KOBE – DNSSEC Workshop (3 of 3)
Wednesday, March 13, 2019 – 13:30 to 15:00 JST
ICANN64 | Kobe, Japan

JACQUES LATOUR: Good afternoon. We're ready to start. I hope you all had a great lunch. I'd like to thank our sponsors for this lunch. Our many sponsors. Before we could name them in 30 seconds, now it's like so – Afilias, JPRS, Verisign, Cloudflare, .auDA, NIC.br, GoDaddy, PIR, SIDN, COREnic, Google, CIRA, and Donuts. So that's pretty good.

UNIDENTIFIED MALE: Okay, thanks, everybody. Lunch may not be free to them, but it is free to us which is very nice.

JACQUES LATOUR: So Marrakech will have lobster and….

Okay, so next up is the is the KSK future discussion. We're going to talk about when and how often we roll a key. That's the discussion we want to have today.

PAUL HOFFMAN: That's just a subset.

JACQUES LATOUR: A subset. And I'll leave it to you, Paul.

PAUL HOFFMAN: Okay, thank you. Hi. I'm Paul Hoffman. I work at ICANN in the office of the CTO, and [Matt] my boss sitting here. You've probably seen us at other KSK rollover related things. So now we're talking about what goes on in the future. I'm having a brain fart. This is a fancy clicker. Okay, very good. It was like, wow, my TV doesn't even look like that.

Let's talk first a little bit – so the idea is I'm going to give three or four slides, and then the floor is open. This really is not me presenting to you. This is me planting seeds. And really what we want is the floor to be open. We want to hear from you. But I'll spend a little bit of time planting seeds.

For those of you who aren't familiar, although since you've already been through two sessions this morning about DNSSEC hopefully a lot of this is familiar to you, the process of rolling the KSK has been going on since 2015. On 11 October 2018, which seems like forever ago but really wasn't that long ago, the old key stopped signing the keyset and the new key started.

At that time we didn't know what would happen. We had been getting a lot of conflicting signals before that. And I'll have a slide on this. But there were a lot of surprising things that were found

during the process leading up to 11 October. The reason I'm bringing that up now is that affects our thinking about what we will do with the KSK for the future.

For those of you who are interested in all of the stuff that happened, we have a recently published whitepaper at that URL that explains the whole thing, including not only the rollover but also the surprising stuff that we have found so far.

So really the reason why we're here is not just to say I would like to do this, but given what we now know about the past, what has happened, what is still happening, how do we want that to affect our thinking about the KSK in the future?

Just to be clear, a lot of people have already expressed opinions. We certainly want to hear more. I'll show in a moment how we want to do that. But there has already been a wide variety. So for those of you who haven't been following the discussion, if you hear two or three people today give an opinion and they all give the same opinion – that seems statistically unlikely that would happen – but if it does happen, don't get the feeling that there's unanimity or even I would say strong consensus within the community now. We would love for that to happen later, but for that to happen discussion needs to happen.

Now let's talk about how these next steps will happen. There is already discussion happening on this mailing list. If you're not a

member of the mailing list yet, that link will get you to both the archives – and you'll see that not a whole lot of discussion has happened so far but a reasonable amount – and also the ability to post to the list.

The mailing list is definitely the preferred mechanism for discussion. We're going to have discussion today. In the next 40 minutes, we will definitely have discussion. I don't want to dissuade people. But what's most important in our mind is that discussion happens or that statements are made in a place where discussion can happen. So it's really, really important that at the end of this discussion process that it's not 30 or 40 people saying "I want this" as much as "I hear that you want this, but the reason I didn't think that I wanted that was for this reason." Or "oh, yes, you're right. I want that." We really want to see how the discussion evolves.

Maybe I'm dreaming. This might be like some of the technical discussions which aren't discussions but they're a constant announce list. But given the importance of the KSK to DNSSEC, given the importance of DNSSEC to the DNS, it would be really nice if there was community agreement that came organically, that came from discussion.

So again, it's totally useful for you to speak today if you want to. And don't feel like you have to. But it would also be completely

okay that if you say something today, that you repeat what you said on the mailing list. That will not be considered duplicative. That's a good way. Because as you can tell, there is a small subset in this room of people who are active in thinking about the KSK. So it would be really good if what you said today was brought to the list so other people can see it, respond to it, such like that.

