PANAMA – DNSSEC Workshop (2 of 2)
Monday, June 25, 2018 – 10:30 to 12:15 EST
ICANN62 | Panama City, Panama

RUSS MUNDY: Anyway, we'll go through the presentations in that order. If there's clarifying questions, or a few short questions about each specific presentation, we'll take those at the end of each presentation, and then take overall questions for the full panel at the end of the panel time. So each panelist is doing about 15 minutes of their individual presentation' that will leave us with 10-15 minutes at the end, so with that, over to you, Matt.

MATT LARSON: Thank you. Good morning everyone. I'm Matt Larson, VP of research in the office of CTO at ICANN, and I am here to give you an update on where we are with the root KSK project. How do I? What do I do? Yeah, okay next slide please. So here is the schedule we have been operating on since early this year, post the decision to postpone the role and you can see where we are in the June ICANN 62 and we were going to hold another session for community feedback, this turned out to be that session, so thank you Russ and the program committee for asking us to present on this.

So, we are going to talking about the KSK rollover several times throughout the week, so if you like these slides, you probably get the chance to see them again. In May, the board asked the SSAC, RSSAC and RZERC to comment on the plans that we had revised and, so in August 10th that's when we hope to hear from those committee's with their feedback. The plan then going forward is in mid August to publish the final plan that would be contingent on rolling the key pending and porter's resolution. Then that broad resolution we would hope happen mid September then of course the date everyone has heard me say over and over again, October 11, 2018 is when the root KSK will roll. Next slide please.

Going backward a bit, we took public feedback, community feedback, to revise the plans to proceed with the rollover and on February first that's when we publish them and just to review the plan that we called for was to roll the root to KSK on October 11th and in the community discussion when we asked for feedback on measurable criteria we didn't really get anybody suggesting anything specific. People that say we should keep talking about it, and they say whatever day you have, you should let us know.

So, next slide please. So, that plan went out for public comment. Public comment was opened from the first of

February through to second of April, we received comments, I will categories them as largely supportive, not overwhelmingly supportive, but largely supportive there were couple of less positive comments, and we published a public comment report on April 23rd and then we updated the plans that had been in place, they are all updated now, assuming in October 11th,2018.

If you could, next slide please. So, everybody here probably knows what RFC 8145 is and you are going to be hearing about it from Wes and Joe later on in the segment anyway, and the ICANN org is receiving RFC 8145 Trust Anchor reports from 11 of the 13 roots of letters and we are publishing high levels of graphs that are updated daily. We are also publishing the source IP addresses, not all of them but the sources that are saying that they only know about KSK 2010 the current key, and the word sources that are reporting that they are not ready for the KSK roll. And information about that is on that same page.

Next slide please. So one of the graphs you will find on that page looks like this, we have them broken up by individual root servers, but here is the graph for everyone going back to as far as we have data. I should add I am not showing it on this graph but over time additional root servers were added but they don't show on the graph exactly when things added. Note that there are two different y-axes here, let's start with the one on the left, that's the number of sources that are reporting trust anchor data that those would be unique sources in a 24-hour period.

And the green would be the total number reporting and then the red would be the total number reporting they only have KSK 2010, and in other words that they are not ready. So, if we look at the far right you can see that we are coming up on 180 000 unique sources per day sending in trust anchor data, but if we go down then to the red line, as of right now about 20000 of those, if I am reading little less maybe, are reporting that they're not ready for the KSK roll because they have only KSK 2010.

So, then if you then divide the red line by the green line you get the black line as a percentage and that would be the right hand y-axes, and those are the percentage of sources that are reporting only KSK 2010, and therefore that would be the percentage that we would expect to not be ready on October 11th.

So you can see there has been some activity the percentages has gone up and has come back down and I am going to leave that to for Wes to talk about, because I think he has cracked the code for probably why a lot of that spike happened. I guess I will point out that might take away from this graph we have sort of come full circle, we are kind of back to where we were in the fall of last year. But I think we have a much greater understanding of what this data presents, I think we have little less confidence in what it is trying to tell us and how we are presented.

Next slide please. So here is what else we have going on, we are continuing the communication we have been doing for the past few years at this point, including public presentations at large events to try to get the word out. It is something we are doing that we are announcing for the first time this week because it's only kicked off recently is that we are preparing KSK rollover readiness survey.

We are going to engage a professional survey firm, and we are going to contact the top 10 000 ASN's worldwide that show evidence of DNSSEC validation and that's coming from APNIC data which in turn is based on their Google add network research, so thank you APNIC and then through the transit property thank you to Google as well for their support of APNIC's research. So we are going to attempt to contact by email and certain ASN's by picking up the phone and calling as well.

We are going to try and let them know about the KSK roll and get a survey if they are ready or not. So this is attempting to actually contact individual ASN's. Joe will tell us how much fun that is and how successful that is, but it's something that we think we need to do, so we're going to do it. The other thing we are doing is kicking off a research project, documenting exactly how popular validators react when the trust anchor changes from what's configured.

So for example, we want to discover timing of one validation failures started happening, so what want to know is come October 11, 2018 if someone is not ready, how does their traffic pattern change and we want to be able to looking for that app that roots servers, so we can understand the traffic that we are seeing and know for example if we should anticipate spikes in traffic.

And then we are also going to answer any questions, are answering any questions from SSAC, RSSAC and RZERC as we hope that they are preparing comments before the August 10th date requested by the board for their feedback on the set plans, and I believe that is my last slide. It is my last slide, so with that I would take any questions.

RUSS MUNDY: No questions for Matt, I guess folks here have seen this several times, so --

MATT LARSON: You can all give the presentation.

RUSS MUNDY: It would be a good explanation. One more call for [inaudible] specific. Okay, next we have Wes Hardaker.

WES HARDAKER: Thank you, Russ. We will have slides in a minute, I am Wes Hardaker form the University of Southern California information science institute, and that's the last slide. That's the first slide, and that's zoomed probably because I think they are wider than that. Why don't you go to the next slide then we see how off it is? Yep. Can you zoom out a little? That might work, I mean we can find out --

Alright, so really quickly, I am going to talk about for the KSK rollover plan and I will illiterate a little bit in what Matt said and skip a lot. And then further some problems with it and a case study of what I looked into some date to figure out what we could actually study, what we could actually measure about who had some of these older keys and then I will talk a little about the impact of success.

Next. So, as Matt jus concluded this is basically the same timeline that is part of the critical plans in October of 2017 a new KSK was generated and it was put into the rogue zone in July of 2017. I said 2017 a minute ago. It was generated in 2016. As we referred to in the rest of these slides of KSK 2017 and the older one I am referring to as KSK 2010, so you understand those dates, because that is when that one was put into place.

So, the expectation was that we would roll the key in October of 2017 with the three-month period between initial publication and the time where it would be put into operational use. But on September 27th ICANN wisely and Matt's team wisely desaid it to sub the rollover plan, because there was a whole bunch of questions regarding, you know people that was still signaling only the old key with RFC 8145 data. And then the new expectation is that this October is when they KSK 2017 key will be put into operational use.

Next. So as Matt's already talked about this, this graph there is a couple of important things I am going to add to it though. You know the black line is bad right, so black line went horribly up at one point, it got up to near 20%, and it was sort of this jump that caused me to really go and figure out what was going on and I think there is more work to be done than what I've done in my minimal research I am talking about today, and this graph is actually from the DNS-OARC talk that I think Matt or Roy gave at DNS-OARC-28.

