ABU DHABI – Tech Day - part 2
Monday, October 30, 2017 – 13:30 to 15:00 GST
ICANN60 | Abu Dhabi, United Arab Emirates

EBERHARD LISSE: Okay. Good afternoon. Let's come to order again. We are proceeding with the second part of the agenda. Francisco Arias from ICANN will talk a little bit about the RDAP Pilot. I personally am interested a little bit in it because the CoCCA Tools software that – and I heard today 51 other CCDs are using is about to implement RDAP. So I personally am quite interested in it. And I'm quite interested to hear about it.

FRANCISCO ARIAS: Thank you. This is Francisco Arias from ICANN Org. Next slide please.

UNIDENTIFIED MALE: Use the clicker.

FRANCISCO ARIAS: Oh the clicker. Okay. So this is an update on the status of the implementation of RDAP in the gTLD space. So gTLD registries and registrars, those that have a contract with ICANN.

A short history of RDAP on the gTLD site. So on the gTLD site this started back in 2011. The Security and Stability Advisory Community published a recommendation for ICANN to replace the WHOIS protocol. And that was adopted by the Board in later the same year. We published a roadmap to make this happen in 2012. In that same year, the work to develop the RDAP protocol in the IETF started. And that finalized in 2015.

Then we started the work in the gTLD space at first. We worked in developing what we call gTLD RDAP Profile that's as some of you may know, those of you that have work with RDAP will know that RDAP is like many of options that you can implement. So individual implementations need to define what of those features need to be implemented so that the profile that we published last year was an attempt to define what should be implemented on the gLTD space.

Importantly, the RySG group that's the group of gLTD registries in ICANN, disagreed with the way thing were done and requested that that was reconsidered. So that was done. And then we started the work on how to move forward with RDAP. We received a proposal from them, from the Registry Holder Group, how to implement RDAP and we accepted that proposal. That's where we are here.

So the proposal that they submitted to ICANN the next couple of slides identify the points that what they proposed to us. They want us to work on developing a profile or a set of profiles and to implement the RDAP in the gTLD space. They wanted instead of going directly to production as we initially were intended last year. They wanted to have a pilot program and that started to run for instance last September. We have at the moment four registries and registrars that are participating in the pilot.

That pilot is set to run until July next year. Anyone can participate from the user site. There is a pilot website. I'm going to show the link later in the slide. So anyone can participate there. You can go on and play with the implementations that registries and registrars have put forward.

They idea is that at the end of the pilot, we are going to have a date, an agreed date on when to turn RDAP as a production service with the set of agreed profiles with the registry. So some point after July 2018, hopefully we are going to have RDAP as a production service in the gTLD space, and we should have it by July next year.

So the rest of the elements here are just details on what they were asking. There may not be a particular interest to this audience. They relate to certain contractor requirements that the gTLDs have and that need to be waived, for example.

Perhaps one thing to note here is that in this pilot period, the gTLD registries and registrars can play with any new functionality in RDAP.

The profile that we have published last year had, as I mentioned before, a set of functionality that had to be turned on, and another set of functionality that couldn't be turned on. What we are doing during the pilot is anything is allowable as long as it complies with relevant RDAP RFCs. So things, for example, like differentiated access that is allowing certain parties to have access to certain parts of the research and data and not to others but something that is a particular interest given conversations around privacy laws in Europe and other places that are going on right now.

What else do we have here? I think this is the most important part about the proposal. This is the timeline. As I said before, we started the pilot last September. We started the work also on the gTLD RDAP profile or profiles. We are starting with the profile that were published last year as a baseline but is going to be modified by the community as needed.

As I said before, the intent is to end the pilot by July next year and by then have the agreed profile or profiles to use in the gTLD space, and also a timeline to have the service in production. And

of course the date for when RDAP is going to end production is still to be defined.

And how to participate, here are the links. The link to the pilot page, I believe these slides are going to be published, yes. So you can find the link in the slides. That's where you can find WHOIS is offering these pilot services and the links to their services so you can play with them. We also have an RDAP page within ICANN with some basic information about what it is and what is coming.

And we have a session here in Abu Dhabi on Wednesday if you are interested in participating that will be a working session with the gTLD contracted parties on advancing the discussion regarding the future RDAP profile for the gTLD space.

And I believe that's all I had in regards to RDAP.

EBERHARD LISSE: Okay. We still have to participate in the pilot ccTLDs. I'll just e-mail you if I have it – if CoCCA Tools, for example, supported this, I'll just e-mail you. You will send me the details or some of your staff will send me the details. And I'll mess it up the way I want it.

FRANCISCO ARIAS: Yeah, I don't see a problem with that. I think it should not be an issue.

EBERHARD LISSE: The pilot – is it only for gTLDs or is it for anybody who wants to participate?

FRANCISCO ARIAS: So on the user side, it's for anyone. On the server side, it was directed at gTLDs. But at least I personally don't see an issue with all the registries participating.

EBERHARD LISSE: The issue is not that ccTLDs have a contractual obligation but the point is if there is a standard that is developed and it can be supported without much fuss from our side, then each ccTLD can decide whether they want to follow it or not. And it's often I think standards are there because they're used by many people and they're also out and therefore the more people, the more TLDs try it, the better it is in the end for you to see from your side what's happening.

The developers are around from [all software] so we spoke about it and they're busy implementing it, so we will connect it. And since it's in wide use by many ccTLDs, if one or two tested

with your system, then the others can just benefit from the testing.

Any questions from the audience? Okay, thank you very much. You will have to give us… There is one more question from the audience.

UNIDENTIFIED MALE: [inaudible]

EBERHARD LISSE: It's working, it's working.

MARK [HENDERSON]: I'm Mark [Henderson] from Verisign and one of the RDAP pilot participants. And I just wanted to know – better? All right. Mark [Henderson] from VeriSign and one of the RDAP participants. And just to follow up on your comments. I think the RDAP participants would welcome participation from ccTLDs. As you pointed out, if there's a standard we come up with that is something ccTLD operators would be interested in implementing, I think that just makes it all the better. So speaking at least on my behalf, I think we would very much welcome participation from ccTLD operators. Thank you.

EBERHARD LISSE: As I said, CoCCA Tools is used in 52 ccTLDs and they're implementing it in the [packet]. So if we test it from one, it works on the other 51 as well. And whoever wants to use it can use it depending on the data protection regulation. But if whoever wants to use it, I think it will be helpful. I'm personally interested so if you're somebody who sort of takes this – if you have a CCT who is interested and takes this on board, that is also part of this Technical Day so that we can then also report in the future meeting what our experiences are from my side. And if it works on CoCCA Tools that gives 51 other ccTLDs at the moment the opportunity to decide whether they want it or not.

