SAN JUAN – KSK Rollover Update
Wednesday, March 14, 2018 – 16:15 to 16:45 AST
ICANN61 | San Juan, Puerto Rico

UNKNOWN SPEAKER: Hello everyone, we'll be starting the root KSK roll session very shortly. Can I have the presentation up on the screen please.

We'll be starting in 30 seconds please.

Hello everyone, welcome to the root KSK rollover session. Hopefully there are people in this room that have not seen me deliver this content already this ICANN meeting, hopefully. Here we are. Let me start by giving you a recap of how we got to today, as I think everybody knows, if you have enough interest to be in the room, the root KSK rollover was originally scheduled for October 11th 2017, but we decided to postpone that. [inaudible] analyzed our CAD 145 trust anchor report data and they found that around 7-8% of the resolvers that are reporting, which is admittingly a relatively small number at the time, had only what we call KSK 2010 or the current return KSK that did not have the new root zone KSK. Something was not right with those 7-8%. ICANNs office to the CTO research team repeated that analysis with traffic from different root servers and we found essentially the same thing. Depending on exactly when you look at it, the percentage is higher or lower. It was still a

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

higher percentage than we were comfortable with, and more importantly, we didn't understand the reason why those resolvers were reporting the old key. We decided to pause the root KSK roll and try to determine why those resolvers had the old key. We took a list of 500 resolvers that in the month of September 2017 has reported only the old key, only KSK 2010, tried to track them down and we found out that tracking down operators based on just the IP address is hard, we knew that, this is existence proof of that. We can only get in contact with 20%, that would be about 100 addresses and of those, the majority were in IP ranges known to host ephemeral things like virtual machines or containers.

There's also an issue with RFC 8145 in that the signal itself is sent as a DNS query, so that means it does things like it forwarded from one resolver to another. We know that forwarding was going on which that obscures the signal, you could be thinking that a particular resolver has only KSK 2010, that resolver could be fine, it could be a resolver behind it forwarding to it, and that reported only KSK 2010.

The upshot here was there was no smoking gun single cause, which would have been ideal, if there would have been one or two root causes we could have potentially talked to vendors to fix bugs had we found them, we could have adjusted our communication message, but that wasn't the case. So, with no

clear path forward revealed by that research, the ICANN organization decided to ask the community for input.

Now we're upto late December 2017. We said that we would accept input and discussion on the KSK dash rollover at ICANN the org list, which is a list dedicated to this project for updates and discussion. I'd suggest that if you are not on that list, please do subscribe. It is very low volume at this point but it's [inaudible] with the project. The results of the discussion, there was general agreement that there really wasn't a good measurement here. The design team that ICANN convened that now has been over two years ago that they made their report upon how to roll the return KSK, made recommendations. The suggested that a good measurement would be number of users affected and they suggested that 1 half of 1% of users affected would be a sign after the root KSK roll that there were significant enough problems to merit rolling back. If number of users, that's a very reasonable measure, but that's also a very difficult measure and that is not what the RFC 8145 data is telling us. There was hope among the people who wade in that there would be better measurements in the future, and they were looking forward to something that is now called Sentinel, which is a draft that Warren Kumar from Google and [inaudible] are working on that would allow user base measurement, but we're not there yet today. The consensus among that group was that

ICANN should roll the key in a timely fashion and keep doing the outreach we'd been doing.

Based on that feedback on February 1st, we published a draft plan, and I want to emphasise it is a draft plan, for the KSK rollover. The components of that plan were first we're going to delay it 1 year, we'll roll it on October 11th 2018, we had hoped that the discussion on that list would yield criteria to measure, we didn't have anyone suggesting any particular criteria. We're also going to continue the outreach, things like this very meeting, publicising the KSK roll and we are going to publish more observations about trust anchor data, mainly the RFC 8145 data, even though we continue to, as each day goes on I have less and less confidence in what that data is actually telling us. If it's really telling us anything that's useful or just giving us an alarming signal that's not really terribly relevant to the project.

