tests, actually.

UNKNOWN SPEAKER: …no longer will be necessary and we will use TLSA to verify….

STEVEN FARRELL: So who knows what the long run will be? But I can tell you that browsers won't do now so it doesn't really matter what RSSAC, nothing for browser, for things like SMTP over TLS between web servers.

UNKNOWN SPEAKER: …certificates, so it would be more interesting…

EBERHARD LISSE: Rick Lamb.

RICK LAMB: Hi. Rick Lamb, ICANN and Fellow, you know, bound and string builder, very much so. Thank you very much for this kind of presentation.

EBERHARD LISSE: Not bound and string, toaster oven.

RICK LAMB: Toaster oven, okay…

UNKNOWN SPEAKER: A little more sophisticated.

RICK LAMB: All right. This is exactly, one click, one check kind of system. [I] find this really interesting, but I think we all know that there is this kind of [inaudible] and figuratively, right? With things getting into those, so I know you're not the let's encrypt guy, but like my question is, are they or anyone else also looking to generate S-Mine certificates, web certificates, or maybe that's already there, right?

So the idea is, of course, secure an email. I can do the same thing and go, my account, like a complete idiot, I wanted an email account and I got a webpage, I get an email account [inaudible]. I know…

STEPHEN FARRELL: So Internet security for email is, you know, another day's work, but I think the, I think it is true that at least the let's encrypt folks

do have the intention to broaden just beyond web services. When and how that will happen, we'll see.

It's kind of, so I don't know their timing yeah. I can't answer for them.

WES: Again, thanks very much for doing this Stephen.

EBERHARD LISSE: Can you identify yourself for the record please?

WES: I did, but I'll do it again. I'm Wes [inaudible] from [Persons?]. You know, we need to get rid of [inaudible] and things like that are actually a leap of faith. And doing stuff from day one is absolutely the right direction. DNSSEC advocate, one thing [inaudible] let's encrypt folk to get them to sign their zone.

STEPHEN FARRELL: Sure, I will.

EBERHARD LISSE: Okay. Thank you very much. And give him another minute. [Applause]

Wait, it's a remote presentation. Remote question?

UNKNOWN SPEAKER: We have two remote questions. Shall I? The first one is from John [inaudible], he's asking, "Most new websites tend to be just [inaudible] ware and are handled by contractors and not regularly [updated]. Is there some kind of outreach plan for these contractors or developers? Or directly to hosting service?"

STEPHEN FARRELL: Sorry. There isn't a plan at the moment, but it's definitely something that we want to do is talk to people that are creating [inaudible] how they work and see if we can get involved [inaudible] as well, with this same…

UNKNOWN SPEAKER: [Inaudible], if we can successfully find some people who did the hosting, I think that they already have those relationships, so I think that should…

UNKNOWN SPEAKER: And the second question is from Dev he says, "What about service providers that rely on HTTP at the end point or proxy servers, to improve performance for their users? Not all content must be…"

UNKNOWN SPEAKER: That's also a fine political statement, I would guess. Regardless of what one thinks of the pros and cons of it, the web is going more and more towards HTTPS, and we need to live with that and we need to deal with it. You know, I think this is a case of just dealing with the weather, than we can't really push back on it. If you look at the figures, I think the [inaudible] are saying more than 50% that the TLS traffic that they see is… But that landscape is changing.

EBERHARD LISSE: All right. Thank you very much. And then Martin and Maarten could come in front, because their presentation is interlinked.

MAARTEN WULLINK: Good morning everyone. Sorry for the delay. My name is Maarten Wullink. I work for SIDN. And today I would like to use this opportunity to talk a little bit about our work with DNS big data analytics. Let's try again.

A little bit about SIDN for those of you who are not familiar with SIDN. We are the registry for the dot NL ccTLD. With 5.6 million domain names, dot NL is currently, I believe, the seventh largest

TLD, and [inaudible] that almost two and a half million of those domain names have been secured with DNSSEC, making it the largest DNSSEC deployment in the world.

I'm also a member of SIDN labs, which is the R&D team of SIDN. So when we look at data at SIDN, we particularly look at DNS data, and what we see is that we roughly see [somewhere] in the neighborhood of 3.1 million resolvers each month, sending us around 1.3 billion queries each day.

And if we were to store all of this data in the compressed [inaudible] format, that would be around 300 gigabytes of data each day. Now in order to be able to capture and analyze this data, we developed [inaudible], which is our big data platform. It's goal is to help us to be able to improve the security and stability of dot NL and Internet at large, through data driven method. The problem we faced is that, that existing solutions for analyzing network data do not work very well with large data sets, and tools also have limited analytical capabilities. Our main requirement was to have a high performance, near real time solution.

And the approach we took was to skip the [inaudible] files, because they are very expensive to analyze. [Inaudible] files to optimize format, and use tools and query engines that can leverage the advantage of the format. So what are use cases?

Well we are focused mainly on increasing the security and stability of dot NL, so we look at visualizing DNS patterns, detecting bot net defections, real time [inaudible] detection, generating statistics.

[We had a] nice page called stats dot SDN labs dot NL, where a lot of [inaudible] NL zone are presented. We do scientific research in collaboration with Dutch universities. And we also use it for operational support for our DNS operators.

I would like to quickly show you two example applications we have developed. The first one is what we call a DNS security scoreboard. It's goal is to visualize DNS [inaudible] by combining external data feeds, which we get from fish tank and net graph, and combine this with our own DNS data.

The architecture looks like this. There is an event analyzer that gets feeds from external [inaudible] such as net graph. These feeds are enriched with our DNS data, which we get from our new platform. Enriched data is safe in the database and it's available for analysis through a web interface.

This screen shot, I don't know if it's visible, but it's an example of a domain name that is effective with a [inaudible]. If you look at the top chart, you see that the bars are a number of queries received by domain name. And the red bar is the day when the phish was reported to us. And it's clearly visible that in the days

leading up to the red bar, there is already an increase in the number of queries received per day, and the number of unique IP addresses, and a number of unique country hues.

So there is quite a visible difference there. And below that, there is also an increase in the number of unique networks that are used in sending these queries. And at the bottom, there is also a distinct difference in the geographical [distribution] of the resolver of these countries. All of these are indicators that could be used in developing a system for automated detection of phishes or other malicious…

Another application we've developed is what we call resolver application, and it's goal is to [detect] malicious activity by assigning reputation scores to resolvers. [It does this] by fingerprinting resolver behavior. [Inaudible] that… [Inaudible] works like this that you have benign clients and malicious clients, and [inaudible] clients can use either shared resolvers, like ISP resolvers, or they can use built in resolvers.

Either way, their queries for dot NL names end up at our authoritative name service, and there we try to analyze the data, [inaudible] or bot net such as the [cartwheel] bot net, or identify DNS amplification attacks. Architecture looks like this, data is our DNS infrastructure. It's being transmitted to our other platform where it's stored for analysis. The resolver reputation

service is built on top of that platform, smart algorithms to identify a [malicious threat].

If bot infection is found, it's being, it will be passed along to abuse hub. Well, many of you probably don't know what abuse hub, but abuse hub is a platform of the abuse information initiative by SIDN and the largest Dutch ISPs. The function is to collect abuse reports for these use desks of the corresponding ISPs.

Though this whole chain, from collecting the query data and license data, and sending the abuse report to abuse hub, and abuse hub will send it to the abuse desk, will help in trying to clean up… These two applications are built with our [inaudible] platform. This slide shows the high level architecture of this platform, don't really mention it.

The goal is to develop applications and services to further increase the security and stability of our [inaudible] zone, the DNS as a whole. Well the main components are [inaudible] services, platform and data sources, a privacy framework. And the platform together with the privacy framework are what we called the [inaudible] plumbing.

In order to comply with the strict Dutch data protection act, we have developed a privacy framework for the DNS big data, which rates the legal, data, technical, and architectural aspects of

privacy management. This is required because the resolver IP address can be considered an identifier.

So every new application development project that starts SIDN and that wants to make use of data stored in [inaudible], create what we call a privacy policy. This policy has to be approved by a privacy board, and this policy contains elements such as the purpose of the policy data the application wants to use, the type of filters that are used to, by application. [Inaudible] how long the data is used, and of course, type of research, application, or is the production application is shared on…

And the privacy policy is approved, and it can be implemented by the [inaudible], I'm sorry, by the [inaudible] component. The component is what we call policy enforcement points, where specific policy can be implemented. We have described all of this in a paper which is available on the website for download.

The architecture of [inaudible], as you can see, we use quite a lot of open sourced components and added some of our own components on top of this. The blue layer contains [inaudible] which is a SQL compatible performance query engine. We also used a, how to distribute a file system, which is a file system which is distributed across the cluster.

And we used the [inaudible] data format to store data on, head of file system, data with [inaudible]. On top of that we have our

own components, such as a work flow component to get data from the name server, following to the database, a DNS library for the decoding the DNS data. Version library for converting the [inaudible] data into [inaudible] data, and services and applications.

Work flow, you can see that data is collected on our name servers, and it's being transmitted to our staging area, functions as a short term storage for raw data. The data is being dubbed by a [inaudible] decoder, which decodes it into memory, tries to join a DNS, with the DNS response from our authoritative name server. Some filtering is done to remove data we're not interested in. An enrichment step is being performed at metadata about the resolver IP address. We add things like geographical location of the IP address, the ASN belonging to the, the autonomous system number of the IP address.

All the while we collect metrics which are being sent back to our monitoring software, and as a last step, we important a form of [inaudible] files in our environment. When this step is done, it takes about 10 minutes from name server to [inaudible]. Data is immediately available to be queried by an analyst or by applications built on top of [inaudible].

So, one slide about performance. This is an example query, which tries to count all of the DNS queries we received from

resolver issues in the IP version four protocol. We did this for a data set for one day, one month and a year. And with a single thread, and with 10 threads. And we look at it, and the skill, the vertical axis is the [inaudible] time, is in minutes, so if you look at the dates, it doesn't even show up, it's only a couple of seconds.

For a month, if you look at 10 threads, it's somewhere in the neighborhood of a minute, and for a whole year of data, 10 threads will be somewhere in the neighborhood of under seven minutes. And one year of data here is similar to almost 52 terabytes of [inaudible] files, which will be, I guess, somewhere around 50 or 60 billion DNS queries and responses.

And these figures are from our research cluster, which is a really tiny cluster. It's only four nodes. So we plan to, when we add more data, to also scale our system. The current status is that we have two name servers hooked up. Each day we add 320 million queries with their answers. And we process about 70 gigabytes of [inaudible] volume, which is transformed into 14 gigabytes of [inaudible].

The database contains about 18 months of DNS traffic, which is about, well more than 74 billion queries and their answers. Volume is about three terabytes. And because [inaudible] replication for three ways, as a redundancy safety, one of the

data blocks fails and one of the servers [inaudible] servers, to avoid data loss.

And we estimate that our loss has a capacity of somewhere in the neighborhood of 150 billion to 200 billion query and response. So our conclusion is that the combination, or technical solution which combines [inaudible] with [inaudible] and [inaudible], is a very good combination, if you want to do DNS analytics. We made some contributions.

We formed interesting research at SIDN Labs combined with universities. We identified several malicious domain names, and bot nets, created the external data feed. The abuse hub of information exchange, which helps leaning up the infections. And we gained a lot more insight in the DNS query data which we receive on our servers.

Work is aimed at combining data from part of name servers we are running with scans of complete dot NL zone, active with passive data. And also getting more data from our name servers and resolvers, and expand our open data [inaudible]. I don't know if there are any questions right now?

EBERHARD LISSE: All right. Thank you very much. [Applause]

Hang on second for a second.

UNKNOWN SPEAKER: We actually learned from dot NZ.

EBERHARD LISSE: Now while we're waiting…

UNKNOWN SPEAKER: …the tools that you've built here [inaudible]…

MAARTEN WULLNIK: The components we built on top of the open source components, I don't know. I can't say at this moment.

EBERHARD LISSE: Christian is sitting over there. Ask him.

Is it a remote question?

UNKNOWN SPEAKER: John [inaudible] is wondering, "Do you use it to check the health of the TLD?"

MAARTEN WULNIK: Yes, well our DNS operators can use this tool to identify anomalies in DNS traffic, and zoom in on particular queries that are causing the anomalies. So yeah, I guess that's a yes.

UNKNOWN SPEAKER: [Inaudible]

MAARTEN WULNIK: Yeah, that was an example query, just counting the number of queries we receive on IPv4. We also store the queries that we get on [IPv6]. You can do the same query and replace the four with the six, but we have a lot more...

MACIEJ KORCZYNSKI: Hi everyone. My name is Maciej Korczynski. I'm a researcher at the [inaudible] University of Technology. And I will present a project in collaboration with SIDN and National Cyber Security Center in the Netherlands. And the title is repetition metrics designed to improve intermediary incentives for security of TLDs.

So what we do is we collect a large and reliable data sets of different security incidents. Mostly blacklists of spam, common control of phishing domains. We also make our own analysis and act measurements. On the other hand, we try to structure

complex DNS ecosystem composed of different players like registrars, hosting providers, their resellers, also their equipment, outdated name servers and so on.

And one of the goals of the project is to create reputation or security metrics for different types of players and different players in general. But we explicitly distinguish those metrics from measuring the performance that security, their performance because it's driven by multiple factors and not only by the performance one type or one particular…

So I will very briefly discuss the types of security metrics, and then an example of security metrics for TLDs, for hosting providers, and I will spend some time on the discussion. So, we propose different types of security metrics at three different layers, I would [inaudible] obstruction.

The top layer is related to security metrics for top level domains. So we compare entire TLDs like dot NL with dot com and so on and so on. And here examples of such metrics, could be the number of bot net contaminations, from TLD or phishing domains, phishing attacks per TLD, or their up times, for example.

The second layer is a refinement of the first one, and here we propose security metrics for market players related to TLDs, and those are a registrars hosting providers, DNS service providers.

Of course, those roles can be played by the same actors. And finally, the third layer, we propose metrics related to network resources managed by each of the players.

For example, we check the number of open resolvers per AS or number of misconfigured or misbehaving alternative name servers and so on. So some examples for security metrics for TLDs. Our first group of metrics is related to the concentration of malicious content.

And here the most straightforward is the number, of course, of the unique domains, but then we can imagine that under one second level domain, there can be a registered hundreds or even thousands of fully qualified domains used in common control, both in common and control communication, or for example, the phishing attacks.

In addition, after also talking with the Dutch police and analyzing the data on child pornography abuse, we noticed that there are a lot of URLs under certain fully qualified domains, containing or pointing to malicious content. That's why also we decided to take this one into account. And of course, size matters here. That's why we cannot, of course, compare directly to the number of, for example, phishing domains in the [inaudible] space, and the dot com space, that's why we

normalize, in this case, in case of TLDs, we normalize number of abuse by the number of domains in the registry.

So here is one of examples. Here on X axis we have a number of domains in registry. On Y axis, we have number of phishing domains. Each dot here corresponds to a TLD, red ones are gTLDs, blue ones are ccTLDs. And here, the first thing that we can observe is of course, with any number, higher number of domains and registry we observe, of course, higher number of phishing domains.

It comes from 2014, and we can see that now the other [inaudible] were less interested in new TLDs. Now we observe a little bit different trend, and the attackers are more interested in both compromising, mostly compromising new gTLDs. So another type of security metric is of course, related to [inaudible] of both maliciously registered compromise domains.

So here on, I'm not sure how visible, but here on X axis, we have the TLD ID. On Y axis we have a mean in days. And each bar corresponds to a different TLD. And here that we can see that for example, the one TLD on the very left, to mitigate, for example, black listed domain, it takes a little bit more than seven days on average.

But of course, we make some more, much more analysis based on medium, and it's also survival analysis, and it's clear also

here that sometimes mean is simply driven by few or several incidents that take much, much longer to mitigate. So now, I will discuss some security metrics for hosting providers. And we made a case study for the Dutch market.

And here, very, very briefly, what we did, we aggregated badness per AS, and concentrations of mostly compromised domains per AS, now we do it also per IP space. And we normalized it in three different ways. First of course, by advertised IP space, but also thanks to the axis to the DNS debrief on [inaudible] security, we could normalize it by the number of IPs used for hosting, and also the number of domains hosted within a certain AS. Why is this so important?

Because we can imagine that two hosting providers, one with just a few IPs and having a shared hosting, and we can see there are hundreds or even thousands of domains per IP, and on the other hand, we can have a hosting provider with a lot of IPs, but for example, their main business could be a dedicated hosting. So the number of domains per IP would be much, much lower.