MARK [HENDERSON]: I think that's exactly the hope of the pilot, to be able to leverage that kind of experience and sharing of information. Thank you.

EBERHARD LISSE: Okay. Thank you very much. If there's nothing from the floor and from the remote audience, then thank you very much. And we come to the next presentation, which is about the Etisalat DNS operations. Where is the presenter, Mr. Albanna?

You can just stay here because your next one is after this one. You may have a seat, the clicker is available for you. Wherever you like.

MOHAMMED ALBANNA: Good afternoon, ladies and gentlemen. My name is Mohammed Albanna, Manager/Internet Core Services from Etisalat. Welcome to our presentation today.

Today we'll be highlighting the following points. First four we'll be having an introduction about Etisalat. Then we talk about DNS setup. History started since 1996 until 2015. What challenges were faced during that time and how we overcome these challenges by having DNS Modernization Plan started since 2015 to 2017. We'll show some graphs for our monitoring systems. And eventually, we finish our presentation showing our future roadmap.

About Etisalat. Etisalat is considered now one of the leading telecom corporations in the Middle East and one of the largest telecom corporations, ISP, and GCC. Etisalat started since 1976 for 40 years and is committed to provide customers with the latest telecom technologies. We are targeting yearly to provide customers or to enhance our customer experience.

Mainly Etisalat targeting business and consumer customers providing them packages such as E-Life Home Entertainment. For mobile customers, we are providing postpaid, prepaid, visitors. And we are providing IPTV along bundled with E-Life customers. Fixed voices. And we are targeting also business

customers by providing them Cloud solutions, messaging, managed services, and hosting. Mobile devices and other packages.

This started with our setup since the year 1996. We started our DNS setup by having one server. Having .ae ccTLD zones with customers along with caching and recursion, which was enabled.

In year 1999, we upgraded our previous setup to have one hidden master and two servers in Dubai and one in Abu Dhabi to provide high availability for caching ccTLD and authoritative slaves.

In year 2001, we have done separation for our ccTLD .ae from Caching and Authoritative to another setup by having again master plus two slaves for .ae. And we introduced also reverse zone in the new setup also.

In year 2002, we had an agreement, secondary agreement with ISC, RIPE, APNIC for ccTLD .ae which was beneficial to provide geographic distribution which was one in Europe and one in Asia Pacific with an [inaudible] service also.

In year 2004, we had our dedicated caching system for our network usage and for public customers, which also improved availability and security. By having the system, we had geo

redundancy in Dubai and Abu Dhabi and went to reduce two VIPs behind load balances having old caching setup behind these load balances.

In year 2005, we had the dedicated setup for our own network Etisalat Caching. It is mainly used for our network services such as mail, proxy, hosting, and other systems for Etisalat, which was also beneficial to protect from public threats for security.

Between 2006 and 2008, we still also introduced new eGRX-Emirates GPRS exchange, which was used also for mobile roaming activation. And later we upgraded our architecture to have latest hardware, which was also beneficial to improve our performance and increase Cache Hit Ratio. Later we upgraded from critical vulnerabilities, which was impacted in performance and resources. This was still 2008.

In year 2009, we started to transfer all our .ae, hand over all ccTLD to UAE TRA (UAE Telecommunications Regulatory Authority).

Starting from 2011 until 2015, we distributed our caching DNS in POPs (point of presence), to be nearby end users. And we introduced Anycast for each location to provide also high availability connected to data centers in case of any disasters. This also was beneficial to decrease network latency and

enhance performance supported local and Geo-redundancy, supported more QPS, and there was no single point failures.

We faced during that time the following challenges. First challenge was increase in DNS traffic by having new subscribers and customers, having new applications and services. And we had a challenge to enable DNSSEC which was more difficult than the old setup.

The second challenge, mitigation against DNS attacks. It was difficult to mitigate against DNS attacks such as amplification attacks and pseudo random domains.

Then later we planned for modernization and replacement for the old setup. Started from 2015 by modernizing our first phase of our public caching DNS, we deployed a new public cache DNS at POPs around UAE with high availability. This overcome performance, capacity, and security challenges by having built-in DPI to protect against known DNS attacks, improved the response time by having customized solution, and DNSSEC is also available.

This setup was distributed in Abu Dhabi, Dubai, and Northern Emirates. Having site failover within the same region and regional failover within all sites.

Currently we are monitoring our systems using our dashboard SNMP traps and other monitoring systems. We are monitoring CPU and memory utilizations, number of requests, recursive queue, traffic trends, and Cache Hit Ratio, and other graphs.

Our future plan, to enable IPv6, enable DNSSEC on caching and for authoritative domains, replacing remaining systems for authoritative DNS, eGRX and internal caching systems.

Thank you very much for listening to our presentation.

EBERHARD LISSE: Okay. Thank you very much. Any questions from the floor or from the remote audience? That not being the case, thank you very much.

Francisco will now tell us about the EBERO test. We have had a presentation about this before, if I'm not mistaken sometime ago in Marrakech about one new gLTD that the registrant relinquished because they decided they didn't want to use it anymore. And that was then used as a test case. And apparently there has been another test case and we are very interested in hearing what's happened.

We are not in a hurry so you can take your time.

FRANCISCO ARIAS: Thank you. Hello again. This is the second time in the Tech Day to talk about these EBERO exercises as we call them. We have done three so far.

So this is the agenda that I have for the presentation. First, a short description on what the EBERO program is and how do we go about declaring what we call an EBERO event that's when the need arises to do these process. And then I'm going to show some statistics about what happened during those exercises. And finally a summary and the next steps.

So first, what is the EBERO program? So the EBERO program is something that was conceived in ICANN as part of the New gTLD Program that was launched in 2012. The gTLDs were launched in a 2012 application window. They had in their contracts a provision that allowed ICANN to take over the operation, the technical operation of a TLD in case that certain thresholds of low performance were reached. That was, for example, if DNS or DNSSEC and the contract made a difference on those two. And if any of those were done for a period of four hours or more in a given continuous seven-day window, then ICANN would have the ability if we wanted to take over the operation of the TLD from the registry that has a contract with ICANN.

So I noted too that we needed to have what we call the EBERO providers. We needed to have experienced providers that will be

able to take a TLD operation in a very short period of time. We focused on what we call the five critical functions. I'm going to describe them in the next slide. The point here is that we in the gTLD space try to have many different ways in which the registries will do their business or services that may add. But there were certain core services that we identified as the five critical functions, and those are the ones that we focused to have the EBERO program be able to provide support if there was an issue.

So we have three organizations that we have contracted for this service. And we have a five-year contract with them. I forgot to mention that the objective of having the EBERO program in the first place is to protect registrants so that we mitigate the impact in the services they get from the registry operator if there was a catastrophic failure.