And, so at this point it was still going way up and the big spike hadn't occurred, nor had the drop. The important thing to take away is that you know the black line is a percentage of KSK 2010 trust only, so they were only trusting the old key so it's bad.

Next. So, I sort of had this question, you know, why were so many new addresses regularly occurring, why were like new sources that we had never seen before regularly appearing and sending out 8145 signals with only trusting the old one. Why were there new addresses only trusting an old key, didn't make a whole lot of sense. So I kind of wondered -- because a decent dive into data analysis give us any sort of reason.

So I analyzed two sets of data, and analyzed the ICANN 8145 data I thank OCTO for providing that data to me, so I could look into it. And you know, that is 20 million records of 1.1 gigabytes of data, and I took all of the B-root data from UFC ISI's B-root for a month and looked at that and I did some cross correlation between the two. And that's a decent amount of data, it's 2.8 terabytes of data, so it took a while to chuck through it.

Next. And because the data's so large, the first thing I went to do is reduce it down to sort of just studying a slice of I that might be helpful. So, I took all of the unique sources that were in the 8145 data, which was 1.2 million addresses, and I looked you know from them, only the ones that wa signaling 2010, which was another 500 000, and then only sending one signal, one of the things I noticed is not only were a lot of these addresses, you know, signaling and coming up new and signaling with the old key, they were only doing it once out of three months, they were only doing it once out of three months, they only sent one 8145 query which was just plain odd considering they were suppose to send it once a day or so.

And so I compare that with all the sources in 8145 that also sent their request to B-root, and the reason I narrowed that down is, I figured it was more likely that if they sent there one query to B-root, it is more likely we would have seen other stuff right, if they had sent it to some other root and they were never talking to B, because we were more latency, we have a smaller root than some of the rest of them, I figures if they send it me it was more likely I would see other traffic from them as well.

Again, mostly just to limit the quantity of data I had to search. And then again, I searched down to only those that were signaling 2010 KSK, KSK 2010, and then only those that sent a signal, and so that was down to 16 000 thousand, and then finally I limited it to how many other queries they were sending. So in the entire month, I was looking at 6700 unique addresses that sent a single 8145 query and only 29 total request, one of them had to be KSK 2010 for ignored, you know a single queries because that's the only thing they would have seen. Why were these supposed resolvers spending out saying only one query, a couple of things and then quitting? Right, that didn't make any sense to me.

Next. And to look into this, if you look at the total number if queries sending anything, 63% of these sources in 8145 data, sending KSK 2010, 50% of them were only sending you know 2 queries.

Next. So, what I wanted to know is the a commonality between these last 6700 addresses, there is something that are unique, you know that makes them different than anybody else, and to do that I looked at you know all the data in March and I tried to see if the is a commonality in the QNames they were sending so they sent a 8145 data, they would send other stuff and I wanted to know what's the commonality about the other stuff.

Next slide. So, I counted the top queries that they were sending, _ta-4a5c is the 8145 query for KSK 2010, and then the next most common was the root they were asking you know for DNS keys and things like that from the root. And asked for records and things like that you would expect form a resolver, and then the next two top ones were for a VPN provider, a virtual private network providers domain, and I thought well that's interesting and not only that, the next top one were also a alternate one for the same provider, I am leaving it anonymous here just so we don's shame or blame the company, because they were actually quite responsive as I will show you in a minute. But clearly this to me and the kid that I might have found a unique source to the problem.

Next. So I downloaded there software and I extracted out of the their software and I did a string search for basically the SHA256 hex code for that KSK 2010 key and it revealed in there a root.key file that contained only that key, it didn't contain the 2017 file. So I found that interesting and along with that the other packaged files also contained the LIBDNSSEC.SO which is a shared library from the DNS unbound resolver which is what the bottom bullet said, it's kind of clipped. So, two things you know happened, one I found the key, the 2010 key in a file and I found a DNSSEC validating library, this clearly proves that I likely found the cause of a lot of our pain.

Next. So, I reached out to the vendor and I actually went to bed at 12:30 at night after founding this problem and writing to OCTO and his team and said I think I found one of the sources and was wiped out and went to bed and I got up the next morning, and not only had OCTO actually discovered the right contact for the company for me that ICANN directly reached out to them, and had it all set to go, so I wrote them you know at 6:30 the next morning after not sleeping much, and they responded you know within a hour, and said ohh, that does look like it's a problem, and they agreed to two things, one it's a problem and two that they will be releasing new software in a couple of coming months just to fix this problem.

Next slide. So the result is that in the graph that Matt showed earlier, that major downward spike that happens at the end is the result of the first software released from that VPN provider happened near where the blue arrow is pointing. And we're actually still avoiding, so the VPN provider has multiple software packages on multiple platforms and including android/iOS releases, and I checked, they said that the android and iOS to be coming out really, really soon, unfortunately I checked a few minutes ago, and the android version, I didn't check the iOS store, but they are probably waiting to ouch the android button at the same time that the iOS release process goes through. So hopefully that will come out soon and we'll see another drop in that black line, how big I don't know, because of automatic updates I expect it to be a fairly steep drop and then level off very quickly, but it will be another bump downward.

Next. So, a couple of conclusions, one, rolling a public trust anchor key is challenging, I think we have learned that lesson, tracking down misuse of keys and you know 1 million plus sources is really challenging, there is a lot of resolvers out in the world and figuring out why some of them aren't updating, is extremely challenging, and though I solved a large percentage of that black line, a couple of them weren't points right, there were only user behind each one of those you know addresses I found, it's not like a resolver with lots of users behind it, it's only one.

So, even though the black line dropped a lot, it's worth noting that the percentage of users affected is not equal to the percentage of the black line dropping, and so resolvers that are sending, have a lot more users behind it will be much more affected than to a set users, and this case all of those resolvers where VPN providers, there VPN would have stopped working but the rest of their internet connection would have been fine, it was only the VPN that would have failed for them in the first place, so single application.

And then of course these users, because a lot them were dynamic IP addresses, that's why we were only seeing it once, you know their cellphone or their laptop in a coffee shop or whatever would spin up, get an address and then go away, so that's why we were saying so few queries coming from them, and they would nothing else for a month, because they'd have a new address next time that they put in their software.

So, there's still a huge number of people behind more heavily used resolvers that could cause a problem. So, my bigger take away from doing a lot of this analyses is that and not to sound to negative but, I think if we enroll the key in October, you know will affect the DNSSEC isn't massively widely deployed yet, so we will certainly affect real people around the world, and in a negative way, and I think it is time to rethink our strategy for rolling the key and how we are doing it, because I am thinking we doing to it generally to frequently, five years I think is probably too frequently, I've done some analyses in comparison with TLS type certificates and we're doing quite differently than them and I think we are missing out lessons learned that we could apply to the DNSSES key rollover.

So I think we are doing it to fast and we don't really have a suitably deployed automatic update mechanism which is RFC5011, as well as software pushes and things like that. And I look forward to being on the next design team, because I think it's time to do some rethinking, so whether we go forward or not at this point, and whether ICANN decides to do that in October, I look forward to -- in November started thinking about the whole process again of where do we go from here. Thank You.