Now once that's done and once we have this discussion, and there's going to be discussion in two other fora in the coming months, in the second half of the year IANA is going to review the discussion. They're going to review what was said today. They're going to look at the mailing list. They're going to review what's said at the other similar meetings. And they're going to make a proposal. How that happens has not been determined, whether there's going to be formal community, informal community, whether they're going to ask for another round of discussion. That's up to IANA.

And just to be clear, I am not part of IANA. Matt and I are from the office of the CTO. We've been sort of the "stuckees" on doing a lot of the public activity around the recent rollover. But in fact, IANA has been the one actually punching the buttons, clicking the mouse, whatever. I guess for us old folks you can turn the dial, do it that way. But that has been IANA, and that will continue to be IANA. So even though I'm running the meeting, it's IANA who is going to be the recipient of what you say in the discussion and

they're the ones who are going to look at this and they will be the ones who come up with a plan for what to do next.

And to be clear, there is currently a plan for what to do next. There is wording in the statement for the KSK that says that ICANN should roll the key after five years, so that clock probably started on 11 October 2018. But for those of you who are good at math who figured out that if we started in 2010 and we did not roll the key until late in 2018, the English phrase "after five years" is being taken literally as don't roll it before five years but that does not mean that your supposed to roll it at five years and one second.

But having said that, that can change. IANA is not run by the same kind of community bottom-up way that ICANN is. But if the community as a whole says we want this, we want to do something soon, we want you to do it later, whatever, IANA very much listens. So that "after five years" wording can change. It's not set in stone. It's in a document that has already been revised a few times in the last eight years. So just assume that what you say will be listened to both by the community and by IANA, and something will come out of this. And like I say, the current general thinking is sometime in the second half of this year after the active discussions are done.

But again, the mailing list is the best place for the discussions as well because that's the best place for you to be able to say –

because I assume no one has today a super prepared and well memorized statement. You can take your time on writing on the mailing list. You can get your wording right. For those of you for whom English is not your native language – by the way, bless you. I'm one of those Americans who only speaks one language – English. So I have deep respect for the folks who use English as a second or third or fourth language. But you can take your time to write up your thoughts and send them to the mailing list.

I have two slides now to explain why we actually want to bring up the discussion of what happened after 11 October 2018. This is a graph of the DNSKEY query volume to most of the root servers. It's not actually all of them, but a close enough approximation. What one would expect is that over time the same number of resolvers would keep asking – and these are under normal circumstances this would be, hey, it's been two days since I got the DNSKEY. TTL is timing out. I want another copy. In a perfect world or in a reasonable world the green line at the bottom that looks like it is at about 1,000 would continue across.

Everyone in this room can see that's not what happened. Where the first set of red bars is where it says KEYROLL, as you can tell from the bottom, that's around 11 October. All of a sudden, a bunch of resolvers started sending DNSKEY queries to the root much more often than they had before. And given that many of

them didn't start sending them, that means a small number of resolvers are sending a bazillion queries.

This was an unexpected result. There are many theories about why this happened. Although the early theories would have said, oh, that's only going to happen for another two days and then it will drop back. You'll notice that didn't happen either. So DNSSEC itself is a fairly regular protocol. How it is implemented in resolvers apparently is a little bit surprising.

So fast forward to the second set of bars which is 11 January of this year, so about two months ago. We revoked the old key. And that means that the old key is still published in the root zone, but it has its revoke [bit set on]. And we said nothing much is going to happen there. Even after seeing this jump, we were like, yeah, yeah. That's going to be a non-event. And that was worse.

We are surprised by these things. We don't know what the resolvers out there are doing. And by the way, that last bar, this slide is actually a little bit old. It is still going up and, in fact, there was an extra jump that happened recently apparently.

So I'm showing this to say that as we think about what should we do about the KSK rollover, if you're about to say we want to do X because we know Y will happen, maybe soften the "because we know." Think about the failure cases. Think about whether this means you want us to roll faster or slower or whatever.

Am I feeling an earthquake? Japanese people, is that an earthquake or not? Yes, it is. Out the door.

YOSHIRO YONEYA: No, don't stand up.

PAUL HOFFMAN: Don't stand up?

YOSHIRO YONEYA: No.

PAUL HOFFMAN: So just to be clear, we do it differently in California. I did the California response. Thank you. The Californians get under a doorway.

WARREN KUMARI: The Virginians who almost never have an earthquake panic and run around.

PAUL HOFFMAN: Okay, thank you. Thank you. So this is a cultural difference. Yoneya-san, thank you for announcing that. But, yeah, we get under a doorway.