So after whole aggregations, we received the overall ranking, and this ranking that, and we identified AS with consistently high concentration of badness. And actually those results were used in another project also, in Netherlands, called Clean Netherlands.

And here, the Dutch police, public prosecutor and also authority for customers and networks was involved. And the goal was to enhance the self-cleaning ability of the Dutch market. And they were, we were searching for some driving factors there, and they were contacting hosting providers with the high concentrations to promote best practices, but also if needed, to pressure a little bit rotten apples, guys that are willingly facilitating some criminal action.

Of course, before this project, we had lots of discussions with different hosting provides in the Netherlands, and we agreed that the metrics are robust and it's a good way to, a good approach. So also, I mentioned that we make, we have also this third layer, and we dare make a lot of measurements. Now I'm not presenting those results, but here we also use Zonemaster, presented by Patrik a few presentations before.

So, discussion. So I encourage you, if you would like to know also how your TLD is doing, if you would like to compare it against the market, then definitely drop me a line. I can, possible by some driving factors so why the attackers are potentially more interested in your domains.

For example, we observe that some TLDs are very high, just because of some single attacks, but we coordinate also attacks, and we could observe that, for example, in one case, 13 second

level domains were responsible for thousands of, fully qualified domains used in coordinated phishing attack at some point.

Other driving factors, we see that it's driven also by price. So how those also results can be used by you, if you for example, introduced some policy, or you were willing to introduce some policy, and that you think that might improve or change the abuse, or concentration, or up times, then you could upload that in the trends that we are preparing.

Also let us know about your policy changes. So for smaller TLDs, as I mentioned before, we normalize concentrations of abuse per day, per size. So of course, you can imagine if that TLDs is smaller than even five or like even one incident in observation period, can result in a very high position in the rankings. So we need to also provide appropriate interpretation of the results.

Also I guess for smaller entities, I was thinking that it might make sense to put some abuse, or one of abuse software… So in the Netherlands, in general, we observe a lot of interesting initiatives, and abuses handling software that is released for free. And those, I guess, could be also used for, by smaller TLDs. And it works simply, there are certain entities that provide, that provide black lists, then they are automatically created tickets, and they are sent to, probably to clients.

So TLDs act here as intermediaries. So we have limited access to domain WHOIS, and also to some resources that's on black list, like to shuttle server reports. US registries, you have access to those, so if you're willing to share some databases then that would appreciate a lot.

Also, the same thing that we did for hosting providers, we were thinking to do for registrars, but to do so, we would also would like to invite some registries to eventually give us the snapshot and the mapping between domains and the registrars, and then we could create some corresponding, of course, metrics. Identify registrars with high concentrations of badness, or particularly, I don't know, outdated software, and so on and so on.

And then we could think about some case, control case studies, and implement some incentives that potentially abusive… So I guess, that's it. I will, and we would like to [hear] feedback from you.

EBERHARD LISSE: Thank you very much. [Applause]

Jay.

JAY: So I have just… Not free, but it's still very cheap. Very interested to know whether it's 20 cents… Get that figure, that would be…

MACIEJ KORCZYNSKI: So we have the problem because we do not have very detailed access, accessed to very detailed pricing. But if you would… For example, we could definitely check it. Sorry? I see, yeah.

Yeah, we could find a way, I guess. Let's maybe discuss afterwards, we could.

EBERHARD LISSE: I was just going to, I think, when was it? On Sunday, somebody said that they had never seen a dot NA address being used for automated spam. This is clearly because we are more expensive than that, I think that's the easiest thing. The easiest thing, if you have a large, you don't have to lower the price, I'm sorry, that much that it becomes attractive.

Just so if you notice the threshold, what's to come unattractive for bad guys, that's probably good to know. Any remote participants? Any other questions? Okay, thank you very much. [Applause]

So I don't see, I see now. Could interrupt myself in mid-sentence. Next one is David Conrad. I would have moved Roy's slot anyway.

DAVID CONRAD: I'm doing fine. How are you doing sir?

Ma'am, no. Okay. Hello everyone. I'm David Conrad, ICANN's CTO. I've now been in the job for all of a little over a year. It has a been a fascinating and entirely non-boring exercise. So I was asked to talk about the ICANN's new technology office, my office, the office of the CTO.

We call it oct-to internally. There is apparently, I'm told, a children' cartoon, the Octonauts, and I've beaten people when they've [inaudible] as Octonauts. So don't do it.

Yeah, wrong crowd to tell that to, huh?

I never will. So, what is the mission of this group? So I was actually going down to, I don't know, going through the Singapore office and told the folks down there that I was going to be there. And they said, "Lovely, you can give a presentation on the office and you can talk about all the mission and the goals, and all of that stuff."

So that gave me a reason to actually write up a mission on the airplane flying down. So the mission in this has been vetted by my illustrious staff. So it might actually bear some resemblance to reality. The mission of the oct-o group is to constantly improve knowledge about the identifiers ICANN helps coordinate, to disseminate this information to the Internet community, to improve the technical operation of the Internet system of the unique identifiers, and to improve ICANN's technological stature.

So what does all of that mean? Other than the management speak that you see there. So basically what we do is we'll research issues related to the Internet system of unique identifiers, as you should be aware, domain names, IP addresses, AS numbers, and protocol parameters. We also support improving the security, stability and resiliency of those identifiers.

And many of you are probably aware of [inaudible] group, John is ICANN's chief security, stability and resiliency officer. John now reports to me within the oct-o group, and his team is been challenged to improve the security, stability and resiliency of all of the identifiers that we have some role in.

And we have a bunch of initiatives in that space. We provide internal and external Internet technology resources. And that's basically, hello. Stop that. Oh, I thought I did it.

Training, we collect data, provide research, provide informational resources and consulting, both internally and externally. So you know, internally within ICANN, you know, ICANN is now, staff is about 320 or so. Perhaps unsurprisingly, not all of those people know how to spell DNS, so one of the things that we do is try to help people understand what it is that resources, coordinate externally.

We provide trainings, in particular our SSR group goes out and provides training to the anti-abuse communities, law enforcement and those folks. We also provide presentations on various aspects of [technology]. So the team, currently… So when I started back at ICANN, I had the vast resources of myself, which aren't very vast, but since that time, the team has grown to a [total] of 11 people.

Myself, obviously as CTO. Roy Aarons, made the tragic mistake of working for me again. [Inaudible], Paul Hoffman who, I'm not sure. Are any of these folks actually here other than Roy? Probably recovering from last night.

And John's team, Carlos Alavarez, [inaudible], Richard Lamb, Rick is here. There he is, hiding in the back. And Dave Piscitello.

And I actually even have an executive's assistant, Kathy Peters rules my life now. So these are the major projects that we're currently undertaking within the group.

The ones that are consuming my time, primarily, are related to the implementation of the transition, you might have heard something about that happening around here. The Board advice registry, I think is actually one of those sort of hazing rituals. When you rejoin ICANN, you get assigned a bizarre project, and the Board advice registry appears to [be] that one for me.

It's basically just recording all of the requests that go to the Board. And tracking their progress. Why that's sitting in my group is one of those questions that will haunt all time. The SSR review team recommendations, some of you may know we have a few reviews that occur now and then. One of those is the SSR [review].

They came out with a set of recommendations, I believe eight recommendations, and I'm responsible for seeing that they're… Some of you might have heard that we are considering planning to roll the root KSK. Anyone here think that is a good idea? Raise your hand.

Anyone thing it's a [bad] idea? Raise your hand. Okay, interesting. So we're coming up with a plan. The plan doesn't

mean, you know, there is any set date on implementing it, but it is currently, we just finished the last, the public comment on some recommendations. We're in the process of reviewing those recommendations, the public comments that will be applied back and recommendations.

Recommendations will then be provided to the root management partners, and will be turned into the actual roll over plan… At some point, we might figure out when we will actually be doing, executing the plan, but that is currently…

Paul Hoffman has decided to actually… Everybody knows middle boxes are broken, so Paul is actually going to go and see that middle boxes are actually broken. Roy is doing bizarre and interesting things to root server data, or at least wants to, and related forensics. [Inaudible] is also helping out in the transition, and looking at Internet health indicators.

The last presentation, I suspect in [inaudible] coming by and harassing people for information. One of the more interesting aspects of what [inaudible] does is an IP for an address market research, just seeing what the implications of [inaudible] IPv4 free pool are, and what it means to the real world.

He's also has the joy of looking at the digital architecture that's the handle system he's just researching. Rick is continuing to do DNSSEC awareness, adoption, deployment, and training,

wondering all over the planet, helping people deploy DNSSEC. And the rest of the SSR team are doing sort of the global security engagement, law enforcement, and anti-abuse folks, to try to help reduce the amount of abuse that occurs within the Internet system of unique identifiers, primarily domain names but we have had a [inaudible] as well.

And Steve Conte is actually coordinating all of that training. These are the ongoing activities, I won't bother to read all of that. The high level is we do a lot of SSR support. We do technical resource support for a lot of internal things, AOC reviews, there are a few of those as I mentioned, we supported [inaudible] ACs and SOs. ICANN's Board, provided a number of briefings to the Board on technology issues. Because, you know, the ICANN Board had a better understanding of the technologies.

Doing a lot of data analysis. We support, if any of you folks here are registries, and you're asking for an enhanced service through the gTLD program, you get [inaudible] and my team is called in to evaluate the [inaudible] of those [R-seps]. Also I get to review every document that ICANN has publishing, at least in theory, for technical accuracy.

And as some of you might know, ICANN does tend to publish a lot. So I don't review all of the documents, I have failed to. I

apologize. So summarizing, you know, basically the office of the CTO here is primarily intended [to help] folks, and we actually mean it. We really do want to help folks internally, within ICANN, we provide support, answers, consulting, or we've even instituted sort of brown bag informational lunch, informal sessions where, during lunch, we explain new and interesting aspects of Internet technologies.

I think Roy did the first one. You know, you happen to be in the ICANN office, we would be, we'd love to have you participate and attend this. Externally, we do the same sort of thing as we do for internal, we provide support, answers, technologies. But the main goal is to try to improve the state of the art, or the Internet system of unique identifiers as much as we're able given our limited technical mission.

And with that, I will [inaudible] before lunch. [Applaud]

EBERHARD LISSE: It's copyrighted though, by the BBC. Anyway, any questions? All right, thank you very much. Then we'll commence at 2:00 and Roy will try to keep us away from the post-meal stupor.

Okay, you can all take your seats. Of course, I must punish the ones that are on time because of the ones that come late. Please find your seats relatively quickly. And Roy is going to

start with a pre-DNS naming something about the history of names in the DNS, in the early days.

ROY ARENDS: Thank you everyone for being on time after lunch. For those of you who are on time. So this is your after dinner dip, so if you fall asleep, I won't blame you. This is, I made a mistake. The initial title of this presentation was early domain names, but it needs to be early host names.

This is about the naming systems before the DNS was live. So my name is… So my name is Roy Arends. I work for ICANN. I'm a principle research scientist in the office of the CTO. And I'll focus on DNS, DNSSAC, anti-abuse and statistics. And so what is this about? This is about finding the very first host name on the early Internet. I got this question when I was giving a lecture in England, in Cambridge University.

And a student came up to me and asked, "What was the very first host name?" And like every one of you, right, who knows this, "I said [inaudible] dot com." They said, "No, no. Long before, because people must have been given names to systems long before the DNS was there."

And I found it a very cool question. So I started to focus on this early maps and these early stories that people will tell. There is

a lot of documentation on the Internet about this. But to be fair, this is not an academic talk, this is not a scientific thing. This not an in-depth research project. That's kind of two out of the three terms out of my job title already out of the window.

So this is digital archelogy at most. I'm also not looking at the early OSI nets, SNA, or Darknet, this is about the early ARPANET. So in order to do that, you need to start by looking at a very, very first picture. So the birth of the very early ARPANET, the [inaudible] connected to a host. This was in UCLA, September 1969. This is an [imp] created by Bolt, Beranek and Newman.

BBN in 1968 in December got a contract to build these things. They were supposed to build four of them, they were to deliver them to four universities during 1969. And the very first one, site one, was delivered to the UCLA, the network measurement center. And it was this box.

This was a scientific data systems, this was a six by seven. To be fair, I wasn't born then, so this is all heresy. The operating system was actually called sex, S-E-X. I think this one of the very first time that I can call out that name without getting in trouble. The alternative name for the site is actually sex as well. Sex stood for Six by seven experiments.

Even in the manual, I understand, at the time, sex user manual. Can you imagine the disappointment after opening the book for

the first [time]? Anyway, so back to the first slides. As you know, one swallow does not make a swimmer, just as one node doesn't make a network. So the Internet, the ARPANET exploded really quickly into two nodes.

This is also the first ASCII [inaudible] I find. This is nine pipe symbols. Sorry, guess it makes a pipeline between the UCLA and the SRI. SRI stand for research institute. I'll use these names later on a little bit more. These are the four nodes that I was talking about. BDN deployed these system at these four universities, UCLA, SRI, UCSV, and UTA.

This is 1971, April 1971. The reason why I include this graph is because of the next graph. You can see only five months, only four months in between, this is August, and already a lot more [imps] were deployed. The way both Beranek and Newman worked was basically to deliver these large [imps] at the site, and every time they connected to a site, they basically noted down some administrator, who to contact, basically your early WHOIS data, if that makes any sense.

And here is the RFC 235 that talks about this. There is also evidence that the scanning was an item early on, on the ARPANET, because both Beranek and Newman would log in every day to make sure that your site was up or not. And every two weeks, it would result in a document similar to this. And the

cool thing about this is the terminology used, right? Status or prediction. Some of them were life, the server, or user only.

Some of them would be life soon, or October 11 [inaudible] or November. So this was really the first host table without any names, if that makes any sense. John Postel, in that same week, if you look at the dates, but of course this is all in hindsight, around the same time, proposed actually the very, very first naming convention. This was basically a four character or three character site, an [inaudible] dash, and an [inaudible] node name.

In this case, UCLA [inaudible], slide one, with the alternate name sex. This is, by my definition, I'm not sure how correct it is, actually the very first host name and the very first node, on the ARPANET, UCLA dash MC. So if you ever do a pop quiz on the Internet, that has to do with these kind of questions, you now know the very first…

Back to [symbolics] dot com, remember that I said that that was the very first registered name. That's actually true, the very first registered name, but not the very first domain name on the Internet. That was [inaudible] dot net. And the reason it was the very first domain name, was because nic dot [inaudible] dot net was one of the early root servers, so that predates [symbolics] dot com.

Sorry for the side node, I forgot to tell you that early on. Again with the maps, this is to say, a geographical map. [Inaudible] this is 19872, this is also the connection to Hawaii. So let's say the first, can you say the Trans-Pacific connection? At least out of US mainland.

1973, this is London connected through Norway, Oslo to SDAC. I thought were significant because I now happen to live in the UK. And you can really say in the UK, it was the first international computer being connected, if [inaudible] not being a computer. The did statistics as well at the time. They looked at the amount of packets being released by the system.

And the average daily inter-node, I'll get in a moment what inter-node is, was about 60K per day on average. The difference between inter-node and intra-node, inter-node would go off the wider ARPANET, intra-node was only packets shared with the other two node you see here from UCLA. This is interesting as well. About this 1974, but actually the byline is in German, I included this for Dr. Eberhard Lisse.

This is 1974. It's basically the same picture, it's a little bit increased amount of nodes. Also, at that time, the very first scaling problems. And 40 sites were connected, and basically everyone was doing their own thing. You would call their systems, whatever they like, and they were duplicates and Peter

[inaudible] wrote this, and I love the passive aggressiveness, and the frustration in that one sentence, which you see there as one sentence over five lines.

And I'm going to read that for you just because it's fun. Now that we finally have an official list of host names, it seems about time to put an end to the absurd situation where each site on the network must maintain a different, generally out of date, host list for the use of its own operating system or user programs. And in that same document, this is RC 606, he actually proposes a new naming scheme as well.

And it went something like this. So you had Bolt, Beranek, and Newman, remember they were rolling out these [imps] and they knew who was connected to what, so they would create a host table and send this to SRI NIC, SRI NIC would get updates to names about these hosts as well. Basically during the day, three, and eventually people could download that host and the text files through FTP and later through a naming service, port 101, yes.

If you ever scan the Internet, I would love to see what server was running through port 101. And SRI NIC would then basically look at, if the context were correct, if the names were unique, if it meant that the network guidelines, and also it kept a coordinated list of technical liaisons, and here is such a list.