So the five critical functions that the EBERO program focused on are DNS resolution, DNSSEC. We know DNSSEC is of course a part of DNS. However, the differentiation was made in the contract in order to allow back in 2017 when the New gTLD Program was launched DNSSEC was not as prevalent as it is now. And there was a thought that we should perhaps we able to give the gTLD registries a break should there be an issue that was only focused in DNSSEC. So we needed to have those provisions in place in the contract that will allow us to make that

differentiation. For what it's worth, it turned out to be a good provision because most of the issues we have seen in regards to the DNS service are DNSSEC issues of incorrect signing or expiration of signatures or brokenness on the chain of trust, that kind of thing. Very few cases in which there is, for example, no response at all from the DNS servers.

So DNS, DNSSEC, the critical function is the shared registration system. All the gTLDs are required to offer EPP. And the fourth critical function is what we call in ICANN the Registration Data Directory Services. That is WHOIS port 43 and Web WHOIS.

And finally, the five critical function of the gTLDs are also required to offer – well, this is not really a service but a function, is to do data escrow, that is basically a deposit backup of the registration database with a third party and there is a contract that allows ICANN to have access to that data in case of a disaster. For example, if one of the emergency thresholds that allow ICANN to declare an EBERO event, if that's reached, that will also allow us to get access to that escrow deposit, which in turn would allow an EBERO provider to take the operation of the TLD in the worst case scenario in which not only the operation of a TLD is failing, but also the registries are able or willing to cooperate to make that transition happen to the EBERO provider.

So those are the five critical functions. And these are the current three EBERO providers that we have. Okay, so that's the basic information of what the EBERO program is.

So how do we go about declaring an EBERO event? EBERO event is – that's the term we use when one of the emergency thresholds defined in the gTLD contract are reached. And we need to transition such TLD to an EBERO provider. This is just trying to depict the high level – the different components that we have. We have at the top on the left, we have TLD monitoring system that is a system that is measuring the performance of the functions. At this point, we're measuring DNS and RDS or WHOIS and Web WHOIS. And by the way, for DNS we are monitoring all the TLDs, so ccTLDs too and the root zone. Of course we don't have a Service Level Agreement with ccTLDs or the root servers that is done in just a matter of knowing what is happening and if we can offer assistance when we see issues.

So we have of course a contractor compliance function within ICANN that is part of the decision making process to define if there is an issue that reaches one of the emergency thresholds and we have a contract that defines those parameters. And we have an EBERO team within ICANN that has the different functions. We have different processes that we follow and there is a decision-making process that is based on the alerts that are generated by this monitoring system.

Well, I guess I'm getting too much into the – they tell this is the general high level process of how does an EBERO transition looks like. We have in ICANN, there is a provision in the gTLD contract that requires the registries to provide ICANN on a daily basis with a copy of their zone file. So we have those available.

In case we were to declare an EBERO event, we will use the latest available copy from there to provide it to the EBERO provider so that they can use it. Unless of course we were able to obtain cooperation from the EBERO provider. Sorry, not EBERO provider, from the failing registry. Suppose the registries in for a natural disaster in the operations so they are unable to provide the services with the required level. But they are still able to communicate with ICANN so maybe they can provide us with their zone file and maybe we do a proper, for example, proper DNSSEC transition so that we can introduce the keys of the EBERO provider in a way that doesn't interrupt DNSSEC service.

What else? The EPP and the WHOIS services, those in the worst case scenario, we obtain a data escrow deposit from the data escrow agent of the TLD that is failing. With that we – well not we – the EBERO provider will reconstruct the services. And again, of course if the failing registry is available and willing to cooperate, we could get from them the data. But if we are unable to do that, the process with the EBERO provider calls for a cross validation of the data that we obtained from the data escrow agent and the

latest zone file that we have available to identify any issues and we have detailed process to define what remains into the zone file and what doesn't in case there are discrepancies between the two.

I think I'm going to skip this one.

So like I said, we have done so far three EBERO exercises. These are real transitions of TLD. We took over the operation of those TLDs. In all those cases, those were gTLDs that decided for whatever reason to not continue operating their TLDs. They decided they gave notice for all the contractor requirements within ICANN to terminate their agreements with us. And so we asked them if they were willing to let us let's say play with their TLDs at the very end of the process once the only step remaining in the termination process was to revoke the TLD from the root zone. Instead of doing that, we invoke the EBERO process and the full transition. Once we have the services, the five critical functions working, then we continue with determination process which was basically request IANA to remove the TLD from the root zone. So that's in general terms what the EBERO exercise work.

I have the information here. So in those three cases, the gTLDs had no registrants. Those were TLDs that were really not launched. They never opened for registration. They were

delegated in the root of course but the registries for one reason or another, they never launched the TLDs and never got registrants.

Per the contract with ICANN, every new gTLD has the requirement to have at least one domain name active which is NIC dot whatever the TLD is. That's where, for example, the WHOIS services leap. They are the DNS services leap. So they have to be up all the time. So that was the only one domain name that was open running in the TLD and we transition that one.

As I mentioned before, we have done three exercises. We did one with each of the EBERO providers. And all three were successful. They took a different amount of time. I'm going to show the times that they took to do the transition. So in all three cases, we didn't actually have a failure. So in that sense, that part was simulated. We simply said at some point in time that was agreed with the agreement internally with ICANN, we said at this moment, we are considering as we had the DNS service failing, for example. And we followed the procedure that we have and we defined the points that are shown here simulated 100% emergency threshold of the service failing.

So at that point, we declared the EBERO event and we were able to I think – perhaps the most important at that point is the

second to the last. So the change in the root zone from the moment the event was declared to the moment where the root zone change was effected, it took us in two of the cases around eight hours, in our case it took a little bit more than a day.

I must mention here that these exercises were done on a voluntary basis with the EBERO providers. And one of the things that we agreed with them is that we were going to be working in mostly business hours for them. So we tried to accommodate them in that sense. So that's why you see, for example, in the second exercise, it took more than a day because perhaps the issues started when their business day was finishing so they came back the next day and finished the task. That's why there was some variability in the numbers.

The last, the bottom part of the slide shows the total DNS downtime. This is of course simulated because – or part of it is simulated because like I said we didn't have an actual failure. But if we have had one and discounting cache effects which will complicate the equation, these are the time that it took us to restore DNS service, DNS and DNSSEC of course. So you can see that in about half a day, we were able to do it. And this includes that change in the root zone.

In terms of SRS and RDDS, what we did there is we took the opportunity to simulate the worst cases scenario which is in

which you don't have cooperative registry. So instead of asking the registry to provide us with a copy of the registration database, we used a copy from the data escrow agent. So we requested the data escrow agent for the release of the deposit. And here are the times that they took to deliver that to us and the times that it took the EBERO providers to restore the EPP service and the RDS service.