RUSS MUNDY: Thanks, Wes. Do we have any specific questions at this point for Wes about his presentation? We'll do an overall questions later.

LARS LIMAN: Lars Liman here form Netnod here --

RUSS MUNDY: Tap the mic, I am not sure it's on. Front one works.

LARS LIMAN: It's on, but --

RUSS MUNDY: Yes, there we go. Kiss the mic, it's a very close one.

LARS LIMAN: I am too tall for that. Just one remark. I would actually advocate the opposite approach, which is we roll the key to -- Yeah, but it took 30 seconds, I would argue the other approach, which is to roll the key more often, because that's the way to force people to do automation, I can understand the argument that there is no cytological reason to do so, but just to get the NIC isn't working I think the way forward is to actually roll it, not to wait longer intervals because that will only increase the impact when it happens.

WES HARDAKER: No, that's a valid point and that's for other reasons why people like to roll their key on a regular basis, it's the problem of doing that with trust anchors that are distributed and devices, it's hard, and automation is the only to make that succeed, you are 100 right, and not only that, the automatic SSO one; what's it called, easy -- you don't know either. So the less encrypt process is sort of proven that automation is one potential way to go forward and if you can get everybody to do it, it's a useful solution.

RUSS MUNDY: Duane.

DUANE WESSELS: Thanks, Duane Wessels From Verisign. Wes, can you say a little bit more about you think would've happened on the roll of the day if this VPN didn't fix their software.

WES HARDAKER: Well, my hypothesis, right, I haven't actually run the software to see what would happen, but my hypothesis is that their computer continue to work and so one of the things I didn't mention about the VPN software, so once VPN comes up, all future DNS request goes to the VPN, that's why we wouldn't see the data, but basically their VPN would refuse to start because it was unable to the get the VPN end point addresses which is really why it was spinning out, it was saying: hey what's my VPN host I can connect to, and that would have failed, so the VPN software would have said: I am sorry, you know, we were not in a network in use.

DUANE WESSELS: Okay. So those users wouldn't have maybe known that it was a DNSSEC issue, their software just wouldn't ever --

WES HARDAKER: So I don't know what the error message says, as I said I haven't tried it, I mean that is a really good question as is what I should do right? Is actually go block DNS and see what actually the software does, and I haven't tried that.

DUANE WESSELS: Okay. And maybe a question for Matt, you said you were doing some tests on the science, is this something, are you adding this particular thing to your tests, are you going to test like this particular VPN software? Are you able to test it?

MATT LARSON: It hadn't been on the list, but we certainly can.

DUANE WESSELS: I'm just curious.

RUSS MUNDY: One more quick question for Warren and then we go on to Jo's presentation and then we do the overall questions afterwards.

UNKNOWN SPEAKER: Mine can be an overall one.

RUSS MUNDY: Okay you want to --

WARREN KUMARI: I guess I had two questions, Barry, what you helped Larson but you don't helped me? Okay there we go, so Warren Kumari Google, can you go back to the slide with the pretty graph? Yes, that one. So, brought me a laser pointer, how geeky is that, so how come you have a very straight line over there and then it has a sudden drop? And then you have the long tail sort of going off, we have not really explained that, have we?

WES HARDAKER: Matt, do you have an answer to the big spike?

MATT LARSON: We don't.

WARREN KUMARI: Okay.

WES HARDAKER: I had meant to look into it myself and I hadn't, somebody did and said that it was weird and you know the addresses seems to have not much correlation between them, but I haven't thrown my data analytic skills at it.

WARREN KUMARI: Okay. And, actually I guess I'll do the same as a Wes question. So you have a conclusion that we are doing this to frequently, to fast blah blah blah. I kind of agree with Lars that the frequency is reasonable or possibly even faster because of automation. But I think the too fast bit is potentially were we serve agree.

WES HARDAKER: Yes, so --

WARREN KUMARI: I think this one does need to happen now. Then after that it should be --

WES HARDAKER: I had a longer presentation in mind when I originally put this together along with conclusion, I have not research to sort of back -- how we are different, and DNSSEC is challenging because we switch from one key to the new on a flat deck. And that's one of the issues with the way we doing it now, and it's understandable why we are doing it that way, but -- consider that last bullet as a preview for another talk to come.

WARREN KUMARI: But we actually have, it's not well deployed, but we have a rollover mechanism which makes us different. But --

MATT LARSON: Can I just add that we really don't have a frequency because we have done this zero times, and we talking about doing one time, and even though we no frequency, I think what happens after the KSK roll, if it happens on October 11th, 2018 is completely up in the air, absolutely up in the air, I think that's we need have a discussion with the community about that, we have to talk about a algorithm roll, if that's something we should talk about, I really, I hesitate to use the word frequency at all., because there is no timing, no schedule whatsoever for the roll beyond the one thing talked about.

WES HARDAKER: Does the current ICANN policy state every 5 years, or after 5 years?

MAT LARSON: It's after 5 years.

WES HARDAKER: So it basically only talks about doing it once.

MATT LARSON: And the DPS of course can be revised.

WES HARDAKER: Without change definitely.

UNKNOWN SPEAKER: Quick comment actually, have you guys actually saw in the past how those TLS clients did the trust anchor rollovers, that was pretty bad. Some old browsers don't follow the new ones, not updated so I guess we are doing it pretty better than them? I know these days they publish trust anchors with 10 years in advance and roll after five years it's published, but you know --

WES HARDAKER: That is what's in the other 15 minutes that I didn't present today, so I have more data on that, that will come out in the future.

RUSS MUNDY: Okay thanks, thanks for that good discussion there, one of the things that Jacques and I just had a little short conversation about, is this is -- we think the most important discussion with this panel we have up here right now, we're going to give the rest of the time to the panel if need be, whatever it takes to get as much as discussion we can generate, so that gives us an additional half hour if we need, so next on out panel of presentation is Joe Abley whose going to talk about some of the data collection that was done. Joe over to you.

JOE ABLEY: Thanks, Russ. Just hang on for the slides. Actually one thing I will say well that's the last slide again, and that's also zoomed. I got too much text on some of these slides, so it's probably worth trying to un-zoom it a bit. So one thing I will say while the zooming is happening, this is a little presentation about some work that ICANN gave me to do as a contractor at the end of 2017.

And at that time I was doing a bunch of little contracts, and I have since then started doing a big contract with the Afilias, so I quit often have a little affiliation with Afilias, but at the time of this work I didn't, so that was why I was a bit vague with my affiliation earlier on so. This is not Afilias work, although I am currently working with Afilias. That's really zoomed in even more. Or we can give it a try.

Next slide then. Alright, there we go, that's better. Good. So I did this talk about a little of this stuff, because very helpfully we've had all kinds of conversation about what 8145 is, the one detail that wasn't in Matt and Wes' presentation, is probably because everybody knows it, is 8145 specifies a couple of different mechanisms, but the one we're really talking about, and the data is saying we are talking about, are the ones where they send queries, where the queries themselves embedded within the labels contain an indication of what a local trust anchor still looks like.

So that's what we are talking about, we are not talking about EDNS options or anything else that might be specified in 8145, we're just talking about this collection of queries, as received by the root servers, because that is the trust anchor we are interested in. So I'm not talking about any of the details beyond that really of what 8145 is or anything else, or how it was collected, you already heard good descriptions of that. This is just a description of my experience trying to contact actual people who perhaps had something to do with some of these queries being sent, so I could ask them why you doing that, so that was basically the entire exercise that I was given to do, and that is what this presentation is all about.