We have, this is by alphabetical order, this is one of the very early lists, and we see a familiar name here I think, it's Steve Crocker, with a telephone number. So anyone with a mobile phone, have a try. I'm not sure you'll connect to him now, but this is what? 40 years ago. The reason I include this…. I'm sorry? Well try it anyway.

So this is, yeah, the reason I include this is because this image, this on the right, this is Steve Crocker, not in 1974, but 1966, that's the guy on the right. Who knows who the guy on the left is? I know you know, I know Russ knows. So okay. This is [Vince Serf?]. To me what I really like, this is 1966, long before the first packet every flowed on the ARPANET, long before the contract to Bolt, Beranek and Newman went out, even long before ARPA had the idea to do this networking stuff.

We see here two ICANN chairs in one picture. Anyway. Also interesting 1974, this was a previous slide, this was also the invention, or basically, publication of [inaudible] TCP document. It then still took nine years before the old NCP was, there was a flag day at the time. I was playing with Legos at the time, so I wasn't really there.

1977, another topological map. The reason I include this one is because it had the first lawyer speak. It really says here below, please note that while this map shows the most, sorry, the host

population of the network, according to the best information obtainable, no claim can be made for its accuracy. 1977, people. Okay.

Some more measurements. This was from a document named the first 10 years of the ARPANET by BBN. Started December '69 until July '77. This is per month, and we see here interesting statistics, intra-nodes, inter nodes, this is all NCP. Right? Before '71, there was the host to host protocol, after '71 was NCP. I mean, all of this stuff, the reason I'm telling you this because I just figured this stuff out and I'm supposed to know better.

So NCP, eight bits, so maximum of 256 since it started with one, they only had 255 addresses. But I think eventually they only hold out 58 or 64 addresses, sorry, by July '77. Continuing, this is 1982.

Okay so 1982. So this is… Remember the network I showed previously? That's now that small circle in the middle, ARPANET 10. And the reason I'm including this, this is February of '82, this was pre-flag day. Flag day, January '83. This show both IPP four numbers and the NCP numbers. So if you ever wondered why 10.0.0.0, I'm sorry, 10 slash eight is a private space, it was actually the original APRANET.

I didn't know that. Anyway from 1982 to 1985, this is a slightly different picture. More links basically. But also between 1982,

grumbling started again. This stuff is difficult and to maintain all of these naming stuff, and at that time 1983, more than 5500 sites, with a centralized flat ASCII text.

The problem was this thing was so large, that small systems couldn't take it into memory. And if you couldn't take it into memory, it had no file. So people were not just updating them. So they use the historic information. Also, the maintenance of a single host will become cumbersome. You can imagine if every one of those sites, they had changes, new administrative here, new sites deployed, renumbering, etc. to do all of that in a single host table, just didn't scale.

And my next slide was supposed to be a DNS slide. I didn't include that because this talk was about pre-DNS, and I know myself. If I include more than 30 slides and I kept on talking and talking, and I only have three minutes. I think I raced through. Five minutes left? Okay.

So, just to highlight pre-DNS hosting, the very first one was UCLA dash MNC, remember that for your next pop quiz. And the first scaling problems were less to the host or text files, which led to a second sort of scaling problems, which really led to the origins of DNS. I would love to talk at future points during Tech Day about the, how DNS was conceived.

I would love to take to [inaudible] about this, to other people who were there at the time. Like I said, I was playing with Lego. But would love to know all about it so I can take, give another Tech Day presentation. Thank you. [Applause]

EBERHARD LISSE: Want to show the missing slide.

…the way it looks now.

No, but that was 1995. The point is, I've used a similar thing, I see that. Just to show, and that's what the Internet looks [inaudible]. That, a slide like this is what you wanted to show, I think. Any questions? Rick Lamb.

ROY ARENDS: Sorry, if could send that to me, I will include it in the next presentation. Thank you.

UNKNOWN SPEAKER: [Inaudible] right through after lunch, that's… In your research, did you, how much overlap or how many references, did you see any references in looking at the Internet to all of the other systems? I know you said at the beginning, you were not going to discuss IPX, [inaudible] net, you know, all of these other things. How much of that did you see…?

Having been [inaudible] and written translators between all of those things, it always occurred to me how blind many of us were to the Internet's predominance. I mean, I spent all of my life through [inaudible] to ISDN router. You know, because at the time, we were grasping it's draws, anything, anything, right?

And the number of people that made tens and hundreds of millions of dollars, not me, just translating between the Internet and something else. You know, Apple Talk, or whatever, right? Anyway. Did you run into much of that?

ROY ARENDS: Yes. As you know, success has many fathers. And there is a lot of people that claimed to have been involved at a time, even though there is some of them at a major involvement, some of them smaller involvement, but a success story gets repeated many, many times. So there is an awful lot of overwhelming, similar text about the early ARPANET, which you hardly see is really good documentation, for instance, of what has been done in the UK, or in other military networks.

The ARPANET was nothing new, computers exchanging bits with each other. So there is very little, really good documentation, at least not as much as there is from the early APRANET. Now in terms of naming schemes and translation between them, yes there were early naming schemes. Remember the bash, the

expiration mark. There are different naming schemes basically different systems had different names, different networks had different naming schemes.

But what I also realized that around 82, 83 the concepts that had been used to develop what eventually is now the DNS, has been, is really not original. Has been borrowed from other networks as well, which makes sense, doesn't it? You look around you, you see what works, and you document that and you implement it. So I hope that answers your question.

UNKNOWN SPEAKER: Yes. I did a fair amount of digging myself for Swedish history. The Swedish ccTLD was registered in September 1986, and there is not a lot of people alive anymore. So I tried to find out what our Swedish domain name was originally, and it was obviously [net dot SE?], which was in our WHOIS database. It was created in 1983, three years before the ccTLD even existed.

And the history of this process, how it came about, there is none discovered anywhere. So how it came to the ccTLDs being created and the first names, it doesn't exist anymore.

ROY ARENDS: I'm not sure if there is a question, but if you would allow me to respond. I don't know anything about Swedish history. So bear

with me for a second. What I found, just an enormous amount of pragmatism at a time, for instance, remembered that I mentioned that a [inaudible] name being the first domain name. Of course, that was because the very first root server, well one of the first root servers was named NIC dot [inaudible] dot net.

The same thing is true for that SE name that you just mentioned, created in 1983, but really, really assigned to a foreign body by [Postel] in 1986. What helped me, is I've been going through the ropes with, I mean, my guru here, he is probably here in the back somewhere, is [inaudible].

I've learned an awful lot from that gentlemen. Oh there he is. [Inaudible] he was really there, at the early stages of the Internet in the Netherlands. So whenever I had a chance, we used to work together at Surf Net when I was just dry behind the years, and he was really, he already knew what we were supposed to be doing.

I had a lot of discussions with him and I learned a lot about the early Internet. And Ted [inaudible], others as well, from the Dutch perspective. I'm sure you have these people in Sweden as well.

UNKNOWN SPEAKER: …and some of them are not alive anymore. But the only thing I have is some samples of [inaudible]…

EBERHARD LISSE: I have done the same, what we're just talking about with [inaudible] from Puerto Rico, Murray who died six months ago. Incidentally, he had his first email still there. I should know what happened in dot NA because I've been doing it since inception, I don't have my first email. My email records go back to 1999, since then I've got every single email in a database, a few gigabytes.

But the first emails, I know only what the first host was, was because I still have it, I have it now. I've spoken with [inaudible], they don't even have the document to be able to contact IANA exactly when dot CL was registered. They can only approximate it to 87, because the exact document was lost.

If we had known 20, 30 years ago what we were going to do, we would have kept much better record. But I must say, I have no clue what I was getting into. And many of us, even computer science researchers, would have never anticipated the, how much it could influence our life, how much part of our life, and how much part of our travel arrangements this could take.

UNKNOWN SPEAKER: …going to be, [inaudible]… I still got the original document, how dot NL got established. It was like 25th of April. And you note that the WHOIS entry at that time, actually had [inaudible] contact, it's called a technical contact, which is not alive anymore, it would have to be me. This was actually changed into a kind of how to set up your ccTLD document, which is kind…

What do you do? Who you should ask? That's been used by most European ccTLDs. The interesting part is that this is also the reason why the, why the [inaudible] 76 ccTLDs had name server on the same machine. And of course, published later. Yeah, but that's how the original document is there.

It's interesting to note that [inaudible] the same problem as the host dot text file, it ran out of memory on a regular basis, so we had to recompile the program. But that's another story. If you want to have it, I've got a…

EBERHARD LISSE: Okay. Thank you very much for this reminiscing thing. It was quite good, especially given the hour. Now Scott Hollenback is going to talk about RDAP.

SCOTT HOLLENBACK: Roy, thank you very much for that introduction. It turns out that I've got a little bit of Internet archelogy on my first slide as well. And so your conversation leads right into this. But first, I would like to start with a question, how many of you have heard of RDAP?

It's actually more than I expected, so very good. I don't feel like I'm introducing an unknown topic here. But I'm going to be talking about federated authentication for RDAP, and how it address a particular challenge that we've had with the Internet fossil WHOIS protocol. So let's step down here a bit.

So let's start first with a little bit of a background on WHOIS and some of the challenges that it presents. Getting back to what Roy talked about, WHOIS is one of these efforts that go back to the early, you know, late 70s early 1980s. And while there was a process to identify hosts, and capture them, and host files which eventually led to the DNS, there was a parallel effort going on to identify and maintain a directory of people who were using the ARPANET at the time.

This eventually found itself, this list of people found itself published in a book. And if you look at the picture I've got here on this slide, this is Danny [Hillis] giving a TED talk in February 2013 where he used a copy of something called the ARPANET directory as a prop. Imagine that every single user of the

ARPANET could be captured twice, both by name and by email address, published in a book, and that book would be no more than about two or three centimeters thick, well that's what this ARPANET directory was.

WHOIS was the online effort to gain access to the information that appeared in this book. And it had its origins in the older finger protocol. So if you wanted to, you know, find out for example, the name of, the name of someone associated with an email address, or the email address associated with a person, [you] could use this WHOIS protocol.

One important point to recognize though, is that the original WHOIS protocol was designed for use within a small community of cooperating users. You know, so you have a bunch of people who, for the most part, know each other, trust each other, and they really didn't have an issue with having this type of information available in a place where they could all find it.

Now fast forward 30 or 40 years, and the Internet has changed dramatically. We know longer have a small community where everybody knows each other. Technology has advanced, and we've been faced with many challenges that have been associated with WHOIS. For example, the issues of data privacy, issues of security, issues of accuracy, which I won't touch on, issues of internationalization.

And there have been many contentious attempts, you know, to address these issues by attempting to fix or patch WHOIS with what, I will describe as band-aides. Though how successful have those been? Not very. You can't patch something that is essentially un-patchable or unfixable. And so I'm going to make the very bold statement and say that it is time to stop trying to patch WHOIS, stop [inaudible] it, and instead, take a different approach.

Think about the problem differently, think about a different way of solving [it]. What kind…? How many attempts have been made and what have we done? Without going into the history of the many replacement WHOIS protocols, I wanted to start instead with a description of the expert working group on gTLD director services that was formed in 2013.

I had the honor and privilege of being asked to serve on this panel, and we spent… First off, a little bit of advice. If anyone from ICANN ever approaches you and says, "Hey, how would you like to participate in something that's going to last about two or three months? What do you think?" No, it doesn't work that way.

This ended up being about two years' worth of work, but still, having said that, this was very satisfying. Because I found that I was working with a bunch of likeminded people who really

thought that we were coming together to do something truly innovative. To truly think about solving a problem in a new interesting way without continuing to maintain the status quo.

And while a number of you might have some issues with the recommendations and the final report that the group produced, I wanted to call out a couple of things in particular. Particularly this recommendation that the EWG recommends that a new approach be taken for registration data access, abandoning entirely this concept of anonymous access by everyone to everything, in favor of a new paradigm that combines public access to some data with gated access to other data. Right?

Bold recommendation since we don't know how to do this with WHOIS today. So the big question is, how? How can we do this?

So, the new approach that I alluded to earlier can be accommodated using this registration data access protocol, or RDAP. And most importantly RDAP is not WHOIS. Occasionally I will hear people describe RDAP as restful WHOIS, and every time I hear that my blood boils just a little bit, because it's important to note RDAP is not WHOIS. Okay?

It's specified in RFCs 7480 through 7485, they were published earlier this year. Those are RFCs document the protocol. There is also an informational RFC that we produced called RFC 7485, that we use primarily for note taking purposes within the

[inaudible] working group. It was an effort on our part to develop a catalogue of existing WHOIS functionality, you know, which types of data were people collecting in the various gTLDs and ccTLDs, and who is publishing what, and you know, which bits made sense being included in RDAP by default.

But it's another important thing to note that while we have proposed standards for the RDAP protocol, there are a number of additional specifications that are needed for real operational use.

The working group struggled in a couple of places to reach consensus on a single way of doing many different things. And so you'll find that there are a number of places in the protocol specs where we talk about being able to do things either this way, or that way, or some other way. Or in the case of search, we came up with a very, very simple basic mechanism, and said flat out, that hey, if you want more than this, some additional work needed to be done in the future.

But we did agree that the core of our RDAP was intended to address a couple of technical issues with WHOIS, including the lack of standardized command structures and output and error structures. So this issue that every WHOIS operator, you know, kind of comes up with their own format for processing queries and producing responses, should be addressed with RDAP.

There is no support for internationalize and localization WHOIS, and we've got that built into RDAP. And again, the focus for what I want to talk about a little bit more here today is the lack of security services in WHOIS, and the fact that you have no way of identifying, authenticating, and providing access control to the clients, you know, the people sending queries.

The last, again, most significant point or a significant point, is that RDAP is designed to be easy to implement and operate. Those of you who are familiar with the previous effort to replace WHOIS, produced a protocol called IRIS. Technically competent, very deep and very difficult, and it just never caught on in the operational community.

So what is gated access to data? If you dig into the final report of the EWG, you'll get all kinds of interesting background information about it, but I tried to summarize it here by noting the difference between the capabilities of WHOIS and the capabilities of RDAP. And with WHOIS, more or less, all clients can see all data.

So if you send a query for a particular domain name, or a particular contact or registrar, you're going to get back everything that the server provider has. With RDAP though, type of response that a client receives can depend on a number of factors, including who is asking, what they're asking for, when

they're asking, where they're asking from, why they're asking, and how they're asking.

In order to make decisions based on any of these factors, the server needs to know a couple of things. We need to know the identity of the client, and we need to know something about what they're authorized to see. What is it about identification and authorization? One way you can do these, you can provide these types of services, and what typically happens today with web services in general is, using user names and passwords.

So if you want to gain access to a particular resource on a website, you have identity credentials in the form of a name and a password. But if you think about how this works for services like WHOIS, or RDAP, on an Internet scale, I as a RDAP server operator, just as I said in the first slide, employed by VeriSign, I'm not really interested in trying to provide or give access credentials to everyone who sends WHOIS queries to VeriSign's RDAP servers.

This clearly does not scale well for us as a server operator, and I don't think it scales well for, you know, for people as RDAP clients, because you would need to get user names and passwords from every server operator out there. And when you think about how many gTLDs we have now, and how many

ccTLDs, and how many address registries, this clearly and quickly becomes an [unmanageable] problem.

Another important thing to note is you really need more than a username and a password in order to be able to make access control decisions. Just knowing that someone, you know, has a particular name and a password, doesn't tell you anything about the purpose behind their query.

And when you're talking about this type of data, it does make a difference. For example, if you know that you are dealing with a query from a law enforcement officer, for example, or a researcher, or a data miner. And so we need to have some capability in this identification and authorization process to know something more about who is asking and why they're asking.

And if you're looking for a little bit more background information, have a look at RFC 7481. That's a document that I co-authored that describes the security services for RDAP and the types of things that need to be considered by these services.

So one solution is federated authentication. And if you're wondering, well, what the heck is federated authentication? I've got a little bit more information here for you. I mean, I'm sure folks have seen what single sign on looks like, where you might be asked to go to a news website, and you're prompted to log in

with your Facebook credentials, or your LinkedIn credentials, or your Gmail address, or something like that.

And federated authentication is a very similar concept. It's a means of identifying and authenticating entities based on mutual trust, between members between a common community or a federation. So you have a couple of key players in this federation. You have identity providers, some type of an entity that is responsible for issuing credentials. You have [clients], you know, the people that will assume these credentials that then use them to gain access to resources.

And you have the operators of the servers on which these resources reside. They are known as relying parties in this scheme. When a server operator receives a credential, you know, given that I don't have access to these user names and passwords, or these credentials, or what not. I need to know who I can talk to, to validate the [fact] that this identity is valid, and the credential was issued by the provider, who the [client] claims to have issued it.