An important point here is that we have a public comment period open, this plan is only a draft and it is subject to input from the community. We really want to hear from the community and get feedback on the proposal. That public comment period closes April 2nd and there is the URL for the page that explains it at the bottom of the slide. If we end up proceeding with the plan, the timetable and the plan it is now, and the draft plan, it would be something like this. April 2nd, the comment period ends. Mid April, the staff report needs to be

published, just like any ICANN public comment. We'll also revise the plan as necessary and publish that. If you're now familiar with the cadence of ICANN board meetings, the ICANN board meets 6 times a year, it meets at each ICANN meeting, so that's 3 times. Then in between each ICANN meeting it has what is called a board workshop where it meets. The next ICANN board meeting will be in mid May at it's board workshop and at that point we will request that the board make a resolution asking ASAC and also our SAC, I have not updated the slides, to review and comment on the plan by October 1st. We will likely have another session in Panama to talk about the root case KSK roll. On October 1st, we hope to have received ASAC and our SACs feedback and we will make any revisions to the plan. If those revisions do not call for a different date than October 11th 2018, then by mid August we'll have published that final plan and at the September board workshop we'll ask the board for a resolution directing the ICANN org to roll the key on October 11th.

That's where we are, that's where we're heading. Again, I can't emphasise strongly enough how much we want community feedback and would like people to weigh in at the public comment. I'd like to send the rest of the presentation talking about the signal that we're getting, the data that we're getting via this RFC 8145 trust anchor reports.

At this point, the ICANN office to the CTO has access to RFC 8145 data from, actually this slide is now out of date based on lead breaking developments from 12 root servers, we added H to the list of servers that are providing us data. The initial analysis in late 2017 used [inaudible] data from B, D, and F, but since then we're using, [inaudible] plugin to DNS cap. The root operators run DNS cap and that analyses the traffic in real time and every 60 seconds this plug in batches up a report of statistics that it's seen as well as trust anchor reports that it's seen. It sends them as DNS queries which is quite clever, to a zone that ICANN [inaudible] operates. You can see examples of what those reports look like, there's a timestamp, there's the source IP delimited with dashes rather than dots. The trust anchors are notated in four digit hex number, 4A5C is KSK 2010, 4F66 is KSK 2017. The you see a note ID and a root server ID. Based on all that, we can compile this graph. You can see how as time has progressed we are getting more and more reports from additional servers. Let me explain this graph because it might not be completely obvious at first glance. The three lines are plotting two different things. The red and green lines are plotting number of IPs that are reporting 8145 data, and you read those on the left axis, the number of sources. You can see that over time, at this point, look at the green line, that is the total number of sources reporting trust anchor data. These would be unique sources per day. At this point about 50,000

unique IP addresses are reporting trust anchor data to us everyday. The red line is the number that are reporting they have only the old key. If you divide the red line by the green line you get a percentage and that percentage is plotted on the black line and that scale you need to read off the right hand Y axis, and you can see that we're now at about 20% of the resolvers reporting say that they have only the old key. Note that there is a big spike in January, where we suddenly got a lot more people reporting and where the percentage number got worse. We are fairly certain that that is a result of an upgrade to unbound, unbound 168 was released in mid January so the timing fits perfectly and that release was to deal with a security vulnerability, so our hypothesis was because this was a security related patch people were motivated to upgrade unbound. However, there's no drop off in KSK 2010 reports after 30 days.

If someone upgrades unbound in place and does not run the unbound anchor tool, let me say this the other way, if someone upgrades unbound and they run the unbound anchor tool, they will immediately update their trust anchor store with the right data and they'll have KSK 2010 and 2017. It's common to upgrade in place and not run unbound anchor, so that would mean you have a version of unbound that is still configured with the old key, but if it's properly doing RFC 5011 after 30 days, it should run that process and realize it should configure KSK 2017

as well. Why hasn't that happened? One hypothesis for that is that if these are ephemeral virtual machines or containers you can imagine one of those coming up, having the old key, reporting the old key via RFC 8145, I say old key, I should strictly speaking say KSK 2010. It reports KSK 2010, it runs for a few hours, a few days, it's shut down. It never has time to complete the RFC 5011 [inaudible] the next time it starts up again, it starts all over and it still has only KSK 2010, it reports only KSK 2010 by 8145 and that process repeats. We need to do further analysis at some point, you need to stop working on your slides and fly to Puerto Rico to present them, which is where we are now, one of the things we need to keep doing is looking at the particular IPs, how often they occur, how often the reports occur. We know that bind and unbound make these 8145 reports on a fairly regular cadence. If we see IPs that are not doing that, one reasonable assumption is that they are [inaudible] coming up and coming down. That is not possibly the only reason, but that's why we're doing research to try to figure that out.