So you can see that the total RDS downtime which is the last service that is still run, it's in this case we have variability between the three cases from two days to five days to hide back up and running. And remember we were simulating the worst cases scenario in which the failing registry was not cooperating and we were using a deposit from the data escrow agent.

There was a lot of waiting time in this process. First to get the release done by the data escrow agent. Also as I mentioned before, we were working on business hours for the EBERO provider and that also extended these times a little bit.

I should mention that in the time we were working on business hours or mostly business hours for the EBERO providers, especially in the second interior exercise telling the providers exactly when the failure was going to be simulated in order to try to at least mimic a little bit real conditions in which we will need

to – they will have absolutely no notice, right? Just the DNS goes down and in four hours, you have to be spinning services.

Here, the only other thing that may be of interest perhaps is data escrow function from – remember I mentioned there are five critical functions of the EBERO program is focused on managing and being able to restore? So the last function that EBERO restored is data escrow because that you do… So you do first DNS and DNSSEC. Then you do SRS. Then you do RDDS. And finally once you have all of that running, you do data escrow. So that you can see it took about eight days. And in one case, we were able to do it even in less time in five days.

The exercises also helped us identify issues in our internal processes. So as the team in ICANN were executing the exercises, we were out in the issues. Most of these were really small issues. We have scripts that are following when we're executing these processes. And there were certain things that were not corrected. Let's say we have the comments that need to be executed in order to give access to the latest zone file to the EBERO providers, and maybe there was a typo in that comment so we need to fix that for the next time. Or we didn't' have access to the PDP key that was needed to sign something or things like that. So it could be issues in the process or in the data that was missing. That allows to clean things up with the EBERO

providers so that should the need ever arise to call, to do an EBERO event for real, then we'll be able to do it.

So that's the number of issues that appears there. And the good thing is we see – how do you say it in English – we have less issues every time we run these process even though that we were running with different providers.

Here's a wrap just so you can see how we were able to restore the different functions. One of the things that we have to consider for future opportunities that we have to rerun these services that came up during the post exercise brainstorming, all these three cases we have done were done with TLDs that only have one domain name. So there is still a question there of the issue of scale, how things are going to work if we were to do it with say 500,000 names or a million names. And so it's something to be considered for the future. The only issues – of course we don't want a TLD with a million names to fail to test this. So we will need to find out a way to hopefully play with such an instance without having an effect on a big number of registrants.

That's all I have in my presentation. Thank you.

EBERHARD LISSE: I have two questions. One the next thing, why don't you load a million domain names into that thing, into that domain name? That's the one question. And the other one I forgot. Why don't you answer this one first?

FRANCISO ARIAS: That's a good point. That's I guess one way that we could do it.

EBERHARD LISSE: The other question was, how much of this is automatable and how much of this you actually sit and speak to the guys at the other registry or the phone, we're doing this now? Which one – or you push a button and the other side says – it's fully automated, you push a button and all these things, the events happen on an automated basis?

FRANCISCO ARIAS: No. Unfortunately, I wish it was. It is not like that. We don't push a button and everything happens. But we do have a process, a script that we follow. But there is always a human that is executing steps. My understanding is on the EBERO provider side, they have the same thing. They are following a script but there is someone that is executing the commands.

I guess in a sense there are things that are automated. They of course just run a command to load, for example, the deposit into database as opposed to do it manually. But not just one button and everything magically happens.

ROBERT: Hi, this is Robert [inaudible] from PCH. I noticed some of these changes in the root zone. And it kind of makes me feel uncomfortable that it actually does affect the security and stability of the root zone. That the change is actually done in the root zone at the end, would not affect everything else when you push the way when you perform your tests. But if people actually monitoring how the TLDs working, they would get DNSSEC errors and stuff like that, and I think you should consider if that's really something that could be avoided maybe and still perform your tests. And as Eberhard says, try to load. Then you can even do a test within an existing TLD.

FRANCISCO ARIAS: I will not dare do that. But –

ROBERT: Without changing the root, I mean, right?

FRANCISCO ARIAS: I see. So a synthetic case, yes. That's an interesting scenario. So to your point on DNSSEC, yes of course. What we were doing here was simulating the worst case scenario which we have on cooperative registry in the case where you have a registry that is able and willing to cooperate. You will be able to do this DNSSEC transition in a way that doesn't break the chain of trust.

ROBERT: Yeah, exactly. In this case it's going to be what I would call a dirty key roll and people will have the old keys cached and they will have – if they do perform monitoring on different things, that probably will fail to some extent. And since I can't [inaudible] about stability that sounds like funny that you would do exactly that opposite.

FRANCISCO ARIAS: Like I said, this is the worst case scenario that we were simulating here where the registries is not –

ROBERT: Except it's not exactly a simulation. But you do effect operation of the root. There are no customers in this, I agree. But you break a TLD, right?

FRANCISCO ARIAS: Well, to be clear the root was working. There was no effect in the root. It was the service in those TLDs that was affected for those tests for sure, because there was breakage in the chain of trust by doing that uncooperative DNSSEC transition. But remember, these TLDs were going to be removed from the root.

ROBERT: No, I agree completely. Yeah. It's just I can't make so many procedures so the new gTLDs have to bend over backwards to invent things that would never break and they did. And then the last thing that happens is it breaks. It's kind of funny. It's not your fault.

EBERHARD LISSE: Okay, next question. We are not in a hurry. So we can take questions from the floor and remote.

DUANE WESSELS: Hi, this is Duane Wessels from Verisign. Two questions, I think the times that you showed to remediate problems, do those sort of match your expectations or you sort of planning on trying to improve on those?

FRANCISCO ARIAS: To be honest, we did not know what to expect at the beginning. I guess now we have a baseline what to expect when and if we ever need to do these for real. And in terms of will I expect this to be better? Yes, of course. If we ever need to do this for real with the gLTD, it's not an easy decision to say it's better to transition the operation of a TLD to an EBERO than give more chance to the existing operator to fix those issues.

So far, we have been able and we had a good feeling that things were going to be fixed whenever issues have appeared in a gTLD that the registry had a hold on things and was going to be able to fix things way faster than we will be able to do it in an EBERO situation. But to your point, I think it will be ideally, we can improve the times, yes, for sure.

DUANE WESSELS: Yeah, I think you sort of already answered my second question, which was going to be, I understand there have been no actual EBERO events so far, right? But have there been close calls or something like that?

FRANCISCO ARIAS: Yes. So we have not done an EBERO event for real so far. But yes, there have been close calls as you call them. I'm trying to remember the exact number we share as statistics not long ago.