So I [inaudible] via an online protocol. They confirmed some information, they give me a thumbs up or a thumbs down. And if all is well, access to a resource is granted. So how does this work? Little graphic describes it. So the first thing that an entity

or a person, a client, needs to do is you have to register with an identity provider.

So in a simple case, this could mean, you know, registering with Google to get yourself a Gmail address or a Facebook credential, or if you're a law enforcement officer, maybe the FBI becomes an identity issuer, and they issue [credentials] only to members of their small community. Once you get the credential, the next step is to use it to gain access to a resource, so that's step three here.

So thinking about this in RDAP terms, you could send [a] query to a RDAP server, and say, "I would like to know some information about a particular domain name, and here is my credential." Okay, so then I as server operator say, "Okay, fine. This is a protected resource and I know, I now have your credential, and I can figure out who issued it, but I need to now go to this identity provider and ask them to actually authenticate it and tell me if this is a valid identity, and if this person is who they claim to be."

The identity provider returns that information to the RDAP server in step five. And where this really gets interesting is that they can also return additional information that was selected by the client. So for example, if the client wants to release

information and make up something completely bogus about their birthday, they could do that.

Why they would do it in this case, I don't know, but there are other cases where that might make sense. If, for example, they're trying to access information where you have to be older than 18 to see it. In the case of the EWG recommendations, one of the things that we said needs to be implemented is this concept of a stated purpose.

So if you are a RDAP client and you're attempting to gain information about a domain name, I'd really like to know if you are a law enforcement officer, or a trademark attorney, or a university researcher, of course, now no one is going to identify themselves as a data miner, but on the other hand, if you cannot present the fact that you are one of these other types of entities, it still allows me as a server operator to make some decisions about what you are authorized to see and what you'll be given access to.

So step seven, assuming all of these steps play out, you receive some type of results, based on your identity, based on your authorization, and based on what you ask for. So an example, I'm going to make a very, very bold statement here and suggest that part of what we do as we deploy our RDAP is not forklift the current WHOIS model of, everybody has access to everything.

And instead, we should start with something that is much more restrictive.

If you issue a RDDAP query without any type of credentials, you should receive very little information in return. And what you should receive, should be no more than what you can get via other public sources today. So as an example, this is one of the kinds of responses you could get if you were to query example dot com.

The response would confirm that indeed, this is a domain name, it has an identifying handle associated with it, the dot, dot, dots are there to eliminate very long tedious types of things. If you haven't seen JSON encoded responses to domain name information, it takes a lot of screen real estate.

And notices is one place where the lawyers get to play their games and tell you, you know, the terms of use for the service and what you can and cannot do with the data, and anyway, it took up a lot of space. But you could, for example, see information about the DNSSEC implementation associated with this domain, and you could get information about the name servers.

This is the same type of information you could get with a DNSSEC query using DIG, if you happen to send something to a root server or something. So no PII, and nothing beyond what

you can't get via other public easily available sources. Okay, running out of time.

So now let's imagine that you were to provide a basic credential of some sort. Like a Gmail email address and you use that to identify yourself. What I would like to suggest is that in this case, you actually be given access to a little bit more information. So in addition to all of the information you get with an unauthenticated result, you might be able to get access to some domain metadata. Things like when the domain was registered, when it expires, when it was last changed, and perhaps some information of associated statuses.

But again, no PII. Okay? Next step would be, I am providing some type of information that I can use to determine that the entity is authorized to see more. So I am authorized to see PII, and I'm going to wave my hands here, smoke and mirrors, about what authorization means in the context of PII. But you would get back all of the basic authenticated message information, and some additional information based on who you are, what your stated purpose is, and what access control decision is made at the server.

So you might, for example, see information about who the registrant is, who the technical contact is, who the billing contact is, etc. All right, so a little bit more about the approach.

I've actually got an Internet draft out there available right now, draft Hollenback [inaudible] RDAP open ID 02, you can find that with your [favorite] search engine.

The underlying federated authentication technology is built on something called open ID connect. And if you want to get a little bit more information about what open ID connect is, I've got an URL here for you. Interestingly, I am working on a prototype implementation of this at VeriSign labs. Hopefully within the next few weeks, [inaudible] maybe even before the [inaudible] IETF meeting, that I'll be able to make some announcements about public availability for the service.

And the idea would be that for unauthenticated queries, you would get back very limited information for some type of [inaudible] would give you back a little bit more. Details are TBD. There is still a lot to do. This is a concept, I know I've been talking to a couple of people about it, [inaudible] dot NZ has expressed some interest.

You'd have to talk to him about how far along he is. But we need to do some experiments to see just how viable this approach is. More server operators are needed, and we need to actually let people start playing with it. We need to find appropriate settings for all of the knobs, and dials, that are available in

RDAP, and we don't yet have you know, any guidance on how the policy ultimately plays out in this.

There is standardization work to continue, implementation work has to inform policy work, and we need to do a lot more work before this is really ready for production. I believe that is my last slide. So if you have any questions, comments, or concerns, I'd be glad to hear them.

EBERHARD LISSE: Thank you very much. Let's give him some applause please. [Applause]

SHANE: Hi, I'm Shane. I guess we know each other. So I wish you luck with yet another, yet another, yet another WHOIS [CROSSTALK]. It's not WHOIS. I think it's interesting, you guys talked about possible similarities with the [inaudible] system? Like is there a chance to a sort of race to the bottom as far as identity providers? Just making it so that it basically adds almost little or no value?

Also, it seems like you're going to need a lot of identity providers. So it seems like there is, maybe you've already mentioned this. I might have missed it. Is there going to be kind of a layer managing the identity providers [as] well?

SCOTT HOLLENBECK: Indeed.

SHANE: I mean, because just law enforcement, there are thousands, tens of thousands of them. Okay. So it's not going to be a separate work? Or…

SCOTT HOLLENBECK: I think so. And I know that there is some policy work happening right now in parallel in ICANN circles. A number of the recommendations that the EWG made are being considered, you know, to be spun up as PDPs. And this is just one of those things.

EBERHARD LISSE: One more.

MARTIN: Yeah, Martin [inaudible]. Stay out of the politics of this completely and talk about the technical side of this. Shane's question is half of what would start my question. Is ignore all of this. When did we get even access to the existing WHOIS system via a RDAP environment in a, not consistent, but in a live environment?

And I say to this you as VeriSign, same way that I've said the same thing to Afilias and the list goes on. I'm digging down to find the URLs to turn up RDAP, and at least looking at the existing stuff. That would be, that would move us away from port 43, at least.

SCOTT HOLLENBECK: Right. Every gTLD operator has a contractual obligation to implement RDAP when three conditions are met. The RFCs are published, it's commercially reasonable to implement it, and ICANN gives us notice to start the work.

MARTIN: Sorry, can I correct my question? Politics and legal aside…

EBERHARD LISSE: We are running a bit late.

MARTIN: I understand, and this would be really great just to stop us using port 43, even if this would be a good idea.

SCOTT HOLLENBECK: Right. So this is one of the reasons that I wanted to do a prototype implementation in a lab environment, because we

have to have some operational experience with these knobs and dials, in order to inform the policy. Policy work, you can't get away from it. It's still happening. There are some efforts underway.

On Wednesday, there will be a discussion led by Francisco Arias, describing the gTLD operational profile proposal that they have. Right, they're trying to gain some agreement on how you can implement RDAP to meet certain WHOIS functional like requirements. And I believe that one of their goals is to get agreement on those things, so that feeds into the policy process, so that we know when we could check those boxes about some minimal things that need to be done to get RDAP implemented.

EBERHARD LISSE: Okay. Thank you very much. We are running a little bit late, so I unfortunately have to cut the discussion down. Thank you very much. Ondrej will now speak about Turris version 2.0.

ONDREJ FILIP: I will start with a question, because I really don't know who changes were done in the audience. How many of you do know what the Turris project was originally? It makes my presentation a little bit easier because I don't have to repeat much.

Thank you very much. So you see some parts of my slide on the screen. I presented this slide, I think, more than one year ago on Tech Day at Singapore. So for those who don't know anything about that, in a nutshell, we do a lot of security research, mainly related to DNS of course, as a [inaudible] domain name registry.

We run the [inaudible] and we do a lot of analyzing from across the network. And we had the idea that we don't know anything about the edges of the network, so we wanted to know a little bit more about that. And also when we looked at the situation of the router SOHO market, it's quite frustrating. Those devices doesn't support like modern technology of everything that it should support like IPv6, DNSSEC. Nothing like that.

And it was really frustrating to, well you buy a router and because it was sent to you, it might be shipped from Asia to your home, it took some well. The software is not very modern, and even if you try to update it, there is no software available, and if it is it is quite complicated.

So all of us have some small router in your home, and it's probably not very good in terms of security and the date of the software. So we wanted to fix this problem. And what we did, or what we decided is to distribute some routers that we made. First we decided, 1000 of them to people's homes.

And those routers would run some security analyzers, provide some data for us. And from that data, we would do some summarize, check some bot nets, and also protect the people. And we wanted to make those routers updatable, so they update automatically so they are secured. And also, one thing which is probably haven't realized, a home router is probably the only device that runs 24 hours in the home, maybe a fridge, but that's all.

That's the only device that runs in your home, and it consumes power, and usually it doesn't have much to do except routing. So we wanted to make a device which would be more generally usable, something that will bring you some additional value. And last thing, we wanted to make a probe, a router which is powerful to get a bit of traffic, [inaudible], and because we couldn't find any capable hardware on the market at that time, we decided to make our own.

So that's basically the Turris project. And we created two versions of those routers, [inaudible] blue boxes you can see on the screen. And as I said, we distributed like 1000 of them, mainly in the Czech Republic. We have focused on the security situation of the Czech Republic. I had some news that hadn't, since the time I haven't presented it.

First of all, we are really working quite hard on the operating system, as a basis of open R2 distribution, so it's a [inaudible] system. But what is the open ability, like automated updates. So that's something we added through distribution, and we are updating, you know, constantly those devices, so we run through 10 major releases, without losing our customer base, which is great.

And in those [inaudible], we are able to fix some problems that appear at that time, for example, the [inaudible], we were able to fix it in days, like two, three days. So we were quite successful. And the system of updates work this way, that sometimes you need some reboots, because for example, you have to change the kernel of the software. So in that case, Turris sends you an email, "Hey, I need to reboot your device.

If you will not do it in seven days, I will do it for you like about three and 4:00 in the morning." Which is usually the time that people won't work. I know you might be different, but some people don't work at that time. So that is all the releases, and some new stuff we added as well.

In the project, we had a lot of discussion about privacy. We have sort of wiretapping the communication. So we had to really [inaudible] this issue very carefully, and I hope we did well, because we wanted a positive [big brother] about, really

positive, about in the Czech Republic, so we are taking as an example how to make very complicated issues [inaudible].

So one of the feature of the device is that it controls the traffic that goes from [inaudible] to when to the Internet, but it never touches anything on the inside, so and we saw that there are some interesting points that might be interesting for the end users. For example, I would like to know what my sort of smart devices are doing, because they are smart enough to communicate with anybody.

And very inspiring, was the case of smart televisions from [inaudible]. They were actually informing HG headquarters what programs you are watching. Yeah, you know [inaudible] and they can do anything they want. Unless you have somebody who checks it for you, you don't know, so that's why we created a module called Majordomo, which is run on the device.

It doesn't report to anybody, there is a [inaudible] web server that you can look at. And you see what your any small devices are doing, who they are communicating with, and so on. We because whatever we do, we really pay attention to the fact that it has to be reusable for others. So especially this module, we uploaded back to [inaudible] guys, so now it became part of opening [inaudible].

In my view, it's not for Turris project, but also for the other home routers. This is an example, how it looks like, you know, you'll see how many bytes were transferred, which port, and to which IP addresses. I think this phone we download it out of emails, and communicated with Google web servers, probably some [inaudible].

And you have, for each such devices in your home, and you see list of devices and so on, so quite understaffed. And it helped some people because they found out, for example, that there home address was communicating with some Chinese IP address, and they realized that probably, that's now what they really want to, and they found some [inaudible].

Quite interesting module. It's something that I wanted to do.

Another thing which is related to the fact that we wanted to reuse the router for other thing, is called Turris gadgets project. We joined with a local company doing alarms mainly, and they have some ambition in the Internet of things. In cooperation with them, we created an interesting like sets of devices that can be easily connected to the Turris.

They have their own [inaudible] protocol.

They use their own [IOS] protocol for communicating with those devices, so we created sets with like motion detectors, smoke

detectors, shake detectors. And you know, with remote control, and also with some power relays so you could just easily plug it into your power circuit, and you can start some devices.

So that's currently ongoing projects. We chose this 100 most active users, and we gave them, again, those sets were free. And we just asked them, do something funny with that and let us know what's going to be the output.

So that's ongoing project and it will finish like by the end of the year. So we will see some evaluation of it. And the main purpose was to find out, what is [inaudible] Internet of things? Is there anything…? What those people will do with it? What they need actually?

I don't know. Will they connect it to the, I don't know, to the radio to make an arm? Whatever. It's interesting, so we will see. And they could request some additional modules, so we will see what they do with it. And as I said, we chose really the geeks that nobody should do with that.

Another interesting thing is honeypot. It's the current security feature. Since we have enough power, we can run honeypot. I hope you know what is in honeypot. It's a device that looks like, it's vulnerable, but it's not vulnerable. It's analyzing what the attacker is doing. And each Turris can run as a honeypot, and he

just a VPN tunnel to form of service, so it's not analyzed directly on the Turris box, but it goes, it is redirected to our servers.

We try to track what those attackers are doing, so this is another module. You can see the situation from, I think it was from August 24th, and this is my home router. You can see how many people were really trying to visit me actually. Belgium, I think Philippians, United States, and they were trying to upload some motion [inaudible] into my router because they thought it's normal home router, which could help them in some interesting work they are doing actually.

So they were not successful, but we know about them. We tracked them. And quite interesting things, we were able to find quite large bot net, which is made from, or collected from [inaudible]. Surprisingly they're using Telenet, there are turn into… Yeah, it works still. I know that it should be old, old face like WHOIS for example, but it is still there.

And what they do, they try to attack as many routers as possible. They have some database on passwords. So on some device, one password is successful, then suddenly all of them they try the same password so they have some central management, and it's about 8000 devices. It's a quite powerful bot net. I don't know what they good for, but maybe they will do some [inaudible] later on.

But they have definitely cooperated. It's quite funny to track them actually. And probably last bit of improvement is, we would like to change resolver. We are currently using Unbound, and it works very well. Thank you very much. It's a perfect piece of software. Since we developed our own DNS, and it's sort of in the final stage, and works for us, a lot to do, move it a little bit.

So the idea is to send it through, not sorry, to [inaudible] as well. It will be in two phases. First it will be voluntary, so [inaudible] there are using it, and people that they are available they are doing, they will be able just to click and change from Unbound to another DNS resolver. And then when we will be sure that everything is okay, that is [inaudible], we will launch it wide. Okay.

Some outputs which are on the webpage Turris, you can look at that. There is a gray list of IP addresses we detected. There are some graphs that shows what are the trends in attacks, what brought attacks, what are not. And also you can look at some response time of some Internet service, the WHO network pings to some important sites and tells you what is the health of the network. It's a little bit similar to RIPE ATLAS. And we publish everything, many in the form of open data so you can download it.

This is an example of how Turris is trying, it's analyzing your connection, and it shows you which speed you download the data, so you can see how you utilize your home line, and whether your ISP is really giving you what you are paying for. So I actually I should have [inaudible] line at home, so there must be some problem because the last week, I never downloaded more than 14 [inaudible].

But anyway, that's, as you see, I was right at home preparing for ICANN meeting, so I was downloading a little bit more than before.

Firewall statistics, this is for example, Turris sees 800 attacks on FTTP, I think they… So you know, each port we track, you can see how many there are. And the last thing, whenever time we presented Turris about people, you know, calling us and saying to us, you know, this is my credit card number. Can you give [inaudible], you know, there are giving us loads of money.

And we say no we cannot sell it. It's not for sale. It's just for the purpose of this project. We're going to give it for free in Czech Republic and we don't want to sell it. But there was quite a lot of demand. Some companies were able to force us to even sell some small [inaudible] was using those devices for network speed analyzers.

So we decided there was probably a market, not a huge one, but sort of community that will like the fact that they can have a really good router with a lot of power that's able to use for something else. So we decided to make a new Turris. At the beginning, we called it Lite because we thought it would be more and you know, less powerful, but at the end of the day, we created something which is more powerful, quicker, and has more functions.