Here are graphs for individual root servers, these graphs are the same, as this graph is for all root servers, these graphs are for the, at the time, 11 root servers that we have data from. You can see that the data, different servers we have data from different time periods. If you look at the graphs, they are relatively similar with the exception of J root, which has a lower percentage

reporting, so J root looks better than all the others. At least in terms of percentage. [Inaudible] reports that we are not getting reports from all the J root instances, so that possibly has a bearing on that. The point here is that there is not really vastly different data among the root servers. What I do think is interesting is the behavior change that happened in mid January when we saw that spike. This graph shows unique IP addresses added each day, the data point here represents on that day how many source are reporting data to us we have never seen before. If you look at reading from the left to the right, we were humming along with just maybe looks like never more than 1000, a few hundred new IPs every day. Suddenly after this hypothesised upgrade event in January, we now see many more unique sources reporting per day. Around 15-16,000. If you were to plot a graph with how many cumulative unique IPs we see everyday, that would be this graph. The green line shows at any given point in time, how many IP addresses have we seen. Obviously that number starts small on the left, earlier in time, the further ahead we go in time, the more unique IP addresses we've seen, whereas right now the number is 730,000. To date, we have had 730,000 different IP addresses, I haven't said this is the presentation, it is a combination of V4 and V6 addresses. About 730,000 have reported data to us, of those about 250,000 at one point or another have reported that they have only KSK 2010. The math on this calculation, when you look at the

cumulative numbers where we are today, the math is even worse. It's about 35% of the total addresses we've ever seen report only KSK 2010.

I decided to take a look at this by slash 24, so you can see that the shape of this graph, it is slightly different. There was a spike in new slash 24's after the January increase and then it has tapered off a little bit. If I plot the cumulative though, we still have a lot of slash 24's. If you look at the green line, it's topping off at about 350,000. That would be an average of, we're losing a lot of resolution here, an average of 2 IPs per slash 24. I was hoping this number would be smaller to indicate that maybe there were entire blocks that had a whole bunch of addresses in them and we could investigate those and a good hypothesis for that would be that maybe they were address boxes that only had dynamic ephemeral machines in them.

That is a lot of addresses, what's interesting is that the total number of addresses that are report either KSK 2010 or KSK 2010 and KSK 2017, that number is larger than the total number of unique IPs. What that means is that there are sources that have reported I have KSK 2010, and then they've reported I have KSK 2010 and KSK 2017. Not necessarily in that order, but they have made both reports. One reason that might happen, imagine a source reporting KSK 2010 and sometime later they report KSK 2010 and 2017. One reason, and not the only possible

reason, is that machine upgraded it. It now has the new KSK, but out of this 750,000 IPs only about 1550 report both. That means there's not a vast number of machines that we can point to and say that one hypothesis is that they upgraded. Part of the problem here is that there is really a problem with the 8145 signal, we know there are forwarders involved, as I've mentioned, so just because you see a report from a source IP, you have no guarantee that it is that source IPs configuration that you are seeing. It might be someone else forwarding their 8145 report reflecting their configuration to that IP which then forwards to a root server and we see it. That's another explanation why you would see a single source reporting KSK 2010 and KSK 2010 and 2017.

We also know that there is at least one version of an implementation that reported 8145 even if it wasn't validating. If it had KSK 2010 it wouldn't matter because it was doing DNS validation and using that as part of its configuration. It was configured for it but it didn't matter.