I think we said 32 cases to August this year in which the emergency threshold was breached, so it was reached in one of the services. The two services that were monitoring at the time so DNS or RDDS. But in all those cases, we were working with the registry and we had the sense that it was faster to keep working with them to fix the issue rather than do the transition.

DUANE WESSELS: Thank you very much.

EBERHARD LISSE: Okay, next question. And then I have got one last question in the queue on the floor.

UNIDENTIFIED MALE: [inaudible]. Do you have a similar plan if a registrar exits its business? I know this is a contingency plan for registry. How about the registrar? What's the rationale if you don't have such a plan? Who is supposed to follow up if one of your accredited registrar sort of exit the business suddenly?

FRANCISCO ARIAS: So if I understand, you are asking if we have something like an EBERO clause for a registrar that is failing, correct?

Unfortunately, no. We don't have operation like that in the ICANN Accreditation Agreement with registrars.

UNIDENTIFIED MALE: So [less] if the registrar exit the business, suddenly our business? So who is to sort out, take care of the registrants?

FRANCISCO ARIAS: So there is nothing like EBERO. But there are provisions that allow ICANN to transfer their domain names to another registrar if it's basically sell the names or – I'm not the expert on that process but it's more of a business process to transfer the names of failed registrar to another registrar so they take care of the names. Now I will think from a technical point of view it will be easier to deal with the registrar failure than a registry failure. In the case of the registry you need to take care of DNS service. For example, in the registrar site, it is not the case. You just need to be able to take over the operation of the registrar so that you can effect changes on the registry and worse, the DNS will not be failing if the registrar is failing.

EBERHARD LISSE: There have been failing registrars and this has happened already. I can recall not even being a member or employed by ICANN, but I've seen e-mails going around that a certain registrar

had failed and was – ICANN had directed that some for disciplinary reasons and some for business reasons and I think for the failing registrar, there were directions given that the clans must be moved to another registrar.

The problem there is who is going to get the business, obviously. It's not that we must protect the registrants because the registrations are running, the DNS is running, so they have got a year time to sort it out. I could anticipate the system that they say like we have in the ccTLD, we have a registrar of last resort. If we have a failing registrar, we just transfer the domain names into the registrar of last resort and we do not renew them. If somebody wants to have the domain name renewed, they have to come and ask for a transfer to another registrar. Something like this could happen. But I'm aware of none being ICANN. I have been aware of failing registrars and directions having been given that the registrations must be moved. Next question.

JIM: Jim [inaudible]. I think what you're referring to is the bulk data transfer process where ICANN sort of says, "Hey we need a registrar to take over this business." And I don't know who gets the business and how this is determined but it has happened in the past.

Francisco, question. One of the biggest issues that contracted parties in ICANN are dealing with right now is GDPR. How does the new Data Privacy Requirements not only for GDPR but also others an EBRO provider in China who has its own unique set of data requirements, how are you planning to adapt this process to adapt to those new requirements?

FRANCISCO ARIAS: Thank you for that question, Jim. That's as a matter of fact one of the things that has come up in internal discussions. And it's at this moment where we in the EBERO program within ICANN are waiting on direction from Legal as the GDPR issue or privacy issues are sort out so that we have guidance on what will we do. Now in that regard from the beginning when the EBERO program was created, the talk was to have EBERO providers from different regions who have diversity and diversity at that moment was not just in general not with something in particular but turns out that having that diversity is [help].

As you have pointed out, we have an EBERO in China that could help us perhaps if we have a failure of another race in China. Perhaps it will make sense to do a transition to that EBERO in China. We have EBEROs in Europe where perhaps it makes sense if we have a registry failing in Europe to transition it there. But to answer your question, we are at this moment waiting on the

dust to settle in regard to the privacy issues within ICANN so that we have clear guidance on what we can do, what we should do in the EBERO program in a specific.

JIM: Yeah, I think one of the issues you probably do need to solve for is when you're moving registrant data from one existing escrow provider who's providing that service to the registry to another escrow provider. You're moving that data around between ICANN the old registry and the potential new registry. So thanks.

EBERHARD LISSE: We all are waiting for the dust to settle on that one. Just a final thing before I let you go, you mentioned you're testing the ccTLDs. You're probably not willing to share the data in public but would you be willing to share individual ccTLD data with their manager? I would be, for example, be interested in how we are doing in this regards privately between us and other country. If you say that can be done, I will contact you with e-mail and we can discuss this.

FRANCISCO ARIAS: Yes. Thank you for that question. So I think I talked about these in the previous ICANN meeting in the Tech Day. And we have a specific session on that topic on Thursday so if you're interested

on – we have an API that we develop to give individual registries access to the information we see in regards to performance from our SLA monitoring system. And it's currently in pilot. And we are planning to move it to production hopefully early next year, so this session on Thursday at 9:00 a.m. but you can find it in the website. And it is going to describe what registries – and here I mean any TLD, ccTLD or gTLD. We are in fact are already in the pilot. We're ready to allow access to ccTLDs. I think we have a couple that have requested access. And we gave them access to this API so that they can see what we are seeing in the monitoring system.

EBERHARD LISSE: API will not really work for a small ccTLD like me but we'll take it offline. We'll discuss it. I'm also not here on Thursday morning anymore. I have to leave on Thursday morning. But we'll talk about it over e-mail and if necessary, report on the next Tech Day.

Akkerhuis is the next guy. We'll talk about the Root Canary and then Arends. Root Canary sort of tells about measurement about the KSK rollover. And then Arends will tell us why this has been delayed. Is Roy around? Why don't you come up in front already and give Francisco a big hand again. We are not going to be in a hurry, don't worry.

JAAP AKKERHUIS: Okay. I was supposed to talk about effects of the KS rollover but well something strange happened on the way to that event and now I don't have a talk. And so we had to change it a bit and we'll talk about how people are looking at the [changes] in this offer, the effects in general and which is known as the Root Canary. Note that I'm just a messenger. I'm not really doing the work myself. I'm just taking care of tea and sympathy. So a lot of detailed questions, I cannot answer I think. So that's all set. The effects of KSK roll over. There's a lot of speculation about what the supported algorithms, how they will behave.

You want to know about packet sizes, and in general, you would want to know, getting knowledge about it. How does the process goes and just in case you want to do it again for an algorithm change or something like that. I don't have a lot of data to show but I could tell you a bit about the mechanism how it went.

It all started at the RIPE Atlas Hackathon in 2017 when [inaudible] came up with an idea. Well, the RIPE Atlas Hackathon was especially based on things you can do with Atlas. And [inaudible] decided, "Let's have a look at how DNSSEC is done all over the world." Especially because of free seconds to play with and it's kind of an expensive exercise. And he was really

interested in what our algorithms really supported by the various resource in the wild.