So, next slide. So, which is what I just said, so next slide after that. So, what was I given, I was given a list of data, which I got from Roy in Mat's team, and there were approximately 500 IPV4 and IPV6 addresses, 500 addresses total, most of them V4, some V6, that had been seen to send or to source 8145 queries, as seen by the root servers, so the package that were received by the root servers had these addresses, that's not to say they were the systems that originated these queries, but those were the ones that sent them from the perspective of the root servers.

And 500 is attemptingly a small number, which is why I was asked to do this, because that seems like the kind of problem, and this was back in 2017 before the number got bigger. Before it was seen to grow in this disturbing way, so 500 seems a little, put a little work on the phone and we can solve this entire problem make it go away in a month, let's try that which is what I was trying to do.

So I had a rough window, it was September 2017, I forget the exactly what root servers were responsible for the data, but it doesn't really matter, and if I particular questions about the data set, I needed more details of course the very helpful OCTO people were standing by ready to answer questions. I think at that stage, it was still a open question as to whether there were privacy implications in some of these data, so there were no great enthusiasm for making the entire set available, if they could avoid it, which I think is sensible. But certainly if I had more questions about specific times were a query might have been sent, I could ask. So lots of good support there.

Next. Next slide. So Paul Hoffman actually gave me a bunch more starting suggestions, so this for example is the workflow, hopefully you can see this from the screen, so you know to determine the contact information, trying to contact them, ask them if they are aware of DNSSEC. It's like a script for a telemarketer. And this was kind of the idea and at various points during this, there were counters I can increase till I get a response from somebody good, tick that box that, put it in that category, did they seem you know, angry. I have to put people on a do not call list.

You know this is things that could have come out this the script and the next slide, which probably much smaller and harder to read, but if you are interested you can see it was a suggestion taxonomy for how do we classify these individual contact attempts and trying to classify what it is we learned from each one of these things and the idea is to obviously try and have some systematic representation of the data at the end so that we can say you know 80% of these people are in this category or which means they would have experienced harm if the key roll had happened or perhaps they would not have experienced harm, and that's good to know as well.

So that was kind of the starting point for the whole thing, so the next slide, this all seemed plausible, and I feel good, this seems like good easy work for me to do towards the end of the year, Christmas is coming, this is fantastic, but as you could tell from the title, I didn't had as much fun as I had imagining I would might have. So, the first part was relatively straight forward, is find some contact information for these addresses, or in some cases for the AS numbers, the autonomous system number that corresponds to these things, ICANN WHOIS for this, because we all know that all the information you need and the WHOIS, and there is no issues there.

And WHOIS is of course always very accurate and always very up to date. And it's always updated as soon as soon as someone leaves the company, they immediately change the WHOIS that is true, that always happens. In some cases I know people, in fact there were some examples, like Rob Seastrom for example, I found one of the , oh that's Rob's server, and I just call Rob on his cellphone and say Rob what is going on with that, and he said yeah I know I fixed that, and I said thanks, so that was a easy one.

And in fact, the WHOIS data you'd expect to be stable over a long time base, are the kind of contacts who in fact have been doing the same job for the last 30 years, and those people work at universities. And when you call the university contact and say I've got this weird question about the DNS, in fact you get straight through to the individual and it's cellphone number is there in the WHOIS, and it has been for 30 years, this number hasn't changed and you can talk to him and he has been doing the same job for 30 years, and he knows exactly what's going on and it's fantastic and that's not the majority of my experience in this.

That's the very, very small corner of this experience. I thought that because what I was trying to do was really trying to match traffic from an address to an explanation that this in my mind is the same thing as an abuse complaint. If I get a bunch of junk traffic that cause me harm from an address and I want to stop it, then I should call the abuse contact, that's what the abuse contact is for, so I thought that makes sense, and most people also thought that that makes sense, then there were one person, and he didn't and he was very angry.

If I had used his abuse contact to call him about this, because this was not abuse and other I had mentioned because the conversation was a bit distressing and lasted a long time, but -- anyway so that happened. So I guess it's always the case I suppose if I was a telemarketer I probably would have had the same experience as a lot of people that I call are very angry, and a lot of people are very confused and I had similar kinds of experiences with these kinds of things, but also a lot of people who were very helpful.

Next slide. So, I'd like to think of myself as someone who has been around a bit, lived in different parts of the world, but of course I'm just alluding myself because in fact I am model lingual, you know, worse than the English speaking person who is hopelessly unequipped to deal with the population in the earth. So, talking to people in North America, easy. Talking to people in Europe, very easy because Europeans speak certainly better English than they do in North America. Talking to a lot of people in East Asia, difficult.

Depending on the place, sometimes you get a contact with somebody who had a lot of impact, a lot of sort of routine contact with operators in other parts of the world, and they used dealing with English, even with complicated topics like this, but in large the people I am trying to contact, are trying to sort out local issues for local customers whose speak Malay or mandarin or something that I don't speak Korean, and trying to get any sort of progress with those sort of kind of contact, I'm just not equipped to do it, so it was a big section of this data was just hard to get any response from just because I am not sufficiently lingual, multi-lingual.

And also even when you try talking English to people, this is kind of a complicated questions, because they are saying so what is the problem and which of my customers I causing a problem, well there's no problem today, it's a possibility that maybe there will be a problem in the future perhaps pending and I'm asking you complicated question how your non-resolver or your unbound resolver or you bi-9 resolver, and they are saying what, what are you talking about? And it is very complicated, so it is not an easy topic to describe to anybody, never mind getting any useful feedback, so this flow chart that we had at the beginning, which imagined this series of structured questions where we get structured answers back, turns out to be a bit fictional.

Next slide. Back one, that's it. So this is something that I hit with some the most interesting sources of traffic that I saw were things like, they resolve clusters at big ISP's, TELUS was an example where, I forget where the big cluster was or what the people in Edmonton or somewhere in Alberta I think, and they were disparately keen to help and they put loads of people onto it, but it was difficult, because I am asking to say who was the end user the corresponds to this address at this time, and the answer is particularly in Canada, I can't tell you that.

In fact, the supreme court has told me that can't tell you that because the motion picture industry has trained everybody that your IP Address is private and I can't call you up and say did you download a copy Batman the movie, because I am not allowed to know who you are, and the courts says I can't. So, this is an example in Canada of why these questions are hard to answer. The operational path that we have to answer these kinds of questions has kind of been killed by the reactions to things like content privacy.

In some cases people were quite happy to say here is a message, can you give me a message and I will pass it onto the end user, and it worked in some places, in particular in Europe, for some reason there were a lot of people that I passed these messages at help desk and I got responses, and that was actually kind of good, but in lots of cases that means noting, because someone is using software and they are not enthusiasts and it's not on an Xbox they configured themselves, and they don't understand it.

Next one. So this one actually interestingly touches some of Wes' findings, which is that TELUS, another good example with all the effort they put in, they put packet filters in place, they don't normally have a way of tracing what a particular query might have been sent in September and who sent it, because this is the case where queries have been forwarded through a resolver, their resolver are not DNSSEC aware, they just forwarding the query on.