If you want to look at these graphs yourself, we're updating them weekly, that is the URL for them. Root trust anchor reports dot ICANN dot org. I just started to slice and dice the data this way, this is based on this particular table, is not based on the full 250,000 IPs that have currently as of right now, reported only KSK 2010. This is based on a smaller number and I don't

remember what the total number was off the top of my head. This is looking at the number of sources per autonomous system and then doing a reverse sort so that the ASNs with the most IPs reporting... I beg your pardon, this is not only KSK 2010, this is all autonomous system numbers reporting. That would be a different chart if you're only KSK 2010. What this does mean, I need to run this again only looking at KSK 2010 and it will allow us to make some contact attempts to find out what's going on with the ASs that have the most numbers of resolvers reporting only KSK 2010.

We have distributed a list of IP addresses that report only KSK 2010 to the ISPCP and the RARs and the goals are twofold. One of course, to get those systems upgraded but also we're still very interested to find out what's going on any why these systems aren't upgrading. If we find, for example, one of the best results would be to find out there are indeed address space where a bunch of [inaudible] machines that happen to be running an old configuration with only KSK 2010. That would be a positive finding. This slide is now out of date, I do have authorization from the powers that be within ICANN to make the list publically available. We're going to be publishing a page, it is going to look like this, I'm going to sort it in reverse order by ASN and you're be able to click on the ASN and get a list of the addresses from that ASN that reported on the KSK 2010. It's going to soon be

much easier for a given operator to find out what's happening in their network. Obviously, then we will be reaching out to operators starting with the ones with the most reporting only KSK 2010 and trying to find out what's happening. The next steps, keep trying to figure out what's going on with this 8145 data, I'm not happy with the signal it provides but it's as of now the only data that we have. The responsible thing to do is to keep investigating, to try to figure out what it's telling us and if we get to the point where we are convinced it's really not telling anything that is valuable, that's in itself a positive finding. We're going to try to contact networks that are reporting large numbers of resolvers with only KSK 2010, we're going to facilitate other people doing that investigation. We're going to keep talking about this so you're going to keep seeing me and my colleagues and we're going to keep listening to the community because as I've said, ways you can help, please comment on the plan, there's the URL again and please do join the KSK rollover list to stay up to date. With that, I will be happy to take any questions. We have a little bit of time before the next event in here at 5.

HOWARD BENN: It's Howard Benn from Samsung Electronics. Just on that slide you had with all of the operators on, that's the one. My guess is the reliance geo network which is probably their LTE network is

run totally on virtual machines and almost undoubtedly will never hit the 30 day limit. It will be very interesting to see there's a large number of mobile operators on there, whether or not they all suffer the same issues. Again, I'm happy to provide some contact information into geo because that equipment is ours.

UNKNOWN SPEAKER: OK, thank you.

MARK: Mark [inaudible] with the ISPs although speaking on my own behalf here. Can we go to slide number 4? That's the one. I'm looking at the bottom of the slide, where it says results of discussion and boy it's one in three. I'm going to put them in my own words here and you can tell me if I've got it wrong. The first bullet says we don't know what the effect of this is going to be on the internet. The third bullet says, do it anyway. Am I misreading it?

UNKNOWN SPEAKER: That would be the glass half empty interpretation.

MARK: Could you do the glass half full reading for me?

UNKNOWN SPEAKER: I think the rest of this presentation is the glass half full. Let me give the serious answer. The serious answer is that this a hard problem to know what to do. When we got to the point we got to last fall, we decided that we needed to involve the community and get the community's input. The summary you gave is what the community told us. I will point out that it is somewhat self selecting who is on KSK rollover, those tend to be people who are advocates for DNSSEC and have personal capital, and who knows what else invested in DNSSEC and want to see the rollover happen. I guess it wasn't surprising in retrospect that was the sentiment that we got, but that's why the public comment is so important so that we can expose it to a wider audience and if people have a glass half empty interpretation of that, they can comment on it.