So we're doing a test against various name servers and try to explain this in a nice and understandable way for the not all too technical people. So he came up with the following pictures and this is actually something I [measured] yesterday. And this shows what the name servers which are used on this network where we are now from the ICANN room, that's where it was, changed its name to ICANN WPA – I don't know why.

But anyway, what is supported and vertical you see the digests being used and horizontal the various algorithms which can be used for doing this stuff. And the yellow line – these are the [inaudible] secure – they cannot be validated. Let's put it that way. People getting [close] to you, bad luck. I mean you cannot validate the DNS. It will still work but you cannot validate it. I mean that's basic.

And so depending on which brand of resolver, you can actually get a different result. As an example and because I wanted to show some red, this is a typical cuckoo example I [measured] that yesterday also from here. And the interesting part is that they actually don't want to deal with RSA-MD5. Their reason is that MD5 is to stop using it. It's deprecated. So they keep a servfail with all the means people using it still in their domains.

The domain doesn't exist. So you won't know that's a wise move. And if I look closely to – I mean I'll understand what is written, yes, it might be deprecated but it's still part of the standard. So I'm not sure whether this is good idea. But again, you can see different things around here. So that is what kind of state was after the Hackathon that [you call] do this type of stuff. Didn't look that nice. But anyway.

What we really want to do is to continue measurement of this stuff all over. So that's why RIPE was interested in DNS as well. So we managed to get this test standard into all the Atlas probes. And Atlas probes – these are these [Katarina's] cookies. The little help everybody can get and help RIPE doing measurements. I [inaudible] five with me so if people are interested in running an Atlas probe, then just let me know and I'll catch you up.

So this is a picture here of random point on the map. And apparently that probe is talking to resolver with [inaudible] which is using the Google 8.8.4.4. And the red stripe on the left actually gives it away, it's Google stuff. So this is done every hour. It takes a lot of kind of energy so we are very happy that RIPE actually saw some stuff and this will help us out.

So the Root Canary was born. And the Root Canary is like the canary in the coal mines so you just hang a canary around and it

will die quicker than humans when gases are around. So maybe you can see from the patterns whether there are serious things wrong with the DNS and DNSSEC – no, especially DNSSEC validation. It's a combined project between – RIPE NCC takes care of giving us all the right probes. I'm not really sure ICANN is doing that but SIDN Labs is taking part of it. And now NLnet, SURFnet, and University of Twente is actually kind of the leading paper there. And the Northeastern University in the U.S. and because they are doing something interesting as well. It took care [with trend].

And now I'm stealing all the slides from somebody from the Root Canary. It's like more molded from NLnet Labs at DNS org and I just kept a couple of slides to show some of the highlights. So the goals are tracking the operational impact of the root KSK, that's the main problem of the rollover and can't see as [avoiding]. These things are going on and measuring the validation, how it behaves, doing the rollover. So we can learn that for these type of events. So that's the whole idea behind it.

This is the other slide which I by accident left in. and here we go. There are four ways of looking at it. This is what it do. Using the RIPE Atlas probes and not [Katarina's] cookies. And Northeastern University is actually using Luminati, which is a kind of [inaudible] network thing and you can actually abuse for doing serious work. This is also some combination with APNIC

DNSSEC measurement. And at least it's also about doing that. I'm not sure how far they are on this moment. And then there's an offline looking at a date, you can collect and see where you can find patterns. And that's where big data stuff from [inaudible] comes in. And looking at every tool root names, things like that. And I guess that's where ICANN comes in.

Anyway, Luminati, it's an HTTP proxy service. And tests do have exit nodes and mainly of residents who uses. So we can do some things also. And this HTTP request a DNS query can actually do some DNS measurement. It covers about 15,000 ASes and 14,000 of them are not covered by RIPE Atlas. So it's completely different public. Which is kind of interesting and it also shows you in the measurements later. But it's good to see the world from multiple points of view and not just from your own cozy home.

Here are some preliminary findings. We actually measured something. And that was when this is [how in] the slide. It says ZSK in zone, they mean when the KSK was added to the zone whether or not it gave any disruption of the current validation. Although not used, you never know what we saw it do. They might get completely confused validates answer. Well what you see there is not looking – what you see there is apart from the little spikes, which they have the explanation for but I forgot

what it was. There is not really a lot going on. We just did not get – kept taking like it's just it. The display doesn't.

Anything happening there? I didn't touch anything. Yeah. But I don't know what slides I want for myself. One second, I kept the slides.

So for remote users, [inaudible] fell out here. So if you still have sound, you can listen in and that's about it. Here my slide – I caught it already. So yes.

Well the other one is also a nice picture. I cannot show you that but the idea is that you might be able to – there it is. Back again. Okay, here we go. Almost. The clicker doesn't do it yet but we will wait for that. And the next one we actually see is after the KSK got introduced, how it is different between the UDP traffic and the TCP traffic.

Because what happened on that moment is that – I'm trying to do next slide but even – there's the next slide. But they don't see it anyway. So after the new key got introduced, the packets got bigger. So the interesting thing was whether or not that had an influence on the difference between the traffic for – these are all the clicks I have – between UDP and TCP traffic because there was some idea that that might be happening. Well the good news is that you don't see it anywhere, at least not in our measurements.

The relative use of TCP and UDP is about the same and the amount of TCP traffic kept the same so it's all – some measurement is no big deal here. It doesn't do it for me either. Not a measurement we have which you cannot see is that this slide which is also on the Canary Root website itself which gives you the efforts of the total which validation algorithm has been used and which are actually used for securing insecure things and which kind of interesting overview site.

The next one is actually more interesting and there is the difference between the Luminati probes and the Atlas probes whether or not – and then you can really see it's a big difference. According to the Atlas probes, about 42% of the resolvers are actually validating. And so that's quite a lot. The Luminati guys at least uses of [inaudible] being done there is only 7%. But again, remember this is completely different public. People who are running Atlas probes are actually the people who are interested in DNS and the people doing Luminati, the uses of them that interested in way other things than DNS. They just want to see things they're not supposed to see and mainly.

So that's one of the difference between those measurements. Around three minutes. Another thing you can do of course we told is trying to fingerprint the resolvers. That's an interesting way to do and I already showed you the Google public DNS fingerprinted and that's more slides which is more – this one

which is very rare is fingerprint the knot resolver and the powerDNS recursor but what happened there was that we actually take it some bulk which is now all being completely repaired. Actually almost possibly asked 12 people about it, they fixed it. But that's another side [effect]. We never sought about it. You actually could do remotely [deback all those] people resolvers.

There is another completely different experiment and it needs some explaining. And what happened was that it was actually an expired signature put into the zone file, which was very quickly removed again but the interesting thing is there's now a tendency for resolvers to serve stale data. So how long will they serve stale data? So this is kind of an experiment. Let's see how long they still serve the data.