So if I say what was the query from an end user that triggered you to do this, they don't have any data to answer, but they were quite happy to throw all kinds of complicated filters onto the network interfaces and do deep packet inspection and try and find subsequent copies of the query and of course they did not find any, because a lot of these queries were probably, you know, they were sent once by one address and were never sent again.

So, even after the fact it is quite difficult for them to instrument and find the stuff, but nobody has instrumentation, maybe

Google has instrumentation because they measure everything, but I mean, nobody at an ISP bothers to record all queries received and relayed through a resolver, because it is just not data that is interesting. If you can summarize that quickly and make money out of it, I am sure they do, but in terms of keeping it for operational purposes, it doesn't exist, and I guess there are privacy implications these.

The last comment here as I expected to find examples like this but using NATs, where you have a carrier grade NAT, that would effectively provide a single source, for a big giant pile of end users, I didn't see any, or at least I didn't find anybody WHOIS told me that's what they were doing.

Next one. Some of the big cloud providers whose names you can guess, they have infrastructure where it's not easy always easy to match an address to a customer, sometimes customer's have address ranges that correspond to all of to all of the VMs they might ever run. But, in lots of cases these are VMs that have spun up for 20 minutes and have then have spun down again, and the address that corresponds to them has [inaudible], it's mixture of customers, and if you say who was possibly sending a packet from a VM in your cloud on Wednesday in September the answer is yes. Is there a way we can bill for that, no therefore we don't have the data, I think is the answer?

So, it's difficult to know about how widespread this is in terms of the total set of data that I had but certainly all the data sources that I saw from inside Amazon for example, which is an obvious giant cloud provider market leader, unsurprising that they would be in this category, are all from VMs and it's almost impossible to track any of them. The ones that you could track, that have predictable addresses of course, I could find from WHOIS-RT, but I mean so I mean that they're not in that category.

And, next slide. Even if I'd find somebody who seems to understand the question, in terms of actually building this taxonomy and trying to categorize the result it's quite hard, because you would say things, "yeah, I've passed it on" and " you say well okay, what do you mean by it and on and pass, and there is no more detail", so it sounds like somebody might have understood the question but it also kind of sounds like you've been brushed off and there's no way of telling between them.

So, conclusion is the next title, next page and then we'll skip on from that.

So, there's no great surprise here. The DNS topology is complicated, it's not as simple as you have clients sending things to servers and you can identify them all, there's middle boxes at layer three, there's middle boxes at in the form of forwarders in the DNS layer. You have resolvers that have chained to other resolvers, it's really complicated, and measuring the stuff from the far end of the tunnel at the root server, it's impossible to tell where things came from. I mean, the easy cases are the ones that didn't really have a problem anyway because in most cases they'd already fixed it, these are the universities. The ones that weren't going to be fixed, are the ones we've already set up in the lab to test what would happen if the key roll failed, so it's supposed to be just the old key.

So, the DNS is a really complicated thing to measure, and the sporadic nature of the queries makes it even harder. I think Mark Andrews told me byte 9 that's running will normally send one of these queries once a month or something, or once a week, I can't remember it was some big interval, that to be honest that server is probably going to be restarted before that happens anyway. The whole thing just sounds like needles in haystacks, really difficult.

Next slide. We accelerating a bit because I feel like this isn't the main topic I want to talk about, but I was asked to talk about it so here we are. The IP NoC, the NoC, the abuse desk that normally deals with things like addresses aren't used to dealing with things like the DNS, they don't know DNS concepts and so you end up being forwarded on to somebody else who is not acting in a NoC like way, and so that's a problem.

Next one. The things that I was bothering people with were issues that were interesting to us but they were really not interesting to the ISPs, they don't relate to a customer being down, the whole thing is a very vague obscure DNS problem, that might hurt somebody in October 2018, nobody wants to talk about that today, so that's kind of a problem.

Next slide. I should learn Mandarin.

And the last one is just my general slide of philosophy, which is that every pressure in the system, operationally, in terms of protocol, in terms of everything really, is pressures all directing this to be hard because of the direction that we're measuring this thing in. it's like the whole exercise, anyone 45 seems like a perfectly plausible proposal when you read it and then when you actually come to actually try to analyze the data you realize it's kind of impossible to draw any conclusions from the data because every single success factor is contrary to real life, everything in real life is making this hard.

So, I think you know the people, I'm not the only one to come with this conclusion and in fact this has been repeatedly mentioned by the people who are working, like Warren and Jeff

Huston and [inaudible] working on KSK Sentinel, which takes the opposite approach, try and tie the measurement as close to the end user as possible because putting it as far away from the end user as possible is apparently not a good idea. That's it, that's the end of my presentation. Any questions?

RUSS MUNDY: Any specifics for Joe? Warren, you look like you're ready to jump up?

WARREN KUMARI: It's not so much a comment as a statement, somebody needs to buy you a beer, this was a huge amount of really annoying work and thank you for doing it.

JOE ABLEY: To be clear, I didn't do this as a volunteer – yeah.

MATT LARSON: Yeah, it wasn't free.

RUSS MUNDY: Okay, go ahead Wes.

WES HARDAKER: So, Joe --

JOE ABLEY: But, I'll still take the beer.

WES HARDAKER: So, a couple of quick questions. So one, I'm surprised I didn't see a bar chart with all of your check marks, with which things were good and which were bad, and how many did you actually, do you have any percentage for how many you failed to contact completely, or just --

JOE ABLEY: So, I do have numbers about that and David Conrad has presented some of those in the past when he gave interim results of how this process was going. I deliberately didn't include any of the results in this presentation because they're really ICANN's results and I think ICANN if they want to make them public, then that's kind of a ICANN decision, but I mean the end result was unsurprising given the style of the presentation of given the gist of the presentation.

There's a large number of address that I didn't get good information about, there's a large number that I wasn't able to contact, of those I was able to contact there is a large number that I didn't get good detailed information about and there's a small number of people that were very responsive and were and were able to give me all kind of advice.

But I didn't see anything obscured by NAT, I did see a whole bunch of things obscured by forwarders, and the number of actual conventionally run as you might imagine from the 90s system administer our DNS resolvers that are up for a long times with stable addresses, was a rounding error on the total number of addresses.

WES HARDAKER: And, that brings me sort of to the next question, which is, maybe Matt can answer this, of those 500 were they randomly selected, were they selected based on load balancing, you know trying to look for maximum impact, or --

MATT LARSON: So, it just so happens that I have the presentation from the last DNSSEC workshop where I gave those results, so it was B D F and L root traffic for the entire month of September 2017, and it turned out that there were, call it, 12 000 unique IP addresses, 500 which only reported KSK 2010, so that's where the list of 500 came from.

WES HARDAKER: Interesting, okay.

MATT LARSON: And the actual result were only about 20% could actually be contacted and of that 60% were elastic cloud provider dynamic IPs, and some also known to be forwarding. So, it's exactly as Joe said, we were down to very few that Joe could contact and find anything definitive about.

WES HARDAKER: Yeah --

Can I just say, even that we did pay you, thank you very much this was difficult work and we appreciate it, it was very insightful.

WES HARDAKER: And --

RUSS MUNDY: John --

WES HARDAKER: Oh, I had one more quick one --

RUSS MUNDY: Go ahead, Cristian.