MARK: I appreciate that and thanks for that explanation. I continue to have the half empty perspective here. I can't and I feel like I'm an advocate for DNSSEC, so I can't get over and I will provide comments individually on this. I can't get over bullets 1 and 3, there is sort of a cognitive dissonance for me that anyone would actually, seriously propose to introduce a change to the root where you didn't know what the effects were going to be. I agree with you and I completely trust your analysis of this sort of poor quality of the signalling you are getting, but as everyday goes on,

the signalling that you do get which is the only diagnostic information you have, the signalling you do get only gets worse. That's my characterisation of the situation. I find all the slides that come after this very compelling, but it's this slide and bullets 1 and 3 there under results of discussion that I can't get my head around. I don't know why anyone in the tech community would say, go ahead and proceed with the change to the root when you don't know what the results are going to be. Thanks, I will make those comments in individual submissions as well.

UNKNOWN SPEAKER: Let me channel what I think another interpretation might be. This is not my interpretation, I am just repeating what I've heard. Which is that there is potential downside to not rolling the key. It could decrease confidence in the key even though there is no cryptographic threat to it at the moment. Either threat to the cartography itself, or the threat to the physical operational security of the key. Nevertheless, not rolling it could undermine confidence in the key and if you combine that with the belief that it's not going to be so bad and people can fix it quickly and you accept a little that there will be some breakage, you could combine all those things to believe we should roll the key.

MARK: I understand the sort of reputational argument right, is that well what we want to do is roll the key because we want people to have confidence in DNSSEC. I understand the reputational part of that, on the other hand, if people are guessing wrong and I am shocked if it is down to guessing, that though it actually is worse than what we thought it was going to be and we have to execute a rollback, right. We haven't one that yet in the live world. I think we're in a reputational problem either way. That puts ICANN in a hard position and I am sensitive to that, but I'm not as convinced that the DNS Sec reputational issue, on balance, is the same as an unknown implementation in October. Thanks.

UNKNOWN SPEAKER: Any other questions or comments? Cathy, do we have anything from the internets?

UNKNOWN SPEAKER: No.

UNKNOWN SPEAKER: Oh, yes.

UNKNOWN SPEAKER: Thanks for your presentation. My question is that we have found out that, in Nigeria, we have four mobile network operators and

large number of ISPs. We have found out that only one of the four mobile network service providers is using DNSSEC, the other three mobile operators have been [inaudible] other than the DNSSEC. What advantage has DNSSEC over the other security validation systems? Is there any need that the other [inaudible] migrate to DNSSEC? Again, the third question is, what the [inaudible] network service providers that are using the other validation system after the KSK rollover? Thank you so much.

UNKNOWN SPEAKER: Can you repeat the last part of your question please?

UNKNOWN SPEAKER: The last part of your question please, yeah.

UNKNOWN SPEAKER: What I said, the last part of the question is that the three of the mobile network service providers are not using DNSSEC validation system. Then what is their [inaudible] after the KSK rollover?

UNKNOWN SPEAKER: I am sorry, what is their?

UNKNOWN SPEAKER: Their [inaudible]. Is there going to be any effect on the users?

UNKNOWN SPEAKER: That's the easy question, if you are not doing DNSSEC validation there is no impact at all from the KSK rollover. It's the position of the ICANN Org that DNSSEC validation is a good thing, so we do encourage people to do DNSSEC validation and I am sorry, now I don't remember the first part of your question.

UNKNOWN SPEAKER: The first part of the question, I said: Is there any advantage for those that are not using DNSSEC to migrate to DNSSEC?

UNKNOWN SPEAKER: It is ICANN Org's position that you should do DNSSEC validation because without DNSSEC validation you have no assurance that the answer that you receive is really answering the question that you asked, and that it really is coming from whom you think it is coming from. You don't have cryptographic assurance that the answers you are getting really are where they are coming from. DNSSEC provides, you know that the answer comes from who it says it comes from and that it hasn't been modified since it was signed. Modern resolvers have got very clever to avoid being tricked but ultimately to get a higher level of assurance you need DNSSEC. Without DNSSEC, to a certain extent, you're vulnerable

to being spoofed or tricked to believing responses that aren't correct.

UNKNOWN SPEAKER: Thank you so much, thank you.

UNKNOWN SPEAKER: Any other questions? Alright, thank you for coming.

[END OF TRANSCRIPTION]