So that's why it's not called Lite anymore, and we Turris Omnia currently. So it's the new generation. It's again, capable for running one gigabyte of traffic. And as usually with everything we do, it's open source, including software, hardware. You can make your own, it's not a problem.

And the goal was to make a board which would be production, I say the production, forget the development costs [inaudible], bill would be $100, so we don't want to pay the production costs. We pay a lot at the beginning. So at least, if you want, you have to pay for the production. So that was the plan and this is the result.

It's very small board, I have it here. So you can see, it has like six ports, one is metallic and also [inaudible] you can plug some modules into it. It has a gigabyte of RAM, it's much more than it

is usual in this market. You get a bit of flush memory for the line of distribution actually. You can run [inaudible]…

So it's really not just router, you can run it in normal server. And it has quite interesting network setup. You know, the CPU has three network interfaces, so one is connected directly to one port, which is [inaudible], and two are connected to switch, and the switch has five ports.

So you can, for example, make some DMZ, whatever you can play with that because the switch chip is capable of making [inaudible], and you can really play with the network quite nicely.

A few [inaudible] to hardware. I know I become nervous, but I will be really quick. It's one of the last slide. So it has enough network capabilities. You have two USB three ports. Many people run [inaudible] server, a [inaudible] server, so you can do it here with USB three. Or you can use the fact that it is [inaudible] so you can put SSH disk inside, and it has three PCI express loads, so two will be probably occupied by Wi-Fi with dual setup, so 40 gigahertz and five gigahertz.

It has SIM slots in case you would like to put more than a [inaudible] backup, or a few connection. It does for example, DNSSEC validation, and it's quite important for the update. So we have RTC chip with battery backup, because we need to

know what is the time of course, and also if we generate some certificates and stuff like that, we have some [inaudible] for better [inaudible]. So we really [inaudible] easily decrypt able certificates.

And some of the stuff, the most visibly important, most, I think the nicest thing is the [inaudible] because they are complete [inaudible] controlled and you can play with that. I think you all know, David has [inaudible] when the car had this flashing thing. Unfortunately, on Adobe Connect I cannot show you animation, but we did exactly the same and it looks really funny.

And last slide. I brought benchmarks. As I said, we are probably the fastest on the market. I don't know what, or the other devices in production. But you know, we have one of the newest CPU, and really, yeah. So just [inaudible] links is a little bit faster in some of the operation, but that is caused by the fact that we're still unable to come to [inaudible] hardware acceleration of some of the [inaudible] calculation.

But anyway, from the network point of view and stuff, we are really one of the best, we have one of the [inaudible] in this market. And I think that's all. This is the first prototype. It's just some bugs which were fixed. We are make second prototype batch in November, so next month.

And we started, you know, inform about the project. We have some 3000 pre-orders for the project, and we would like to have a little bit bigger number because the larger, but you [inaudible]. It can be… So we all started [inaudible] campaign probably next month, to get some real orders, people that really show that they are able to pay for that.

And if that's going to be successful, we will make it in the beginning of the next year. So this is the plan. And yeah, if you are interested, please look at the website [inaudible]. You will see some more technical information about that. And if you will, you know, order some of them, more happy.

As I said, we don't cover, you don't [inaudible] development costs, it's just for the production of those cases. So that's all from my side. Thank you very much. [Applause]

EBERHARD LISSE: Thank you very much. Even if Rick hadn't stood up, I would have asked him to come and ask a question.

RICK LAMB: As a proud owner of one of these things, this is wonderful work. You know I'm crazy about this stuff, but I'm not going to waste anyone's time here. This is all great, simple, really good. It

would be, one of the things that comes to mind here is, I trust you completely because I know where you live.

So, I have this sitting on my network, and everything in the world is gong through it, and you're collecting all of this information on me, and I really don't care. I know you said you're not. Do you have some sort of, and I understand maybe you do, something that one could point to as a certification or an audit process that says, yeah, there are no backdoors in this?

So the first thing my wife said when I brought this home was, "You're going to stick it on the home network?" I said, "You know, I trust this guy. He's cool." And she met you, good. So, that's one. And the second question, is there some [inaudible] comfort that you provide? And the last thing is more of a comment.

If you could make some of those interfaces, like Majordomo really just, much stupider than that, right? I think it would be really useful. Your average home user could see this and see red, that they're going out, I think now you're really hitting the masses.

ONDREJ FILIP: Good point. And maybe would be nice if we could get some advice on that, because we know, we do everything to prove

that there are no backdoors. We publish all the source code, everything, including how to, as I said, we really try hard.

We don't have any other certification, even some sort of [inaudible] certification has to flag it, of course, but there is nothing that we can improve it, I would love to hear it. And we are happy to look at it. One important thing, if you [inaudible] you don't have to, by default, all of the security analysis is not enabled. It's you need to enable, you don't have to be connected to the network that collects the information of course.

But it would be more than welcomed if the people would join in.

EBERHARD LISSE: Okay. Thank you very much. I have to close the queue, as they say, because we are running a little bit late. But Ondrej is really well-known, so we should be able to, any questions could be also discussed offline. Okay, Shane Kerr is the next.

SHANE KERR: All right, cool. All right, hello everyone. I'm going to be talking about the Yeti DNS project. A project I've been working on for about six months now. And right now, I'm working for BII, the Beijing Internet, yeah that's pretty much it for the introduction.

You may have seen slightly different versions of this earlier. I've given similar talk to this to the IPG before the last IETF, and I gave one elsewhere. I can't anywhere. But anyway, I'll go through it. So the background of this effort, is that my colleague Davie Song was at the WIDE camp this year, and he was talking to Paul Vixie and they were trying to figure out, is there a way that we could do research on the root server system without worrying about process issues and political issues and things like that?

So the question was, if, it would be great if only there was some way to look at the technical questions in a scientific way, you know, a way to just do research on the issue of the root server system. If only. And so, they said, all right, let's make such a way. And the way they made it was the Yeti project.

So a quick digression into what the DNS root server system is. I guess most people here are familiar with it. There is 13 servers lettered A through M, and we call these A Root or F Root or whatever. There are right now 12 organizations that run these servers. There are, nine of them are based in the USA, two are based in the EU, and one is based in Japan.

So those organizations, actually any cast from more than 475 different locations. So if you don't know the details, any casting is just where you take a given network prefix, so in this case a

single address, and you advertise it for multiple occasions on the Internet. So it's a way that these 13 different addresses, well 13 different IPv4 addresses and I think 10 different IPv6 addresses, can be presented from these 475 locations.

Now, that's the current root server system. And I mention the idea behind this whole project was to try to address the root server system without delving into issues of politics like that. And normally I've tried to focus on strictly the technical issues.

But I thought, you know, if I'm going to talk about the politics anywhere, ICANN is probably the right place to do it.

So there are some political issues around the root server system. One is that basically the root server system does not evolve. It's roughly identical to what it was 15 years ago. And basically what happened is the guy who was giving out these addresses working with the operators died, and nothing has happened since then.

Just, it's a bit, I don't know how we ended up here, but that's where we are. There have been, I think, three major changes beyond adding, or removing, or re-numbering… Well, re-numbering happened, but adding and removing root servers. One is that root servers now support IPv6, another is that they support DNSSEC, and then the final one is that any casting that I mentioned.

Those are kind of the three actual real changes that have happened at the root server system, and those have all been done by the existing organizations within the resources that they were given so many years ago. So the places where these root servers are, and this, I put in scare quotes, ownership of the root servers is very important to many people.

Countries perceive, I think rightly so, access to the DNS structure as critical, and the root, of course, is at the top of any DNS queries. And frankly, a lot of government officials in many countries are concerned about the US dominance of the DNS in general and the root servers in particular.

Now, I'm happy to admit that these feelings are not really grounded in technical issues. You know, the current operators of the root servers will come, and they'll say, you know, if you want a root server anywhere in the world, just ask us. And there are now, I don't know. Is it, is the local loopback RFC yet? Or is it a draft?

There are technical ways to just avoid the whole root server system completely and run it locally. So as far as concerns about stability and delivery of this information and things like that, it's not really a technical issue, but I think we have to also recognize that all because something isn't technically important, doesn't mean these feels aren't real, and that

[people] don't actually, that these concerns don't have any importance at all.

That's all I'll say about the political side for now. So what is, given that background, what is yet DNS, as I try to say it before, it's a testbed for the DNS root server systems. There are many things about the DNS root, which would be nice to test. There is a whole bunch of tweaking, what we could do with DNSSEC. We can look at different ways of signing the root.

We can look at different key sizes and things like that. We can look at renumbering the root server system, and either in much more aggressive ways and things like that, check out the impacts of that. We can add and remove root servers. We can look at the very scaling limitations of different technologies, and of course, we could look at IPv6 only.

Right now the Yeti DNS is a IPv6 only network, which we like. It's pretty cool, it all works. And you can do all of this kind of testing in a lab, set up a few servers, send some packets back and forth, see what happens. The reason that we do Yeti a lot of ways, is that we want to do the large scale. We want to see what happens to these operations on the Internet. And that's not only by the technical questions of what happens when we see a specific pattern of queries, it's also about the organization and operational issues of running things at scale…

So to be clear, Yeti is not an alternate name space. This is not… This is… We tried very, very hard to draw a clear line between the publication of zone data, and the zone data itself. So what we use the ICANN root zone. We do not change it in any way. We have to make minor changes, which I'll go into a slide or two, but we don't have any interest in the process of that data that is created, or updated or anything like that.

This is just about the contents. As I talked to, in general the project is not about policy or political subjects, but of course, working in such an area, it's very difficult to…

How does it work? It's pretty straightforward. We've got, at the top there, we take the IANA root zone as I mentioned, right now we pull it from F dot root servers dot net, it just happened to be the fastest way to get it. They allow zone transfers from anyone. We've got three different distribution rest, run by three different organizations.

So one is run by my organization, my company, BII. The WIDE project runs one as well, and Paul Vixie is the mysterious TISF there in the middle. Does complicate things to run separate distribution masters like this. We try to run them all completely separately. The reason we do this is because we wanted to avoid people's concern that we're trying to hijack or replace IANA in any way.

There is not any one organization that's trying to run this project. It's a collaborative work. We wanted to make sure that it's not a single group, or single organization, or single person trying to work on this. It's just a group of researchers working on [it].

So what do they do? We take the root zone, we do the minimum modifications necessary, which are basically, we change the NSR set, and then we strip out the old signatures at our own DNS keys in, and resign it, and that's it.

So basically we do the minimum amount of work that you can do in order to change the… So we have our various components. We have those distribution masters, which I mentioned, they produce the zone. We've got the eddy root servers, the ones that actually get this modified zone and publish it so that people can use it.

We publish a hints dot text, which is what the resolvers use as the bootstrapping, the way they start their queries. You can just download that on Github, if you want to configure your own resolver to use this stuff. And then they'll talk to these Yeti root servers.

We do capture all queries that we get, and we collect those in [inaudible] files, and those all get put on a disk in BII servers. Right now, no one has asked for access to those, because we

haven't run any experiments yet. We have a data policy which is that people are going to have to publish a statement that they would not publish these captures further, if they want to look at them.

We don't have a contractual relationship like DNS [inaudible] has to get access to their data, because there is no Yeti incorporated. There is no, there will be no one to sign that contract. So what we've done is a kind of best effort measure that we can do is, people are just going to have to declare that they're going to treat it properly and they won't do bad things with it.

So we'll see how that works. A real live lawyer did look at it, and it seems like it's not completely crazy. So that's good. Not our lawyer, but a lawyer. So who does what in this? We've got the coordinators who kind of started this project. That is Paul Vixie, Davie Song, and [inaudible] from the WIDE project.

And they initiated the project, and they administer the stuff. We've got Yeti root operators. Right now, we've got 12 operators. We have, at least, we have one application in the pipeline right now. We've got a kind of light weight onboarding process that we do. Then we have a couple of more people that we've talked to and are quite eager, and hopefully we'll set up their servers soon.

Right now, one of our servers is in the USA, one is in Columbia. We've got seven in Europe. Actually, a lot of interest in Europe. I don't know why that is. And then we've got three in Asia right now. We've also got the resolver operators, that's actually the biggest missing piece in terms of our infrastructure right now, is getting lots and lots of resolver operators.

It turns out that a good caching resolver, what doesn't send that many queries to the root, so we're looking at other ways to get a query data to our root servers, but it's going to be an uphill battle for us the whole time to get enough query data. And then of course, we also have the people actually defining and running particular experiments on this platform, which is the whole point of the exercise.

So what are these experiments and other investigations? We have a huge laundry list of things that we're going to run. So for example, as I said, we're at IPv6 only, we are going to be looking at the minimum packet sizes, and what the impacts of this alternate fragmentation model is. If you are unaware, in IPv6, all fragmentation is done by the end hosts, unlike IPv4 where fragmentation may be done by the end host, or may be done by something in between.

So it's a different model, and we think it may have different impacts, and we'll see how that works. We also, like I

mentioned, what to look at DNSSEC. We want to look at what's going to happen when we roll over the KSK. I was going to do, as our first experiment, trying the roll over mechanism that was defined by the KSK ICANN design team, but I don't know if they're actually going to do it that way, and it's kind of cumbersome and a lot of work, so we may just do a simple rollover, see how that works.

We also want to look at KSK rollover frequencies, different algorithms, signature sizes. And of course we want to look at changing the set of root servers itself. What happens if you have a lot of root servers? What happens if you only have [inaudible]? We're trying to come up with ways we can turn this set, either the whole set or subset parts of it and see what happens.

Tons of things and we can just play around and see what happens. I think, right now, we don't have kind of a guiding theory and a set of things that we're looking for. We kind of just want to try a bunch of stuff. I think we're still a bit early in the science, and we need to get to the theories of operation later after we've seen what happens. So for example, well, I don't know. We're a bit behind time.

I'll skip the detailed example. You can see it in those slides and ask me about it. Challenges, what we have going forward right now. We would like to get more root operators. We have a kind

of a success that we think would be enough, we'd be having 25 root operators. I think we'll get there, it's taking longer than we had hoped. We already have more than the ICANN root, which is kind of cool. And it seems to work. As I mentioned, our real big challenge right now is getting query traffic.

We have a few ways that we're trying to do that. So one example that we have is we've talked to one of the local universities in Beijing. And in their buildings, they set up an SSID on the Wi-Fi, which says something like Yeti experimental network or something. So it basically gives students a chance where they can opt-in to participate in this experiment.

It's kind of clear based on the SSID on what that means. And if things break… And we've gotten actually a significant number of queries, which is still less than 10 per second, but it's actually an improvement. And we're going to try to reproduce that model in other universities in China, and also in other universities around the world.

Another option is mirroring real time, real world traffic from servers that aren't using Yeti. That's kind of a less ideal case, because if things break, we don't actually find out because the answers don't actually get back to the end users. But we can monitor and look for differences and things like that.

And the final option that we're looking at is things like the [inaudible] network, where we do explicit probes. And again, this is a less ideal case because those queries aren't based on what people actually want to do on the Internet, they're just based on us, you know, checking and probing the network.

Final thing is, an issue that has come up a lot. Is I talk to people about this, is that, is Yeti a testbed? So by this, I mean a testbed is a place where you can run experiments. It is something that can very clearly break. If you thought it would work all the time, if you knew, then it's not an experiment. So the reason we're trying this is because we're not sure if it's going to work properly. And a testbed is also temporary. You don't expect for it to be around for 100 years from now or whatever.

This is just something that you're setting up to try things out. Prototype, on the other hand, is where you're pretty sure what you want things to look like, but you're not 100% sure, so you just want to kind of want to get it ready for the ultimate production version of it. Today, Yeti is a testbed.

And I try to make this clarification to people, because a lot of people, especially people not involved with the technical side, view it as a prototype. Yeti may one day become a prototype, or we may make a next version of it that's similar to this, that

becomes a prototype, or we may make a subset of the people work on a prototype, but for today, it's not.

So if you're viewing this as a way to, as a model of how things are going to look in the future, I would say that's not actually the case. We don't expect any proposal to ICANN based on this right now. That will have to be a later piece of work, which may be related or not. Yeah. Basically right now, we had finished the last couple of months, stabilizing the platform. Pretty happy with that work. We have good monitoring and things like that.

And we have enough root servers now that we can actually start doing experiments. So that's our next step. We are going, we've done some lab experiments in preparation for our Yeti experiments. So I think in the next couple of weeks, we'll actually start our first experiments, and then hopefully we'll start publishing our results after that.

These are the guys who did it. That's it. Thank you.

EBERHARD LISSE: Okay, thank you very much. [Applause]

I can take one question.