CRISTIAN HESSELMAN: It's not too much. I'm Cristian Hesselman, I'm with SIDN, the registry for [inaudible]. We also took that list from ICANN and we took out the Dutch IP addresses, Dutch resolvers and then we mapped it, we compared it against the number of queries that we received on our DNS infrastructure and now we're currently contacting all these people that are operating these resolvers and the difference I guess with your situation Joe is that of course we speak their language it's all Dutch and most of these people we know because they're registrars and hosting providers and that's our thing. So, it's a comparable exercise but I guess in a more homogeneous environment. I wanted to share that for your --

JOE ABLEY: I think that's great, thanks very much for that and I think it's fair to say my mandarin comments, these extreme example of the same thing when you have personal contacts operationally that you've worked with before, it's much easier to know who to cal and even without the language barrier so I think if there was another exercise like this or there were channels that one could set up to make it easier to talk to resolvers, distributing this amongst operational communities is much better than trying to do it centrally.

CRISTIAN HESSELMAN: Yes, I agree.

WES HARDAKER: So, I guess my last question Joe, is it sounds like in some cases you were told, you can't get contact information beyond the AS contact information or whatever it was you were looking at. That seems to point that there is sort of a end user or I mean, I'm curious about that layer, why you're being denied you know access, why you were denied contact information, when normally I would think that would go to a ISP that should be easy to find the end person.

JOE ABLEY: So, in the case of people who are access providers and the people that I'm trying to reach, like home users. The barrier that I ran into is that the ISP doesn't want to reveal the identity of the home user or tell me their contact information, they're happy to relay the question to the end user in some cases, and in some cases that worked and they contacted me back, their contact information is not in WHOIS, they have one address or they have one /48 at home or something.

And the similar ones where people have VPSs, they've go their own personal Linux box running in someone's cloud at some point, and in those cases the people who are good real enthusiasts have reversed DNS, they have a domain name which itself is another source of contact information and I was able to contact people that way. But in some cases again they didn't, they didn't have addresses that were in WHOIS with the end user who is responsible for that Linux box, it's within a subnet with 1 000 Linux boxes or virtual machines, and the ISP is not willing to tell me exactly who operates this particular one.

RUSS MUNDY: Warren.

WARREN KUMARI: So, Warren Kumari, Google. So, Joe's presentation was fascinating and I really enjoyed it, but does it help us get any closer to how will the key roll go or what will happen you know in late October? We started off on that with a bunch of presentations and this and I think we might have got distracted by the shininess of the last one. Wes had some conclusions on whether we should go ahead, are there other views on that?

JOE ABLEY: So, I have an opinion. I could say my conclusion to this exercise, I could say I didn't find a large validating resolver with many end users behind that was in a bad state, however, I also didn't really find very much so I don't think that that conclusion is worth anything.

I think my conclusion through all of this is that there was not signal, from this data that I could discern using this particular weird manual analysis technique, you know I mean Wes found a lot more signal in the way that he concentrated on very small senders, and no doubt there are other ways of dicing this data and finding other things from it, but this particular approach, this manual telemarketing approach, didn't yield any useful results I don't think.

WES HARDAKER: I keep remembering a question and forgetting it, oh yeah now I remember it again, sorry --

MATT LARSON: Can I answer that question first?

WES HARDAKER: Yes please.

MATT LARSON: So, first I'll say I think we should go ahead on October 11th. We are operating with less data then we would like, right in a perfect world we would know the status of every end point and how many users they had, and you know we could make this decision with perfect knowledge but we don't have that. But using the 8145 data as a proxy and even not knowing how representative it is and suspecting that it's indeed not, we just have no indication that any of those that are reporting the old key, have a lot of users behind them, that's just not what we find.

Jeff Huston's done some interesting working in intersecting his data with that data and he came to the conclusion based on some extrapolating -- well, for one thing, he doesn't see the 8145 addresses overlapping much with his data, and his data corresponds to [inaudible] recursives, right? So he did some extrapolation and his conclusion was 0.5% of internet users might be affected, which is a very very low number, and [inaudible] magnitude lower than the threshold that the design team came up with, we've followed.

WES HARDAKER: So, I get to answer your question too, right?

WARREN KUMARI: Yeah, but I want to respond to Matt.

WES HARDAKER: Okay.

WARREN KUMARI: So, I mean I keep hearing things like 0.5%, which sounds low, but if you actually have a look, that's around two million users, suddenly sounds a lot scarier, so you know just wondered every time we hear the 0.5% number, I start to twitch slightly --

MATT LARSON: Well I don't know that we've heard 0.5% until now, the design team --

WARREN KUMARI: Jeff's intersection from a while back, but anyway.

MATT LARSON: Okay --

WARREN KUMARI: I think, I --

JOE ABLEY: Warren, I think you've got to put that number in context, because I think the percentage kind of is the low number that you're interested in. The internet is not this finely oiled machine where every, exactly, where every day it's like one query gets dropped somewhere and everybody puts their hands up and says stop and find out why that happened because that shouldn't happen again.

I mean the whole thing is a chaotic dirty beast of a machine, it's a wonder that it works at all. So, I think could quite well be two million, you know, contrary to your example sounds like a big number, but does that mean anything either? Perhaps there's four million people on the average day that are down for some entirely unrelated reason and in fact two million is nothing.

WARREWN KUMARI: Yes, but there are lots of people that die every day from heart attacks, that doesn't mean that I can go out and stab someone and add to it right? One of them is something under my control. Well we're really off topic and into delegating the thing.

JOE ABLEY: I think you're saying is it's okay to stab people if they're already having a heart attack.

WARREN KUMARI: Is that on record? So as one of the original authors of 8145, yeah it seems to have a lot more what's than we were hoping. So, a bunch of us have been pushing for a replacement but only what only occurred to me while Wes was doing his presentation, is the replacement thing or potential or at some point replacement than KSK Sentinel also wouldn't have actually found the problem hat Wes had come across; those set of users would be completely invisible or that application because it's not a web browser, would be completely invisible to KSK Sentinel.

JOE ABLEY: Just on the impact of that though, I'm looking because I don't want to monopolies the mic unnecessarily, okay. I mean in that particular example though it sounds like, and again I haven't used the VPN client either, but if what fails as your VPN stops working, then perhaps you contact the VPN provider and say it stopped working and they fix it, because it's actually a bug in the VPN software.

It doesn't actually stop you from using the internet, it doesn't actually stop you from you know, you don't lose connectivity in an un-repairable way, and I think probably what we're going to find is the end user have problems with 1 000 things everyday and they know that things get broken and they have ways of fixing it, they call the teenager in the house who knows how to fix things, or they call the ISP and complain or something. So, I don't know that it's reasonable to say that just because we've found something that might break that user is off the air.

RUSS MUNDY: Wes, did you want to this also?

WES HARDAKER: I don't want to make David stand too long, I'll do it in a minute.

UNKNOWN SPEAKER: On the 0.5%, if I remember correctly the community design team came up with threshold of 1%, as the threshold at which what point if 1% of end users were impacted, that we would roll back. That 1% seems to be a reasonable threshold for not moving forward as well because of if you move forward and you have 1% than you're going to have to roll back immediately, so what's the point in going down that path. So, a 1% threshold, depending on how you measure that could be lots of people or few people.