And this is a graph about it. It's about two hours and then everything was going. But it brings it to the interesting question, if people are doing this stale data stuff, which is actually kind of protocol violation but on the other end, the root disappears because if you're [in] an island you don't want to have complete island without any Internet.

So that's an interesting amount of research we probably can do as well. So the Root Canary guys… There's at least one URL you've seen already and this slide with all the URLs where you

can see all the details. For the rest, we are still twiddling our thumbs and coming up with new things to measure, which is interesting system.

That's about it.

EBERHARD LISSE: Thank you very much. Exactly on time. While you are sitting here, I can embarrass you further by mentioning that you have been recently introduced in the Internet Hall of Fame and I think that also deserves an applause.

Any questions from the floor?

DUANE WESSELS: Duane Wessels from Verisign. You had a slide about serving stale data. Was that really serving stale signatures? It was after a DNSSEC signature expired, is that right?

JAAP AKKERHUIS: I don't know the full details because not everybody notices.

DUANE WESSELS: Okay, but the slide of course said something you messed up signing or something like that.

JAAP AKERHUIS: Yeah, the messed up automatic signing and so on and so forth was an expired signature in sort of a – that's what they understood. And they were looking at how long they saw no error in the resolvers after it was cleaned up. Actually, it smoothed out. So that's what happened in that.

DUANE WESSELS: So I guess the implication is that some validators may be sort of stretched out or don't exactly honor the signature expiration time stamps?

JAAP AKKERHUIS: Probably yes. That might be what's happening. The other things also we note all the probes do it once an hour so they might not even have seen it.

DUANE WESSELS: Or their clocks are wrong.

JAAP AKKERHUIS: Yeah. So the details more it's looked at it very quickly but it's an experiment you don't want to repeat. But it happens.

EBERHARD LISSE: Any questions from remote? Okay, thank you very much. And we then want to mention again, if anybody wants to have an Atlas probe, Jaap has got five Atlas probes available and you can just approach him and he will give you one. I've got one running at my house for many years. I don't know what it does to you but it's running.

JAAP AKKERHUIS: I think it's one of the three in [Namibia]. Here in this area, there are very few so it will be very nice to have a couple of these ones in this area as well. It allows you to do various experiments. Looking at your own network for the outside world, things like that. Run your own tests. But that's complete different talk to what you can do with Atlas.

EBERHARD LISSE: Okay. Roy is going to tell us why they are delaying the rollover now.

ROY ARENDS: It's too bad you actually can't see the slides. They're beautiful slides. I've worked on them a long time. Anyway, my name is Roy Arends. I work as Principal Research Scientist at ICANN in the Office of the CTO. And I'm just looking at the front key. Let me know if this comes up again so I can use the clicker. Thank you.

So when you validate DNSSEC, you need a Trust Anchor. And a Trust Anchor is nothing else than a DNS key – sorry, Public Key. And this Public Key can't live forever. And the community decided a long time ago that we should roll this once in a while. As we all know, Public Keys shouldn't live forever.

There is no way for us to check if you have configured these things either manually or using RFC 5011, that's an update mechanism that you can use. There's no way for us to check. So we knew that. And so a multi-year effort to roll the key. We started some years ago. And we had a design team. We had blogs, outreach, presentation in various venues, plans, vendors. We talked to governments, etc. in order to alert people that on October 11 we're going to have a new key.

There's a process in place. I want to make that very clear. It's not something ad hoc that we do. On July 11, we introduced a new KSK. We looked for changes in the root server traffic. We have access to B, D, F, and L root traffic. And we monitor if there are fundamental changes. There were not.

I'm going to use the clicker now. I'll just say next slide if that's okay. Thank you. So next slide, next slide. That'll work.

So then on August 10, this is 30 days after we introduced the new key, there's an element in RFC 5011 that basically says you have to sniff the key for 30 days. You have to observe the key for 30

days before you can configure it. This is called the hold down timer. So you monitor if there are any changes in root server traffic. And if not, you can continue. And if there are many changes, then you might be able to fall back.

Nothing happened. We monitored the data. Everything is fine. And then on September 19, the DNSKEY response size increased due to the standard ZSK roll. This is where Verisign do their regular ZSK roll. And Verisign is a root zone management partner, and they do this every three months. And they've done this since 2010. So there's nothing strange there. But of course we're in the middle of our own KSK roll. So the DNSKEY response size increased. Yeah, that's fine. Next slide please. The clicker works. Thank you.

Those dates, this is a beautiful slide created by Duane Wessels. He's sitting over there. He's from Verisign. He's a researcher at Verisign. I'll come to this slide back later. Just the two lines that you see here is July 11 where we introduced a new key and then a few days later on August 10, 30 days later, you can see there's massive swap. It's basically beautiful to watch. The red above is at 2010 key tag that the resolvers are signaling. And the green below is the 2010 plus 2017 key tag. In short, green is good. Red is bad. Okay.

So who has KSK 2017 configuration as a Trust Anchor? Because by now, after the 30 days and a lot of outreach, we were assuming that a lot of people had. There is and this is referring to the data that Duane Wessels got from A and J root if I'm not mistaken, was this from both? Yeah. From A and J root. I don't want to say anything wrong so all mistakes are mine and all the credit goes to Duane for that information.

So this RFC 8145 allows you to basically sense Trust Anchor information to the roots. Now this is very, very new. This became an RFC in April 2017. I looked at it a little bit before but I got through like 100 IP addresses that were circling this and this was long before we did the introduction of the new key. So there was nothing to see there. Implementation BIND 9.11 started last year. B3 means Beta 3. So a beta implementation existed with the signal.

Anyway, Unbound 1.6.4. The cool thing is in BIND, it was on by default. Unbound is on by default since version 1.6.7. The reason they do this a little bit later is because [inaudible] labs had a policy to only switch something on by default if it's an actual RFC. We don't know if any other implementations – validate implementations as a sent signal. Except for – okay, not this as well. Thank you for the information. We need to talk afterwards. Okay.

So looking key tag signaling, it was so new that we didn't expect to get any data. Keep in mind there's on average 4.2 million unique IP addresses talking to the root servers daily. This is on average. It's not every day that much.

So I'm going to go to the next slide. What you see here is the same slide that Duane sent us. But that we're hoping to be a little bit closer to zero than basically between 5% and 10%. So we analyzed the data as well. We again looked at B, D, F, and L root. And it then showed that this is data until October 24. We have about 27 instances sending us this data. And 1631, this is 6% showed that they still had the old Trust Anchor configured. And 6% is of course non-zero.