Good. Thank you very much then. Next thing would be the DNS incident response panel, Janelle McAlister and Dave Piscitello on that.

JANELLE MCALISTER: All right. So today we're going to be talking about a recent case study of a ccTLD security incident. I talked about this, about ccTLD security in Buenos Aires in about two years ago. So I wanted to talk again and give an update on what's recently been happening, and also some additional resources that we're trying to bring involved to help ccTLD registries when there is a security incident.

To introduce myself, I'm Janelle McAlister. I work with MarkMonitor. We are a corporate registrar working with large corporation, being Google, Yahoo, Facebook, kind of those large corporations. And I've worked with ccTLD registries for the past nine years, and with ccTLD security in this for about the last six years.

So that's been a hot topic. So what I'll be going over is, who is the target in these type of ccTLD security incidents? What MarkMonitor does during an incident? And we'll go through a timeline from a recent incident. And then Dave will go through some of the more tech information from that incident.

So one of our lessons learned that we over the last six years, in addressing ccTLD security incidents, is that anyone can be a target. And it really is not exclusive to ccTLD registries. It can be registrars, it can be DNS providers, registrars, everyone has been a target, and it's not necessarily if you're a target, it's when you are a target and how [inaudible] onto that.

We tend to focus, in this discussion, on ccTLD registries, just because of who the audience is and just because of the information that we have around these types of instances. So that is where this is coming from.

MarkMonitor gets heavily involved in ccTLD security instances, primarily because who are client base is and they tend to be targets when a ccTLD is compromised. Usually these hackers are looking for a way in to high profile domain names, and they're looking for a way to exploit any type of vulnerability that's out there.

And so our process, when we are notified of a ccTLD incident, we do have an internal system that is monitoring all domains under our management, to check for any unauthorized updates. So if anything does happen at the registry level, that mismatches our own system, our system basically sends out both an internal alert to ourselves and to our clients that says, hey this is mismatched.

And then we confirm that it's a legitimate mismatch, or it's an unauthorized update. And then we start contacting the registry, or registrar, whoever has been impacted, to try to get it corrected as soon as possible.

During the incident, because we have worked with so many registries, we do provide some technical expertise. Generally when a registry is going through this type of situation, it's the first time that they're experiencing it, and for us, we [inaudible] quite a few times. And we can give them so insight on what maybe happening to their system, or what has happened with other registries, and start guiding them through the process of how to, what steps they need to take in order to correct the situation.

So there is a slide with a timeline of this recent incident. And what we saw is on one specific day, we received [inaudible] DNS monitoring alerts on our own system, that showed about six or seven high profile domain names had been, had their DNS updated. Now the new DNS was not configured, so the size essentially just went down.

So we contacted the registry, they fixed it on their end, and the process kind of stopped from there. And then what we saw is a week later, name situation had, where there was an unauthorized update, but again the servers were not configured

and the sites just went down. Then two full weeks later, the exact same thing happens, but this time the hackers did configure the DNS, and started redirecting all websites to this website.

And we generally see in this type of incident is variance, some type of warning going on before the hackers take this type of action, and this is what we're trying to get registries ahead of, is you receive any type of warning, or any type of suspicious activity on your system, to be looking for that and taking that seriously and not assuming that it's just an error or a mistake on their end.

So during this recent incident, that we'll get into more detail, this was a very small registry that did not have full-time [tech] support, or any support for that matter. And with their permission, we reached out to ICANN security staff to help them out, and then eventually the network started resource center was involved as well. To help outline what exactly was happening to their system, and what they needed to immediately do in order to fix that.

So with that, I'll hand it over to Dave to get into more detail.

DAVE PISCITELLO: While he's getting that up, I'm Dave Piscitello. I'm the VP of security and ICT coordination at ICANN. I report to someone who is probably familiar to most of you, John Crane, who is the chief security officer. And John now reports to David Conrad, who is the chief technology officer.

So I'm trying to think where I can begin while these are getting up.

While it's being queued, there are some things I can tell you about this. So our team has one component of our service delivery that we call threat intelligence awareness and response. And a large part of that is working with the registry community, the registrar community, security and operations communities for DNS and also law enforcement.

And we spend, we spend a good amount of time, you know, in daily operations, each of us in the team, often we're in different time zones. So we get some notifications from various points of contact. And in this particular case, normally John fields most of the ccTLD inquiries and incident reports.

John was on vacation, fortunately we had some redundancy, I was not. And Janelle copied me and we started an inquiry. We got on the phone call, we began talking about in general terms what had occurred, and it became obvious that it would be very beneficial if what we could is communicate directly with the

registry operator and assist the registry operator in resolving the problem.

So this was a collaborative effort, as Janelle indicated. The ICANN security team…

The ICANN security team, MarkMonitor, obviously the affected registry operation staff, network startup research center, and we did ask for some assistance from [inaudible] who is now on staff at Far Sight, Jose [inaudible], he is formally part of many of the end run activities, IPv6 network, a research network, and was at the University of Oregon.

So let me make a disclaimer here. What we did was a high level investigation. We didn't want to get involved in connecting to a box, doing any sort of recovery research. We are providing guidance. And we're not providing any legal advice, even though the report that we provided gives some general guidance, obviously. We're also not experts in all of the many jurisdictions. This is not a policy document, and this is not an attempt for ICANN to start reaching in and dictating what TLD registry operators do in their own operation.

Often, we're a little bit hesitant to do anything unless directly asked because that is a sensitive point. However, this really was something that really begged for several eyeballs from different perspectives. And I think a very, very good outcome.

What transpired was about a four week procedure where we initially spoke with the registry operator. We talked with the registry operator about some of the symptoms that he saw the attacker exhibit. He shared us with some of the intelligence he gathered from his initial examination of the machine. It turned out, that in this particular instance, there were a number of very fundamental problems that many operators might scoff at.

For example, the authoritative name server was being run on the same machine as mail for the government. So I don't have to go too much further with most of the technical people in this room, to tell you that this is the sort of registry that really does need help from the community. And we were fortunate that they were very open and willing to discuss with us, and they were very patient, you know, very [inaudible] coordination.

And Janelle did a marvelous job of just gathering people together. She [inaudible] coordination. My job was really rather simple, frankly, because she carried all of the heavy weight. In the course of all of the discussions that we have, either through email or through teleconferences, I started to observe that there were a bunch of really, really valuable messages and insights being shared, you know, as we ask questions and then provided more details for the operator to go in task.

So I like to write, and I said, you know, if we could just sort of put this in an immediate interim, long term, time horizon, framework, this might actually be valuable to other registries. Many of the registries here probably know how to do all of this much better than this particular operator, and this is no suggestion that you all would need this, but it's nice to have something for those registries who are a little bit less sophisticated.

So what we cover in this is some investigation basics. You know, and a break in or a compromise of an authoritative name server, in my mind, is a criminal activity. You know, it should be treated as a criminal investigation. And so we talking about preserving the scene. We talk about basic forensics, we talk about reporting criminal acts in very generic terms, and we try to explain, you know, what that sequence of events and what you need to consider entails.

We also talk about restoring service. There are some complexities in restoring service. As an example, one of the complexities that the registry operator [inaudible] is that the government email was being run on the same machine, and the machine offline was a little bit difficult. So we had to just describe ways to migrate over and leave the mail running.

We also had to kind of explain why building from scratch while hardening the machine during this process was important. We go into detailed levels about the need for considering all of the different degrees of exploitation, right down to the BIOS on the machines. And then we have a separate section where we talk about, you know, in circumstances where you are running multiple services, we give some guidelines and provide some fundamentals about compartmentalization.

We also then take the opportunity in the report that we prepared to describe and cover the kinds of event monitoring that the operator should consider, and some intrusion, mitigation considerations for hosts and network. Mistakes have the power to turn you into something better than before, I think that in this particular case, the registry operator came away with much more insight into how to run this operation.

Do you know, we probably know better than I because I don't do real time monitoring, how much better there performance is, or whether you've done any sort of testing to see if they're a little bit more resilient from attack. So we can make that comment in a moment.

The bottom line here, I think is, I'm sure a lot of you in this community, especially here today, know some of us in the security team. We want to try to make certain that the rest of

the ccTLD communities, especially those who can't afford to come to these kinds of meetings, know that we are available. We're not a red team. We're not an incident response team, but we have an amazing contacts list of very, very good technical people.

And most of them, as part of the community, are very happy to step in and intervene, because it's their Internet too. So if you run into a problem of this kind, by the way, we're also extremely sensitive to confidentiality, and we try very much to be nonpartisan, and to make certain that the [inaudible] environment that we walk into and walk out of.

So if you do have a problem and you contact us, we will generally ask you if we can invite some trusted experts. We'll identify them for you, if you want to vet them in any way, we'll try to do that. But there are also a lot of, there is a lot of assistance that may already be in place. There is [inaudible] community out there that is doing the monitoring, and not only just MarkMonitor, but many companies like them that are aware of hacks and trends and nuisances, and they're also many investigators who have contacts in various law enforcement agencies.

So we have an interesting kind of pivot point at ICANN because we work with the [law] community, so we may be able to help

you identify an investigator or an officer in a country where you believe that the attacks or emanating from. And I think that that's the end of my slide. So I'm going to turn it back to Janelle. Thank you.

JANELLE MCALISTER: Thank you. So what we want to do next is actually open up some questions for the community. First being, you know, Dave and I have been talking about, we have a lot of this information, we have insight on what trends are, but what is the best way to share that with the community? And what information is the community interested in knowing about?

Like Dave say, we are very respectable on confidentiality when a specific registry is compromised. Again, we're seeing trends or when we see information that may be relevant to other registries, what is the best way to communicate that out to the group?

EBERHARD LISSE: Can I just abuse the prerogative of the chair? And worrying about confidentiality, what continent was it?

DAVE PISCITELLO: I think if we tell you anything about the geography, it would be fairly to understand what, and pinpoint. So I'd rather not. I also, I fail to mention that the slides were prepared before the report was approved by ICANN legal and posted on the ICANN website. So there is a blog post that announced it, the report is available.

And I believe that we're going to be distributing it on the ccTLD mailing list. That John would be in that discussion. So you'll have a URL for it, it's not a 300 page tome, because obviously we can't go to prescriptive, and especially for some of the audiences who are relatively novice [inaudible]…

So we try to give general guidance. There is a lot of links to very practical material that has been around for a long time. I actually have a question for the community which is, what's your reaction, or opinion, or interest in having ICANN make this kind of service available?

I mean, is this a good thing for us? Is this something we are treading on toes? It's always hard for me to quite understand, especially because in the gTLD space, there are, you know, there are some things that we can do, and there are some things that we can't do. So I'd really appreciate that kind of feedback.

WARREN: Warren, Google. So I've done a number of presentations about DNS hijacking and listing a bunch of them. It seems that in general, folk are broken up into two sets, those who don't have a problem and those who you just can't reach. [Inaudible] is going to be useful, but actually managing…

DAVE PISCITELLO: Absolutely. And in fact, that's one of the reasons why we usually get contact after something is broken. But at least, at this point, we say, you know, here is a report, you know, and here is some things, and we'll be happy to get online with you, but maybe while we're online, and we're working with you, you could also hand this to your staff and possibly pass it over to your legal counsel or whatever.

So yeah, you're absolutely right. And that is the hard part.

EBERHARD LISSE: To answer your question Dave. I don't care if somebody contacts you, if the service is available, it's good. If I don't want to contact you, it's good. So [inaudible] in medicine, when you prescribe something, more helps more, yeah. The more resources are available, the better. I mean, I know, I'm a little bit familiar with the African continent, and this is what Warren was

mentioning, we have [inaudible] registries, I'm unable to reach most of them.

A friend of mine from Germany runs the registrar, he wants to pay two domain names that he registered regularly in Sudan. There is a boycott, so he just wants to get in touch with these people, find out what to do, how to do it, pay [inaudible], they just don't answer the email.

I happen to know somebody in this particular country. I said go to the office please and ask them to answer the email. He goes there, he says yes, yes, and they never answer the email. That's what Warren is saying. Some people don't know what is happening. Some people are just not [inaudible]…

The ones like us, and as I said, some [inaudible] or some automated system, both under that goes for some form of [inaudible] and fairly receptive to a report that something is happening, and you also will find that usually we are not the ones that are affected because we separate a man from, a separator system so that you, if my system got down [inaudible] the whole TLD is effected, not the whole service effected.

But to come back to a circle, the more resources are available, the better. ccTLDs have what is called a secure contact, but this, the ccNSO has a [inaudible] and Christian [inaudible] is one of the people working on it. I don't know where he is in the room.

Where individuals from each ccTLD can use their email address, and they don't want a generic, we are not participating.

But you have got a way of securely contacting your colleagues in the ccNSO. So you have a reasonably secure way to, ICANN in many countries you have [inaudible], in more countries you haven't. But the more resources are there, as far as I'm concerned, the better.

We haven't had many problems so far, but we usually respond very aggressively if we hear something is not right. We get an email we need to look at this and we immediately go very deep, because it's also a business decision. If you guys don't think we can do the job, you won't pay us for it. Some small ccTLDs haven't figured out that one out. That it makes good business sense to have American companies pay us money.

Anything else? A smaller ccTLDs here. Is this something that you will only restrict to ccTLDs, or what happened if a small gTLD has such a problem? What is the course there? We are not really strictly focusing on this issue anymore?

DAVE PISCITELLO: Well, I probably should not say very much about that. I think that gets into a really, really sensitive issue of scale and, speaking entirely from a personal experience and not

necessarily on behalf of ICANN, it's a very odd sensation where we would be [kind of a] help desk or incident response team for the gTLDs.

I don't, an incident, and they ask us for some help, I don't think we're going to say no, but there is a fine line between us, having a business partner, who is also a contract partner, who is asking us for support that starts to fall into the kinds of activities that you would normally have a private actor go and hire or engage in a forensics team to do.

So somewhere along that line, we'd have to say, you know, we can't do this much for you. We can point you in the direction of somebody who is a paid professional, because I think in some respects, that puts us in some sort of liability as well. And again, I'm not an attorney, but I wouldn't want to be responsible for going in and trying to resurrect a system where it's in a commercial operation, and then be held culpable if businesses were offline for meaningful amounts of time.

EBERHARD LISSE: Because some gTLDs are very small, and they have not so far, I'm impressed, smaller ones in particular, I'm not so impressed with their competence. Some have already failed, some are busy failing, some are just wondering about it. Anyway…

UNKNOWN SPEAKER: [Inaudible] from dot RS. There are two things that I would like to say. It is not only registry, most of [things] that happen, happen to registrars. So it doesn't matter how big is [registry]… Talk about registrar, and I really appreciate that you have that…

I think that some people will use your help because not all of these… Things to do, forensic and… We should talk… About registrars and also about… About MarkMonitor… Because in past five years, we offer registry lock. Haven't used that for high profile…

…almost three months after the incident… …lock those domains.

JANELLE MCALISTER: And we definitely try to support registry lock whenever we can. We send out the information to our clients immediately, but some of it is waiting on them to figure out their strategy, what's [inaudible] lock. But we definitely support registry lock.

DAVE PISCITELLO: Let me answer about registrars, and maybe Warren can actually help me. I know that the subject has come up in SSAC about trying to provide some guidance for registrars. The anti-

phishing working group has written a document, some time ago, that is a candidate for revision at this point, but it's very specific for registrars to assist in the suspension of phishing attacks.

And again, it's a commercial enterprise. I'm perfectly open if the registrar community would wanted to have us put together a report or some sort of guidelines for registrars, to have either ICANN or to bring it up in SSAC, and create a multistakeholder initiative that focuses on, these are the sorts of things you ought to do because my feeling is, most registrars are very similar to electronic merchants.

They have the same issues. They have a web application. They have a SSL based merchant. They have some sort of credit card payment system, processing. So they are as vulnerable to any other merchant. And some of them are small, and are not prepared with proper web application testing, and some of them are not.

So I think that we could, if we get enough people like you who want us to do that, we could look into trying to do something like that. So I think it's a great idea…

UNKNOWN SPEAKER: …because we are a small registry. We really don't have [inaudible] to ICANN, will bring those [proposals] to high level that registrars…

Another question maybe, this is not place for it to ask, but there is DNS list, security issues. And I really don't understand how it should…

DAVE PISCITELLO: I'm not in a position to answer how the [inaudible] list works.

UNKNOWN SPEAKER: …and do people really expect that they'll send some, not a sensitive, [inaudible] I'll send to you, because you work for ICANN, and you're a security guy, but not to the list.