As Joe points out the reality here is that we're actually trying to probe a [inaudible] chaotic mass and you know it's going to react somehow, we have no idea how, but we're at least, I am reasonably confident that the screams of outrage are going to occur because we did do KSK rollover, will be drowned out by the screams of outrage of you know, ISPs going down, and fibers being cut, and all that sort of stuff. So, I'm reasonably comfortable that the internet won't end as we know it unfortunately.

On the topic of KSK sentinel one of the advantages that I see with sentinel, is that it'll actually provide a tool that end users can use with appropriate encapsulation and explanation to actually be able to establish what their trust anchor is and what's actually functioning for their resolvers.

One of the constant questions I've received every time I've wondered around talking about this KSK rollover is how has, I as an end user, you know make sure my ISPs have done the right thing, and the current answer is well you have to call up your ISP and talk to the people who run the DNS and ask them if they've set up the things appropriately and oddly enough, A that really pisses off ISPs when you do that and B most people can't even spell DNS, so that's not really a option. With sentinel there is significant deployment then we could actually have a tool that people could run on their laptops or cellphones or whatever, and establish whether or not the KSK is what they think it is and whether they will break or not.

RUSS MUNDY: So, this has very much sort of flowed from the questions about Joe's presentation into our overall questions set for the panel, and I'd like to also reinforce that this is the opportunity at this ICANN meeting for folks to ask questions of both ICANN staff and our panelist and make comments about whether or not we should resume the KSK rollover plan. And I think Wes had another question or two he wanted to bring in.

WES HARDAKER: No question, I've been queuing questions to various questions. A couple of things.

One of the issues with 81 45 signaling data is that it's uncoupled from the rest of the requests right, it's a separate signal from looking at the DNS key or looking at the DS record, and knowing that because it's separated it might go to different destinations, so it's very hard to do that coupling. And one of the reasons I started my research project looking into the data was I was curious you now, could with data analysis you know, could I find an answer?

And I believe that I could find more, I'd like to throw, I have a researcher in question on our staff I'd love to throw at, but I don't have the resources to pay her, so if anybody wants to make a donation to the university let me know.

I think one important takeaway is that we've prepublished this key for a year, I think three months was probably too short in the original plan and you know I said earlier that I think that we've been doing it too quickly, that was the too quick element, three months I think in my opinion was too quick to publish and then switch. Now will be a year and three months, and I think that's more reasonable time. If somebody hasn't done it in a year and three months then that's likely their own fault or their ISPs fault.

And so that brings me to my final question that I think everybody in this room ought to think about. Let's say October 12th does come around, let's say your ISP the thing that you're sitting behind goes down a hour, a day, I don't know, how long does it take for them to figure out that they have a DNS problem, probably not a day, probably an hour right, and let's say that the impact of that is that they either fix it or they potentially turn of DNSSEC right, if you think about the answers to those questions then you can answer to yourself should we go forward right. If the answer is okay, if my ISP goes down for an hour it will be lesson learned and you know good for them.

And if they turn off DNSSEC are you be okay with that, they won't ever turn that back on, maybe they will who knows, but that type of thinking is going to get you the point of you know what I could suffer an hour, if we think that this is an important event that we should go ahead and do this key roll anyway.

I don't have an answer to those questions for everybody but I can only think about it in terms of myself.

JOE ABLEY: So, I have a small anecdote, two small anecdotes with that's pretty much the same example. I remember going back a few years now, there was a big failure and it made the front page in Canada. It was a big cable company is Canada, had a coast to coast outage that was big enough that made the front page of the national post and all that sort of stuff. And, it was not very well described in the media because these things never are, but then when you looked at forums and people were trying to help themselves about how to fix this.

Their answer was, "Yeah, you can fix the problem by changing your DNS settings to 8.8.8.8", and it turned out that the problem they had was a cascading failure to do with load balancing and the way that the resolver clusters that really didn't stand up well after a hard shutdown with empty cache. I think that's an example of a lot of end user's reactions to what happens when the internet goes down.

It kind of common knowledge amongst the uninformed nontechnical end user population, that they way you fix your internet being down is you change your resolver to 8.8.8.8 and today there are more examples as 1.1.1.1 and there's 9.9.9.9, but I mean there's a lot more ISPs I think actually are really bad at running resolvers generally, not necessarily the big ones, but Comcast is a good example; they have a whole team of people running it, but there's a lot of ISPs where the resolver's a shitty old box under somebody's desk and nobody cares. And it goes down all the time and end users fix it themselves.

So, I don't know that this is going to be a significantly different situation, if a small ASP fails to suddenly resolve anything; end users already know how to fix this.

LANCE LIMAN: Lars Liman from Netnod. I agree with Wes that you're asking the right questions but I think that we need to try to respond to these questions from different mindsets. I've heard responses from the mindset of an end user, and that's fine and I fully agree with the result of that mental exercise, but what if you are Amazon.com and Alibaba or some other company that really depends on your internet connectivity for your daily business, you will lose millions of dollars potentially.

So, that's a different mindset that you have to set, put yourself in. These are also a group of internet users that we might need to reach out to somehow or that need to be involved in this and to look at how their ISPs work. I will say upfront that I firmly believe that both of these have very good internet connections, the probably do their own resolving and they have everything under full control but it's a different mindset.

RUSS MUNDY: Warren?

WARREN KUMARI: Two things, one responding to Lars. So, yes many of them do, but there're also a huge number of for example government folk, and large enterprises like GEE, and similar. And you know they have a large number of employees and they do a lot of business and having their DNS break, because they don't have the expertise, is also potentially exciting. And, wow apparently I'm very jetlagged because I can't remember my second question.

RUSS MUNDY: Jim, go ahead.

JIM REID: Jim Reid, just to follow from what Liman was saying there before. Yes, you would think that some of these organizations have got critical dependencies on the DNS resolution services but have the correct sort of procedures in place, that isn't necessarily the case. And sometimes lots of enterprise networks, the DNS configuration is not properly documented except all that's happened is that one system administrator was told "don't ever touch that box, because that's the DNS server, if that server ever was down bad things will happen, so just don't touch it", so it never gets updated, nobody ever looks at it, and eventually gets forgotten about.

And the problem is even in enterprise networks, we try to get these things sorted out, and you've got potentially the cloud or corporate IT to say "follow corporate standards or do what you have to do". It doesn't necessarily work because sometimes administrators on the ground, their attitude is, the production system is running properly, the data base is running, the SAP servers are running, the emails are working, the websites working, I don't give a toss about the DNS leave it alone.

WES HARDAKER: I'll add one more data point. So, back 10 years ago, I don't know, I read this application called DNS check and it was actually presented here probably 10 years ago. And the whole point of that application was to determine your ISP and you want to use your upstream resolver and be able to query though it and be able to do validation on your desktop. So, it's just a check to see you know how well DNSSEC would work, and one of the conclusions that I got out of that and I ran it on a bunch of networks, I got a lot of other people to run it on networks, and then I graft the results.

And it turned out that NSEC 3, came out just after the Kaminsky bug, and so bind 9.6, was deployed in mass, with everybody you know magically deciding well I have to upgrade my bind because I've heard about this horrible bug, everybody has to upgrade their DNS software. And 9.6 was the point where the entire world had to upgrade. 9.7 was where NSEC 3 was implemented.