Then there were few implementations that sent us KSK-2017 only, which is very interesting because it's not live yet. And analysis is complicated. What you see when you actually look at the live data to raw data is you have validating forward just forwarding to resolvers that do not validate at all. You have IP addresses – sorry, you have implementations that basically walk around in a single network. They have different IP addresses every day. And so these things can inflate and deflate the number. So it's never really an accurate number unless you do the research and that's what we're doing now.

So we did some research into this. This is the latest. There are multiple reasons that we think you only see KSK-2010. This is a cute one, BIND reports a Trust Anchor even if it's not validating. So the idea is you can configure a Trust Anchor and then decide not to do DNSSEC at all. So you have DNSSEC validation now and it will send you that it has KSK-2010 configured. Now there's no problem here because it's not validating at all. But it's kind of a false signal, if that makes any sense.

An old configuration – not old configurations. Historically, if you look at for instance BIND, you use the configuration statement trusted keys to configure a key locally to configure a Trust Anchor. Later on when RFC 5011 came out, there was another configuration statement that you can use and this is managed-keys. And in managed-keys you can basically say this Trust Anchor will be updated using RFC 5011. So people have that confusion between the two statements.

And then there were bugs – bugs is a big word. I'm shooting from the hip here. There's a few instances where you are not able to write the new Trust Anchor to disk because after you dropped privileges, you don't have privileges to go to that file anymore. We've seen that a few times. And then there's an operator error. Things like this one we've heard before. Docker container keeps booting up with only KSK-2010 and then everyday when you boot it up, it starts 5011 all over again. And that didn't work.

And then there's an omission in RFC 5011. Akkerhuis pointed it out. It's not a bug in Unbound in my humble opinion. But if you're close to using the new key, less than 30 days, you will never be able to fill that 30 days. So I think that's an omission in RFC 5011. I'd like to be told otherwise but that's the case. And as I understand that Unbound now does the right thing.

So we always knew that these issues were there. But we just didn't have any objective data until now. So we were worried that these bugs and operator errors were possible. Sorry, I'm just repeating a line here. Sorry. We have hired a contractor to try to figure out the reasons for all these misconfigurations. He's well known in the industry. I'm not sure if I can release his name. There's nothing secret here. It's just that I don't know currently if I can release his name.

A progress indicator of the 500. We started with a sample of 500. We have not contacted 39.5%. We didn't receive a response yet from the remaining 48.5 response and the remaining 12% have responded. And that breaks down in this slide. Sorry, the 0.8% of resolved. That's basically a few configurations. One what I said before, KSK-2010 configured manually, didn't use 5011 update and thought that the system was configured to do 5011 so it will fix.

Then the other one is that BIND journal file, it updated the Trust Anchor, could not be written due to privileges being wrong. And then we have a whole bunch of resolvers forwarding to other resolvers. This is basically the forwarding behind resolvers – validating forwardness. And so we would like to know who is then validating before that but we don't know yet.

So back to the plan and process. How am I doing in time? So like I said before, the September 19 DNSKEY response size increased but there was no alarming increase in DNSKEY queries on the root servers. And then we got Verisign's report. We corroborated it with our data. And we asked ourselves what to do. So now we have this signal. And we don't know if the signal is user errors or if the signal is configuration docks or bugs in the signaling, if that makes any sense.

So we want to have a good look at it. But we were close, October 11 is fairly close to September 27. And we decided to consult our operation plan. And the option in operation plan that we have is and I'm going to read this literally, "The Root Zone Management Partners, Verisign and ICANN, might also decide to extend any phase for additional quarters. For example, if new information indicates that the next phase may lead to complications, the current phase would be prolonged. This is referred to as an extend scenario." So that kicked in on September 27.

When this was decided internally, we immediately went to the public. We went on Twitter, we went on various DNS-related mailing lists etc., etc. We got a message out.

So still, we do not know how representative this data set is. Also keep in mind that validators are not end users. These are two different things. And of course the impact on the end users is what's most important. We are working with APNIC with Jeff Huston and his team. They have this Google Ad based system where they can measure who is validating and who is not, etc. and we will align that with our data. Is that the right word, align? That's the right word, yeah.

The problem is that mitigation is hard. We've already done this for two to three years. And these implementation specific problems don't make it easier.

Next step, this is important. We postponed the root KSK roll until we have more information and understand the situation better. Delay will be at least one quarter. That at least one quarter has to do – that means one quarter or more. That delay has to do with the way our DPS work, DNSSEC Policy Statements. Basically we have key ceremonies every three months and we need to arrange these keys and signatures through a key ceremony.

We will at least partially mitigate. Like I said we have a contractor looking at the 500 resolvers. The 500 resolvers I think

that translate to about 100 or so autonomous system numbers. Another message I've heard recently, should we remove KSK-2017? No. Don't remove KSK-2017.

Anyway, that's it for my presentation. Before questions come up, also tomorrow DNSSEC deployments, there will be this presentation. But Duane will have more information on the – sorry to put you on the spot there – will have more information, far more slides than I have on the glass and how he came to research all of this. Anyway, thank you.

EBERHARD LISSE: Okay. Questions from the floor? Paul, have you got a question or are you just – Paul Waters will make his way to the mic.

PAUL WATERS: I just wanted to tell you that I'm not sure if you know the exact dates when you – if you can see it in your monitoring. But on August 1, Red Hat pushed out all the updates for the new KSK. So maybe you can see how fast people update.

ROY ARENDS: Well, there is a signal that we like to public about. There is a bunch of implementation and this might actually Red Hat, we don't know. There's a bunch of implementation that have

automatically configured Q3ZK. And we don't know who that is. It's a fair amount of implementations. So it's not just a user doing something on a script. I'm sure it it's not Red Hat. You guys know what you're doing. But I'd love to talk to you about the signal and we'll look in our data if we can find your signal. Thank you.

EBERHARD LISSE: Okay, thank you very much. My view on this is that most people wouldn't have understood if something had gone wrong. Why? It doesn't work and it's not squealing preferably if they can't reach Facebook anymore. But it is obviously the best thing to do is to be very, very careful on something that must not break.

As you all know, I'm a gynecologist. I like to when I'm operating to try to leave my patients on the table. So you have a plan and the same thing here is you don't do anything unless you really, really, really – to all engineering standards that are available – sure that it's going to work. And if you're not sure, it's better to postpone and gain more security before you do something.

Thank you very much, Roy. And now, we have the host presentation.

HAMED: My name is Hamed. I'm from .ae [inaudible].

EBERHARD LISSE: Can you speak a bit closer to the microphone please. Sorry, it's my Freudian slip you know.

HAMED: Can you hear me now? My name is Hamed. I'm from .ae domain administration. I like to welcome you to UAE Abu Dhabi. My presentation is short. It's just an overview about .ae. And it's a quick overview about .ae as we established in 2007.

[END OF TRANSCRIPTION]