DAVE PISCITELLO: Yeah, and I understand that. And I think that some of the lists, it's hard for me to even parse which of the lists have the vetting and have the PCG signatures, because I belong to so many at this point. I think that those lists should have some vetting, and I think that most of the lists that I am on are vetted lists, and I can find out who is on the list at any time, from the listed [inaudible]…

And there are lists where I've left because I am, like you, uncomfortable that the list has gotten out of scale, and I don't know all of the parties. So I think it's a problem we do have to solve, and maybe there are some distribution lists that we can start to use. Are you getting up to comment?

EBERHARD LISSE: Woody is first.

JACQUES: I'm Jacques. I'm with dot [CE], and I'm actually on the secure mailing list group, now it's called the TLD ops. And basically we have a list with about 180 ccTLDs, other security contacts, the challenge with the list is, nobody has ever sent information to the list. We're not exactly sure of, I think you said, how or who, who is going to go first to disclose the other issue.

But one thing the list does is on the, every two weeks it emails all the contacts with everybody's contact information. So we don't know if people are using it to contact each other offline…

DAVE PISCITELLO: So I think, did John have a conversation with you?

JACQUES: Yes.

DAVE PISCITELLO: So one of the things that John's has asked me, and it's a fine idea, is one of the first things that might be your content would be this incident response guidelines document. If we produce a registrar document, we can put that out. I'll go back and I'll look and take to see whether the APWG, a report for registrars is appropriate. And we can have a dialogue to see what it is that you like to have, and if we can produce it, we'll produce it.

And you can publish…

UNKNOWN SPEAKER: So we're looking at expanding the scope of the list? Adding more [inaudible] on the list, more registrar, different input…

UNKNOWN SPEAKER: [Inaudible] with PCH. I was just going to live up to my reputation for being helpful in the least politic [inaudible] way. So because there are people in the room here… I just wanted to try and tease out a little bit of the history behind your being a little shy about offering certain services to ccTLDs. So go back…

DAVE PISCITELLO: Oh no. Let's be clear, I didn't say that about ccTLDs. I said that about generic TLDs.

UNKNOWN SPEAKER: Okay. All right. Fair enough. Being a little shy about offering certain services, period. Whatever. So if we go back four or five years, your [veto] was brought [inaudible] to ICANN to investigate the possibility of setting up [inaudible] search, write a global search that would handle DNS related issues.

And ICANN thought that… From the outside, it appeared that ICANN, as an organization, thought that it would be a good idea if that, at least, investigated. I certainly thought that would be a good idea if that was investigated [inaudible] full support to the notion. And there were a couple who thought that instead of there being a not for profit thing doing this, they should have a business opportunity there, and yelled loudly about how ICANN shouldn't be overstepping its role in doing…

I don't care whether it's ICANN that does this or not, it still seems like a really useful thing to me. PCH operates a cert, we're noncommercial, we have about one case a month for ccTLDs, and about a quarter of those eventually wind up, you know, bringing ICANN in on…

You know, we're happy [inaudible] doing that, but it would be nice also if there were some more formal DNS cert with [inaudible] whatever on that kind of thing. It may… I feel like it was a good idea then, it's a good idea now. There were people

who spoke up against it for very self-centered reasons at the time, maybe it's worth looking into it.

DAVE PISCITELLO: So I think five years, you know, five years of reflection gives people an awful good insight into some of the mistakes made, and some of the maybe more [politic] ways to introduce that concept. If I were to speculate, if a proposal like that came from the community instead of, you know, from ICANN, it might be better received.

You know, there are people in the community, and we all know, who just view almost anything that staff proposes as overreach. I don't happen to believe that. I believe that there are an awful lot of staff who care very deeply about, you know, just keeping the net running well.

And everyone on my team is that way, or our team, not mine. So yeah, I mean, if that's a conversation that we should have, then we should have that conversation again. And you know, maybe it needs to not be called the DNS [cer]. I always cringed when that name was proposed, I said this is a bad idea, people are going to, not even read anything more than the name and react badly to it.

But there is an incident response role, I think, that is demonstrated by what we already do without a label, and what you do a lot of the other [certs] do. So I don't know why we're so afraid frankly.

UNKNOWN SPEAKER: …automobile industry globally can have a cert, and the aviation industry can have a cert, and ICANN…

DAVE PISCITELLO: Right. And maybe it's like, a [sack] bit of a cert, right? Who knows what you want to call it, but you do an organized [sack]. You've got one for research and education networks, you've got one for the financial services, automobile services, and maybe that's what it needs to be. So I'm happy to say, anybody wants to talk about DNS [I-sack], and model it after the [I-sack], let's have a conversation. And I don't know whether it's just here, but you're right.

I mean, it doesn't make any sense for us to flog around, you know, and call people on a Saturday afternoon, and pray we get the right permutation and combination of talent to resurrect something that breaks.

UNKNOWN SPEAKER: …have to be really shy about saying, oh well, but there are also commercial people who might be able to help you…

DAVE PISCITELLO: I still have some shrapnel in the heel of my right foot, so I limp a little bit as a result of that….

UNKNOWN SPEAKER: Exactly. Well anyway, I think Dave and I would both very welcome anybody who wanted to talk about…

DAVE PISCITELLO: One more as a group, and it doesn't violate any trusts.

EBERHARD LISSE: Okay. Any other questions?

DAVE PISCITELLO: Thank you very much for having us. We appreciate it. [Applause]

EBERHARD LISSE: Next one is Warren.

WARREN KUMARI: Well in the meantime, hi I'm Warren Kumari, this is Sarah. We're [inaudible] about DNS privacy stuff happening in the IETF, [innovation] and some work in the [inaudible] working group. I will mumble about QName minimization, because that's easy and non-technical and Sarah will do the actual more interesting part.

So you've already heard all of that. I am going to go through this quickly, because hopefully there will be time for questions [after this]. First off, what's the problem? Well I hate doing expense reports. This means that I don't actually do them and instead I procrastinate, and I tidy up my desk, and I clean all of the crumbs out of [my] keyboard.

Eventually I get desperate and do the laundry. Eventually I start reading Wikipedia. And presumably everybody is familiar with this, the problem with Wikipedia. In particular case, it started off with the song 99 Luftballons. The English translation, 99 Red Balloons, it's all a song about sort of people who launch balloons, it gets mistaken for a [nuclear] launch and there is a nuclear retaliation, and it ends up being this whole new [inaudible] thing.

Many hours of fascinating [clicking] later, I'm reading up about centrifugal enrich of uranium. So what's the actual problem with that? Well first off, I never actually finished doing my

[expense report]. Also, all of the URLs that I went to were [HTTPS], which means that all of the data is actually encrypted.

So anybody watching what I'm doing is not likely get the wrong idea. Unfortunately, none of the DNS is encrypted. This means that if anybody is actually monitoring what I'm doing, many of the domain names were potentially suspicious. And as we've seen, governments and similar folk are using pervasive monitoring to suck up large chunks of data, data analysis, extract signatures, and so potentially that's an issue.

Am I actually really concerned about this particular case? Not really. [If] I'm being monitored, I'm sure it's actually for a bunch of other stuff which we won't talk about, but it makes a good example. So a [couple of years ago] actually I think it was last year, the IETF RFC 7258, which says that pervasive monitoring is an attack, and this is a statement from the IETF's sort of community, that this pervasive monitoring, which is large sucking up of metadata and similar stuff, is actually an attack on the privacy of Internet users, and we all agreed that it is something that is worth trying to mitigate.

So whenever [possible], we're trying to design protocols to make it harder for pervasive monitoring or if not unfeasible, at least really expensive. One of the ways we're doing this, is QName

minimization. There is the draft. I don't know if Stephan is in the audience, the work of Stephan's.

The document is completed working group last call, and possibly IETF last call? Anyway, it's getting along in the process. In order to understand how QName minimization works, a very quick refresher on how the DNS works. These are DNS server, www dot example dot com, assuming we're all familiar with this.

The DNS server asks the root, hey where is dot com. Root replies dot com is there. DNS server asks dot com where is example dot com. Dot com replies example dot com is over there. Eventually it asks example dot com name servers. And replies and you get back the answer. Except this isn't actually how it really works.

The way it really works is user asks where is www dot example dot com, and the DNS server asks the root the entire question, the entire name. And this is all of course, assuming empty caches, etc. Root only knows about dot com, so it replies with the only information it has. DNS server then asks the TLD, the entire query name again, www dot example dot com, eventually asks dot com name servers, all fairly obvious and clear.

So where is the sort of attack surface that QName minimization tries to address? Well specifically, there was no reason to ask dot com the entire query. You know ahead of time, it's not going to know the full [inaudible]…

You're also fairly sure that dot com isn't going to know the full answer. And so if there is an attacker on this link, they could sort of look and see what's happening, they could links towards TLDs and obviously, you know, this is a fairly simple example, if you have a name that has [a lot] of labels, you're leaking a bunch of information all the way down.

So the really short summary of QName minimization is that it makes the DNS work in the way that it sort of naively explained, all the ways [inaudible] the way that people initially think it does… You only include the labels that you need to when querying this specific little… So that provides a fair bit of privacy, but what it doesn't address, specifically over here, from the user to their DNS server.

And this seem like, you know, just a single link, but it turns out to be a fairly regular place where this sort of stuff happens. A lot of censors and stuff use this particular thing. If you're at a coffee shop for example, you're browsing the Internet, chances are anybody who is sitting near you with a wireless sniffer, etc. will be able to collect information. So what it's trying to do is to solve that problem.

Basically we're taking the privacy further, trying to encrypt the DNS messages themselves, and this is more active attacks. I do want to stress name minimization. It's not something that you

do instead. Well, here for example, no deprive, users looking up AA dot org, potentially something that they think is private, [inaudible] all they get is a big blot of encrypted stuff.

And for actual [details] on how that works. I will had this over to Sarah.

SARAH: While these are loading, I'll apologize to anybody who was at [inaudible] a couple of weeks ago, because this is the same material that I presented there. So Warren took you through some of the problem statements, and also touched on QName minimization as a solution. What I'll do in this part of the talk is drill down into the activities of the [inaudible] working group, and what they're doing specifically in trying to encrypt DNS.

So this slide is a graphic that tries to capture the activity that's happened in [inaudible] since about a year ago. So it has already produced a RFC describing the problem statement, and I encourage people to read that. It's a very well-written document. In the solutions space, it has adopted two drafts to date.

One has been through quite an evolution. It started off proposing [inaudible] as a mechanism, combined with a separate draft that proposed using a separate dedicated port

[inaudible], and has evolved further. And at the most recent working group meeting the consensus was that they wanted to [inaudible] just pursing doing DNS over TLS on a dedicated port, and not pursue [start TLS] further as a mechanism.

Separately, there has also been adoption of a draft which describes doing DLS over [D TLS].

So in this slide, what I'm trying to do is summarize why it's non-trivial to decide what mechanism to actually [inaudible] DNS. As I mentioned, the first mechanism that's proposed [start TLS]. So of the things that made it attractive where that it could run over port 53, it's a known technique that's already used with other protocols, and it can be deployed incrementally.

However, there were concerns about how middle boxes would react to suddenly seeing encrypted traffic on port 53. There were concerns about robustness of existing implementations. And also, start [inaudible] is actually susceptible to a downgrade attack during the negotiation which happens in clear text.

And additionally, that negotiation adds a latency to the setup of the [inaudible]. If you compare that to doing TLS on a dedicated port, then what's nice about that is that' it's not going to interfere in any way with listing DNS services offered on port 53. That's nice. And also, there are also existing implementations that do DNS over TLS, and I'll talk about them later.

The downside, of course, is that you have to go through the process of getting a new port signed, and then also of deploying on a new port. We also mentioned that [D TLS] is being considered… And you might think that's a more natural protocol to use for this, because it's UDP [inaudible]. However there are a couple of issues to mention here.

To my knowledge, there is no running code that actually implements DNS over [D TLS]. And also there is one issue which still needs to be resolved for it, and that is, that just like UDP, if the DNS message response size is greater than you, then you have to deal with the problem of truncating that message.

Now when you're dealing with privacy, what that means is you have to consider either falling back to clear text, in which case you lose all of your privacy, or you have to know the mechanism, for example, TLS, or you have to look at extending either the [D TLS] and/or the DNS protocol itself to offer some new functionality to cater for that…

So without one of those solutions in place [D TLS] can't be deployed as a standalone solution with a standard protocol that are available today. I mentioned that the working group consensus was to pursue [D TLS] over separate port, and actually the working group decided to pursue an early port

allocation, and hooray. IANA has actually assigned a port as of two weeks ago, and that's port 853.

So that's now available for a year, I believe.

As a port, and this allocation covers both TLS and [D TLS]. I'm going to spend a few slides talking about the DNS over TLS option, because to do this, obviously you need to have TCP underneath. Now DNS over [tier] has historically been just been used as a fallback mechanism, so often in a one shot mode in response to TC equals one, or for zone transfers.

Although there was an RFC in 2010, 5966, we made TCP a requirement for all DNS implementations. There has also been some significant research by the folks at UFC, showing that connection oriented DNS can be as [performing] as UDP. So following that, there is not a [biz] version of 5966 going through the DNS working group.

It's now in last call. And that goes into the engineering behind how you achieve that performance, and it also tackles some of the issues related to security and robustness when using TCP in scale. One other draft to mention is EDNS TCP key [inaudible], which is also in last call in [inaudible].

This is a mechanism where clients and servers can negotiate PCCP connections for DNS over TCP. I'll just spend a couple of

slides talking about some of those details of the performance recommended in 5966. So firstly, we want DNS servers to handle many TCP connections robustly, in a fashion very similar to [inaudible] servers have been doing for many years.

Also you want to use all the techniques to optimize the TCP and the TLS set up and presumption, both involved handshakes. So TCP fast open can reduce the [inaudible] on the TCP handshake. TLS presumption is also well-known technique, which can do that for the TLS handshake. And I also want to mention TLS 1.3, which is working its way through the TLS working group in IETF.

And that does includes both one RTT handshakes, and zero RT handshakes. So that's on the way. The other thing that you want to do is [inaudible] the host of that setup, in other words, send as many messages as efficiently as you can. So a couple of things that I'll mention that are outlined in 5966.

One is that it recommends that clients pipeline their queries, and by this I mean, instead of one shot TCP, or simply reusing the connection, clients send all their answers, beg your pardon, send all of their queries without responses. And this means that you can gain up to a RTT in receiving all of your replies.

Now for this to work, you also need improvements on the server side. So what [inaudible] to do when they receive those pipeline queries, is to process them concurrently, and then send the

responses as soon they receive them, which means that the client might actually receive them out of order compared to how it is sent to them.

But all of the engineering how to do this is detailed in 5966. To move on to what implementations there are for DNS over TLS, well many people are quite surprised to find [unbound] has implemented it for several years now. It was originally implemented out of the DNS trigger tool, but it's a fully-fledged implementation tool which…

There are patches available for L DNS and N DNS [inaudible]. And I also want to mention that bind has made a number of improvements to its TCP engineering, although it doesn't have TLS yet. Another area of active development is get DNS. If you don't know about get DNS, it's a modern asynchronous DNS second enabled API. [Inaudible] by recursive mode, and in stub mode it implements TLS along with a flexible privacy policy and fallback.

So you can specify that you want to do fully authenticated TLS, or fail, even fall [inaudible] TLS, which is less strict but still encrypts your traffic, or you could fall back to a clear text mechanism. This is just a reference slide showing you the status of some of the features I've mentioned in a subset name, survey implementation.

If there are any questions about this, please grab me afterwards. My last two slides, just very briefly touch on activity related to the use of TLS. There is a working group called using TLS in applications UTA, which just this year produced a best current practice document, covering both TLS and D TLS. Now this is important for DNS, because obviously it's a green field deployment, so we want to be starting from the place that they recommend, which is using TLS [inaudible] and good, strong, [cyber] suites.

So on my last slide, I'll just mention that we also have to think about the question of authentication here. Now, we should be doing secure discovery of authentication credentials. And one thing that has been proposed in the deep five drafts, is a pre-deployed configuration profile. The other thing that we should consider, I mean, this is DNS, right?

So we'd like to use [inaudible]. The minor hiccup is a boot strap problem, where you have to do the [dane] look up over either a clear text or an unauthenticated TLS connection, but that is actually that is done in a few other…

To summarize, there is some very active work being done on encrypting DNS on a deep [inaudible] working group. DNS over TLS performance is key. I actually very carefully considered the privacy policy as we move forward. The appendix names a few

slides which shows how get DNS unbound can interact in different modes. And evolution of TLS and best current practices will also be…

…my last slide. Thank you for your time. And I think we'll be happy to take any questions.