And so in my DNSSEC check application I was actually checking for NSEC 3 support. And I found that a huge percentage of those resolvers that were being measured didn't have NSEC 3 report. And it was very clear that the Kaminsky bug pretty much made everybody update and then they didn't after that.

And I think we're in this magical transformation period where finally DNS software is getting sort of more automatically updated by binder now probably being run out of packages on Linux boxes, then you know built by hand. But we're sort of behind in the curve, in terms of automatic updates go, especially automatic updates with config as well.

So, I think if this was in 10 years from now we would find that everybody is always updating all software right, we don't have this latency issue that we do right now with DNS software. The web browsing community sort of fix this, web browsers self update. We don't have automatic updates on all every platform these days; it's too much of a critical service.

RUSS MUNDY: Okay, I think since the mic lines are -- oh, go ahead Cristian, last question.

CRISTIAN HESSELMAN: Yeah, thanks. Well it's not so much a question, but I just had a look at the autonomist systems that we checked and that we're sending KSK 210 queries and there's a couple of ISPs in there but there's also hosting providers so they also run let's say resolvers, and I'm not sure if this is something that we actually took into account during this discussion. Because we talk a lot about ISPs but there might be hosting providers that run the wrong resolvers as well.

RUSS MUNDY: Okay, well thanks everybody, thanks very much to our panel members in particular and thanks for participation by all our audience members and these sessions are recorded so it will be up on the website for this session in a while, I don't know how long it will take, probably a couple days, maybe next week. But if anybody wants to review what's been said, you can hear that and listen to it again.

We cut the presentation, it's sort of the ending summary but let me just do a summary of the summary and say please, everyone it's here at the workshop, think about what you've heard today particularly with respect to the KSK rollover and take whatever actions you can; don't even bother with these slides, Kathy. Take whatever actions you can that affects your part of the internet and your part of the DNS world, to make sure your part of the world is ready for the KSK rollover. Whether it's in October or whether it's sometime later it will need to be done, so let's have that as our summary of the workshop focus today and turn over to Jacques for our great DNS quiz.

JACQUES LATOUR: Alright, so let's get the -- alright, welcome to the DNSSEC quiz, DNS/DNSSEC quiz. I did this one on the plane on the way here without internet access so, I'm right okay. So everybody you should find the answer questionnaire on your table next to you so you fill that in and then I'll try this. Alright so the rules for the quiz is, I can change the rules, I'm always right because it was late when I did this quiz, we keep it simple, it's one point per question, one answer per question, there is 10, and I'm right so right is right, I'm right that's the way it works okay? So, next slide.

Question one, so a lot of this you should have seen today and the morning presentation. So, which ccTLD is the most recent to be signed using DNSSEC?

A: GW (Guniea-Bissau)

B: AX (Aland Islands)

C: VC (Saint Vincent and The Grenadines)

D: IT (Italy)

Next slide.

Question two, so what year did Panama first sign their TLD?

A: 2013

B: 2015

C: 2017

D: N/A

You're not allowed to search the internet or do any queries right.

Question three, so what does DANE stand for?

A: DNS-based Authentication of Named Entities

B: DNSSEC Authorized Network Entities

C: DNSSEC Algorithms Numerical Entities

D: DNS Automation of Named Elements

And it's also a dog I discovered.

Question four, so the DNSSEC uses public key and cryptography to sign and authenticate DNS resource record sets (RRsets). The public keys are stored in which resource records?

A: KEY record

B: Key Exchanger record

C: IPSECKEY record

D: OPENPGPKEY record

E: DNSKEY record

Question five, does DNSSEC protect against a BGP hijack?

A: Yes

B: No

C: Maybe yes

D: In certain corner cases

I'm the source on this. I think we could have a discussion later on that so you have to think like me.

Question six, would a DNSSEC DANE enabled website be less vulnerable against BGP hijacks with a TLSA record?

A: Yes

B: No

C: Maybe not (which is totally different than maybe yes)

D: in certain corner cases

I'm always right so whatever I think is right is right.

Question seven, so what is the percentage of DNSSEC validation for world, according to APNIC DNSSEC stats?

A: 11

B: 13

C: 18

D: 45

Question eight, what does the 'A' stand for in the TLSA records?

A: Authorization

B: Accountability

C: Authentication

D: Anonymization

That's an easy one.

Question nine, so what is the approximate percentage of second level domain that are signed globally all TLDs?

A: Between 0% - 1%

B: 2% - 3%

C: 4% - 5%

D: 6% - 10%

E: 11% - 20%

That's according to the DNSSEC stat from Rick. Second and, I'm right. Whatever Rick does, so I think it's all second level domain not third level. Dot UK is the highest percentage of signed domain, so you know that's right forget it we talked about it, never mind.

Question 10, which of the following TLD is not in the root zone? That took me a long time, there's like 15 000 in there, I tired to find something interesting here.

A: .uno

B: .uol

C: .ubs

D: .uae

E: .ups

One of them is not in root zone. So somebody has to memorize entire root zone for this. Alright so now you exchange your sheet with you neighbor and then we'll go through the correct answer, my correct answers. Okay, so next slide.

So, it's one point per question and maximum of 10 points and then in the end we'll ask people to raise their hand, we'll start at five and up and then see who is the grand winner. Okay so, question one is Saint Vincent and The Grenadines. They were signed in June. So we cover that in the DNSSEC stat this morning. So one point per answer.

Next one, Panama is not signed. So, I don't know what the plan is but they're out of time.

Next one. Dane stands for DNS-based Authentication of Named Entities. Question 4, DNNSEC are stored in DNSKEY, so, that should be easy. Question 5, does DNSSEC protect against a BGP Hijack. No, it's got nothing to with it. I am right. Question 6, would the DNSSEC DANE enabled web sites protect against hijack. Absolutely. So that's were it can -- but with BGT protocol and all that.

UNKNOWN SPEAKER: [Inaudible].

JACQUES LATOUR: Yes. You're right, but I'm right. There's more yes than cornered keys. Or no than -- right, but wrong. Yeah, I'm right, right? I'm right, I'm always right. So question 6 is yes, DANE will help, Question 7, what is the percentage of DNSSEC validation worldwide. 13%. We talked about that this morning. Question

8, what does the 'a' stand for in the TLSA. Authentication. So transport layer security authentication. You got this right, right. Good.

Question 9. What percentage TLD's are globally signed at the second level, so it is 4%, that's a number. And number, so 10, UAE is not there. Ahh, of course, yeah. It would be a region. Could be. See, I did not know that, well I knew that, but I didn't know that. Alright, so let's start with the next slide. So how many have five or more good answers. Alright six? Ohh Seven? Eight? Nine? Eight? Alright so we have three winners. Congrats. Oh did I say the winner they get to do they quiz at the next ICANN meeting. Thank you.

RUSS MUNDY: Thanks very much, Jacques, and thanks everybody for participating both in our conference workshop today and in the quiz, and now it is lunch time, and this afternoon in the same room is where tech day will be and I know many people that come to our workshops also go to the tech day, so this is where you'll come back to.

But this is your DNSSEC workshop lunch ticket, so don't forget to take that with you, and you have to go out the door, down to your left and then to the left again and the lunch area should be visible there. Okay, thank you very much everybody, enjoy the session, thank you.

[END OF TRANSCRIPTION]