EBERHARD LISSE: Thank you very much. [Applause]

Any questions? Speechless. All right. Thank you very much. Bring us to our last presentation. For [inaudible], Jan-Hendrik and Linda will come to the floor and Christina uploads.

LINDA BREUCKER: Hello everyone. We will do the presentation about the atom feed of registry data. Just a short introduction of ourselves. This is my colleague, Jan-Hendrik Lochner. He is a grand master and software engineer. I am Linda Breucker. I am a project manager. And we both work for Knipp. Knipp is a Germany based registry backend provider. And we offer several services around the DNS.

And one is actually the registry system we're working with, and that is the reason why we're here today. That's our agenda. First of all, we're going to go into details about the atom feed

and talk about the mechanisms that we've implemented. Then we're going to do a short case study about dot Swiss. And then we will talk about benefits that such a feed has, and possible opportunities to offer content via this feed.

JAN-HENDRIK LOCHNER: We start with the introduction to get a common understanding.

I have some slides, some graphics, that will demonstrate what an atom feed is, and how it works, and how it can be used. And the general idea with an atom feed, or a news feed in general, is that publisher wants to provide some information, and to this end, he implements a technical news feed, or atom feed in our example.

And to this feed, he could publish the data that he wants to provide. And on the other side of this feed, there are several subscribers. Subscribes may choose to this feed, and then will get each piece of information, each piece of data that's added to this feed by the publisher. And then we have this kind of data flow, information flow here, to the publisher to the subscribers.

Each subscriber is actually [inaudible] a module and there are several kinds of, where pieces or approaches how to consume such a feed. Most common one is a feed reader software. Today, many email clients and browsers already have integrated

such feed reader components. Here is to read, and to display it to a human user, and a human reader.

Another use case of newsfeeds could be a system that has a feed connector component, which subscribes to the feed. And doesn't display the data of information to the user, but processes it further. And last example, how lead can be consumed, actually a generalization of the first use case presented, is a feed aggregator. It does not subscribe to one feed but to several feeds. [Aggregates?] the information, transport it via the feed, and displays this information to a human user as if it were a single [feed].

This is an example of the technical format. Basically an atom feed is not more than an XML based document, and this other document first contains some metadata about the feed. For example, a title, last updated date, when the last additional, last modification was made to this feed.

And for example, link to the feed itself. And after this metadata, the actual feed entries follow, and each entry represents one piece of information, that is [inaudible] of the feed. And entries, also contain some metadata titles, links to the sources of the information, any ideas updates, the summary of [inaudible] is part of the actual feed data, the actual entry data.

And this example on this slide, is the most common way how a [inaudible] entry is structured. Summary, yes, as the name already says, contains a summary of information to be displayed to the user. Actual content is not part of the entry, but the extra content is linked. The other element, three positions above. So we only have the summary in this entry, and the actual content is only linked to this entry.

The summary in this example is in XHTML format. There are other possibilities [inaudible] can be [inaudible] in HTML or even in plain text. So far the introduction, and now I'm going to present the case study, which illustrate [inaudible] the registry in the domain…

The case study is the registration, the domain registration workflow of the dot Swiss domains. This is one of the new gTLDs, and we as a part of [inaudible] of registrars implemented [the] backend for the operation of the dot Swiss. And as you can see here in the top of the graphic, the registry is the [inaudible] com which is the federal office for communications in Switzerland. And it's part of the government structure of Switzerland.

The dot Swiss workflow is a bit different from usual registration workflows. When a registrant or registrar tries to register a domain for dot Swiss, this domain [inaudible] directly allocated,

but always an application is created first, not only in the sunrise period but also in the general availability. The dot Swiss domain is first creation, and enters 20 day publication period. Is due to the Swiss jurisdiction. There is a law which describes this process.

And so the dot Swiss registry has to adhere to this law. Within this publication period, an interested person may choose to apply for the same domain name. And when the [inaudible] period is over, contention resolution will take place, and only in this contention resolution inside it, which of the applicants gets this domain name, then the according domain is actually registered.

And for the publication period. The [inaudible] implementation gets an effect here. I wanted to prepare a live demonstration, but this is kind of difficult in this Adobe Connect setting, so we decided to provide some screen shots. A real live demonstration, but gets pretty close.

At first, you will see the dot Swiss control panel of the acceptance system. The acceptance system is kind of a [sandbox] for the registry to play with the functionality. And the control panel is basically [inaudible] to this registry system. For example, applications on domain can be created.

Later slides, we will switch to the web WHOIS, which is a base implementation to the WHOIS service. And the WHOIS service, we integrated the registry feed by using an atom feed implementation. But you will see this on the next slides.

This is the login page of the [inaudible]. It's a bit hard to see on the slide. When we log into the system, we get into a dashboard, where we can use to create a domain. In the domain creation process, the registrar has to select the basic data, the launch face data, has to enter the domain name to apply for, etc.

And it has to create the context for the domain registrant, registrar, billing contact, which can be quite conveniently done with the drag and drop in this system. And after having entered all data of the domain, the domain would be created, but not the actual domain, but the application as I already told before.

So in this example, we have the domain test dot ICANN dot Swiss. And this domain has not actually been created, but only the application for this domain name. Next step we switch… I forgot about one thing in the workflow. Do not switch to the WHOIS, but there is one workflow, the [inaudible] for the domains will not be directly published with the feed.

First there has to be performed some manual [inaudible] on the domain name. So this is also done with registry system. In the system, there is an integrated issue system, and in this issue

system your issue will be opened each time an application is created.

Here on this slide, a basic view of the issue. And yes, and this webpage, registry stuff can, yes can… [Inaudible] the application and can decide whether the name, the applied for name is valid or with the registrant is legible to apply for this name. Solved, the issue will be solved as valid in this example. Consequence, the domain name application published to the WHOIS, sorry to the atom feed.

And see the next slide, we use the Firefox browser here, which has the basic feed reader [inaudible], and we can see on this slide that a feed entry has been created for the test dot ICANN dot Swiss. And see here is actually the summary of the feed entry, and content, content that is linked to the entry is for this use case, the WHOIS entry for the created domain name.

So when we click on this link, we will redirect it to the web WHOIS, and with the web WHOIS, each interested person can see the details of the registered or the applied for domain name, can decide, wants to apply for the same name or not. This is the function of the atom feed. And [inaudible] the case study. In the next section, we have some perspectives, how an atom feed implementation can be used for other purposes within the registry system.

Case study which already been implemented, and now Linda will tell about some benefits of the atom feed, and show some opportunities for future use.

LINDA BREUCKER: Well on this slide, you can see a picture where we illustrate a communication problem a registry has. A registry wants to send data or information through different channels, maybe it uses Facebook, Twitter on its website. And the different interest groups, such as website operators or registrar journalists, and they actually have to pull this information out of these different sources, and they also have to filter it.

So they have to do a lot of, they have a high amount to do. And this is actually not optimized at that point, as we think, because the registry does not have a possibility to overlook this process. So maybe a feed could optimize that channel, and if you put a feed in there, it's an additional channel that actually links to the other channels, so you don't have to add additional data or information because it's already there.

You just link to that information. And it's a push process, because of the interest groups are subscribers and they can actually get this information right away. You can actually optimize this further if you, as Jan just said, if you use one feed

for different registries. That could actually be a well optimization for the interest groups.

Then the other benefits we've just put on that slide, it's a standard format. It's easy publishing with nearly no additional effort, as I said. Actually the information and data is already there, you just have to link to it. And it's an immediate and direct communication, because your interest groups will be informed right in the minute or second you push the button.

And you probably could reach a wider audience because it's publically accessible. You don't have to be linked into Twitter or Facebook or anything like that. So possible opportunities to use a feed for, well we did a little mind map, I hope you can read it. Well, for example, you could just publish domain record [inaudible] data, or news could be published, such as staff changes, company changes, or possible policy changes that are upcoming, such as when dot AU announces that they want to open up the domain space for SLDs.

Or announcements such as legal requirements, or when you change actual periods like sunrise periods, or if you want to introduce services. For example, dot UK, announced that they are going to use a privacy service. Or you can inform your interest groups about upcoming events, such as auctions or meetings. For example, if you have a press conference coming

up, you can provide that information and then link to it, or you could provide promotions, or at least a link to your promotion.

For example, if you have the promotions on your website, I just picked one. There are several, for example, dot cat right now has a promotion on their website, and maybe a feed could be used so they could reach a wider audience with this. So actually we are done. I hope you feel well-informed. If you have any further questions, don't hesitate to contact us now or later. Okay.

EBERHARD LISSE: I have a question. Given the Swiss usual efficiency, couldn't they have this, done this more complicated?

I personally think for [inaudible] for example, listed pending the [inaudible] so that it might be something that is [inaudible], which is what peaked my interest. I didn't know that the Swiss had such a complicated way of dealing with an application.

How much effort to, the tool that you developed? Is this going to be proprietary? Or will the underlying software be made available so that for example, [inaudible] as a registry system, or other registrars, registries can use this?

JAN-HENDRIK LOCHNER: Software system used for the dot Swiss registry is, of course, proprietary system. The actual benefit of the atom feed that it is, yes, it's basically an open technique, an open standard, which can be implemented quite easily. There are many [libraries] that can be used, choose to implement it on our own, not to use any third party libraries for the sake of [inaudible], and so we can control the whole system. We don't depend on third party libraries.

And so, the answer to your question is, this system can, of course, not be used for other registries, but the concept as such, to use an atom feed for publishing registry data or other data, as Linda [inaudible], is of course, usable by other registries as well.

EBERHARD LISSE: Anybody else? Thank you very much. [Applause]

So, Jay will give us a short presentation on what dot NZ has done with their data collection, presentation. And then he will close the proceedings, as usual.

JAY DALEY: Hello. My name is Jay Daley from dot NZ. I'm going to talk briefly about our Internet data portal, and then finish up the meeting today. So we have this website at IP dot NZ, where we now publish open data. We collect that's useful for the Internet,

useful for people. So it's about any Internet data relating to New Zealand. It's all open data. It's a place for you to explore data.

In fact, some of the data is not just about New Zealand, it's more global. It doesn't have a lot on it though. A place to explore and a forum to discuss, and it's something we're building [inaudible]. I'm not particularly presenting this as a way that I suspect other people should have to do things, but I'm interested in how we're all going to be presenting and making our data available to others over time.

Understanding that initially, we will probably use very different ways of doing it. And that will take some time for us to stabilize into a more standardized way. This is a chart taken from our data portal, taking from our zone scan, showing TTL distribution in one month for the domain.

So we have, we scan all of the domains in our zone. Approximately once a month. And we collect a variety of different data. I'll explain more about that later. And then on our data portal, we can turn that into a chart. So here is a section of that chart magnified to make life easier for you.

You've got a frequency from zero to five minutes. Then exactly at five minutes, and five minutes to 10 minutes, and then exactly at 10 minutes. TTLs are a very interesting data set because

people, nice numbers for them. And then there are lots of little gaps between where there aren't quite so many numbers there.

And so that's why the overall distribution looks like this, showing the first little bits then it looks there for us. Okay. This is a bit of that underlying data, taken from the portal. These are all MX TTLs. And you can see then the buckets, and then they'll be counting the number of domains in those buckets and then given the figures there.

So, we chose a cloud based product called [inaudible] to run this on. So the software is a service running on Amazon web services. From a US company that normally provides this for local governments that want to do open data. The requirements that we had were for a good catalogue, for a chart designer that is built in to a [inaudible] visualization on their own, and for API access, particularly old data so that people can pull it straight into spreadsheets.

We'd looked at a number of alternatives, open source alternatives, [C CAN, D CAN] and various others, and we rejected all of them. [Inaudible] is not the best product [inaudible] maybe 60 to 70% of what we want and it's fairly expensive, but it's better than any of the other products we found, so we went with this.

So here data sets that we have on there, not a lot. So we have our dot NZ scan, it's run monthly looking for 50 plus indicators. I think under the hood, it runs DNS Check. So it's just DNS Check things. We do some things such as rate limiting the number of queries that are made to a particular name server and other things, other bits and pieces.

And we then upload aggregate statistics for each run of the scan. Then we have our registration data. This is a daily count of transactions, and daily count and number of domains in each of the states. And we upload those statistics every day, up to the third Internet data portal. We then have two other things that we've taken from elsewhere. One is the web index and the web foundation, which shows the contribution the web has made to society. Which is slightly unusual way of describing it, because it could perhaps more described as the contribution the Internet made to society.

And it contains some very interesting different measures in there, from press freedom through to something that packet clearinghouse, [inaudible] collect for them, which is whether or not the country has independent peering exchanges, Internet exchanges. And we then have an ISP survey from our local government statistics [inaudible].

So not a lot of data sets, but we're build it up over time. Now those of you who deal with data, know that there are two ways of representing data, known as long versus wide or narrow versus wide. A wide data set will have one column per measure. So you'll see, you know, say we're talking about the number of counts of different types of access technology. Will say yeah, then account for ADSL, account for VDSL, and an account for cable.

Whereas long has far fewer columns than that. It will just say yeah, then access type, and then the count. And then you'll have multiple rows of in that, one that says that, you know, 2015 fiber and then a number. Another one says 2015 ADSL and then a number, that type of way. That's a long data set. We use long data only in this data portal.

It means many more rows, and it means that you need tools that handle clustering, but it is much more flexible for people to do things with, a better way. But it does make it more difficult for ordinary people to use it, so it's more specialist based. But here is the data structure dot NZ zone scan. The fields that we have are the classification, the metric, the county type, and then the count. And here are some examples. Classification of domain areas, for example, and the metric domains are broken delegation.

But those are standard things from DNSSEC. And then count type could be whether counting domains, hosts, name servers, or DNS Keys, particular states. And then here are registration, statistics, data structure. The date, each day, the parent level domain, because we register under dot NZ, dot [inaudible] NZ, dot org dot NZ and others. The measure, they measure type and account.

And the measure, the measure type could be something such as due for renewal, which is a state that a domain could be in, or created which is a number of transactions that have occurred over that day. And so that's it. That's our, there is some mirror text there, wonderful. So that's our Internet data portal.

And so you can find it at HTTPS IDP dot NZ. So before I go on, are there any questions about that at all? Okay. So to finish up then. Over the years, this session has changed enormously. When we had… We've had a couple of times today when people have asked the audience whether they know about something, and almost everybody has had their hands up.

That's been remarkable. So we really now have an audience here that are an audience of very strong technical people, compared to audiences that we've had, you know, some years ago. And I think that's very useful. That means that when people come up to present, they don't need to start at the

beginning. We're all getting much more value from it, and it's a much greater policy for us.

And so that's a useful facet. And I think that as we now make this much more of a CC day, and it's clearly much more CC day from the number of people that we have here, I think we've succeeded well and we have a good technical group coming along now about this.

In terms of the presentations that we're seeing, we're also seeing some maturity there. We have a number of people bringing second generation box and second generation developments, which is very useful. So we're seeing progress, people that have done something and then repeated it. So if you look at this session just two years ago, we would have eight presentations during the day on how people had implemented EPP.

Now we're down to just one a day, that's much better, you know? We're getting… We went through a period where people were reluctant to share things. I think everybody thought there was something that they needed to keep a bit secret, and we thankfully got passed that, and we're back to sharing things. And that's very useful. So you can see that people are now sharing their technology, sharing their algorithms, sharing their data, publishing these things.

I mean, it's still a lot for us all to take in. There is lots of different sources of data for us to go and find and things, but that's working very usefully as well. And so, I think this is showing real progress for us, and we're doing very well with that. Looking forward then, we have Marrakesh, is our next meeting.

We are aiming to do some sessions there on some general themes, and have a call for contributions there, where we're going to be looking at next generation DNS architecture. So we'll be looking at new ways of doing any cast. We will be particularly interested in containerization, and anything else that people are doing there.

We'll be interested in how people are taking DNS up the stack into the cloud [inaudible] and how we are going to embrace the cloud and infrastructure as a service for DNS. The other thing that we will be looking at is more data quality. We'll be interesting in how people assess data quality within their registrant data, what they do about it, type of tools they have for assessing that, and how they share and publish that information.

We think that's very useful and interesting. So these are, first one of these is a very technical one, and the second one is the usual one where we do things correctly as technical people and teach the policy people how to do it later. So, that's just going

to be two of the themes, but we'll be putting out a general call otherwise as well.

So that's it for us for another year. Thank you all for coming. Thank you all for staying so late as well. We may need a bigger room next time, which is really very good. And thank you all and we'll see you all in Marrakesh hopefully.

[END OF TRANSCRIPTION]