UNIDENTIFIED MALE: How did you get the seismograph up there so quickly?

PAUL HOFFMAN: Well, exactly right. So thank you. Let me show you another slide that is related to this which is specifically after revoke. I can't believe I'm being this calm after an earthquake. We Californians sort of do the same things afterwards that Warren was describing of running around in a panic.

After the revocation, this is a log-log graph which basically shows the green bar in the middle is what you would have expected to happen after the revocation. It shows the change in query rates by individual address. So the green band shows the hosts that are asking for the root DNSKEY at about the same rate before and after. That's what you would expect. Those are the big ones in there. But there are a lot of hosts asking for the DNSKEY that were not asking for it before. So something in the revocation caused hosts that had been fine before to start asking for it a lot more. And that's why they're on the upper side of that band.

People have theories about this. That's okay. We don't actually know. Some researchers are working on it. But again, this is another example of what we've learned by collecting data from

the recent KSK roll and shows that we really don't know as much as we thought.

So let's come to what people have been saying so far, and then we'll get directly into hearing from you folks. This is a random selection of what people have said, both on the mailing list and at previous meetings. Why should we roll at all? Why are we rolling at all? What are the motivations? For many people there's an obvious answer to that, and for many people those answers differ. So that's a very reasonable place for you to start in your thinking. Why are we rolling the key at all?

And then if you think we should be rolling it, how often should we roll? Every X years, and there are various numbers for X that people have given, or roll when it's needed? We will do a roll, but wait until it's needed.

Some people have said we're not sure yet. We need better tools to say when they are ready for a new rollover. Other people have said we don't need tools. We don't need to know when they are ready, as we have seen, because we're never going to know and our tooling is never going to be good enough. So both thoughts there have been expressed.

Other people have said we need better bootstrapping for the resolvers so that the ones that are running up-to-date are sure to have the new key. And other people have said, no, that's not so

important. As we saw from the last roll, there was almost no noticeable damage afterwards and, therefore, it really doesn't matter. Should we be nice to people and make sure? Should we just say they'll catch up? Again, two extremes and possibly something in the middle.

Should there be standby keys? Currently, there are still two keys in the root keyset. Although, in a little bit less than a month there will be only one – the new key. Should we have standby keys? This is not about rolling, per se. This is going to the fact that this session is called KSK futures. Anything to do with the KSK is reasonable to bring up in this. If we want to have standby keys, why would we have them? What are the considerations? It will make the DNSKEY set larger. Is that good? Is that bad? Should we care? Things like that.

And then other people talked about, should the signing algorithm change? What are the important considerations? Not everyone is ready for an algorithm change, but an algorithm change might be useful. It might have useful side effects. It might have dangerous side effects.

So all of these things are things we can talk about today. So let's discuss it. We have a roaming mic for the people who aren't at the desks. For the people at the desks, raise your hand. Actually, Kathy, am I doing the choosing? Are you doing the choosing? Russ

is saying I'm doing the choosing. I'll point and I'll actually scoot around this way – you can tell I'm not a Californian really or else I would be scooting toward the door – so that I can see the people here who are raising hands.

Be creative. It's okay for you to bring up something completely new. You don't have to say what other people have said and such like that. Think about what you're going to say, how will that affect not just what you want but how will that affect the whole ecosystem. The last thing we want to do is to make a change in DNSSEC now that causes less adoption, scares people, does things like that. What would be wonderful is if we make changes that, in fact, increase adoption because they are giving people more faith, more stability, things like that.

And then again, I'll beat this drum one more time, whatever you say today please bring to the mailing list. So with that, I'll leave this up, but do not feel limited at all here. Who would like to start whose name is not Warren?

UNIDENTIFIED MALE: The only thing we ask is that people do say their name before they ask their question.

PAUL HOFFMAN: Oh, yes, even if your name is Warren. Okay, so no one is jumping in. Warren's going to scare everyone away. Warren, would you start?

WARREN KUMARI: Sure. Happy to scare people away. One of the obvious discussion points and there's already been a small bit of discussion on it is the whole standby keys solution. From talking to Mike St. John who wrote the original 5011 stuff, the plan was always from his viewpoint that there should be a standby key. And it does seem as though potentially if it had existed, future changes might be less scary because in theory if something goes wrong, there is another key. Of course, that's based upon believing that we have an understanding of how the system works, and currently our track record of predicting what's going to happen when anything changes in DNSSEC is less than brilliant. But this has already started to be discussed, and we should carry on discussing it.

PAUL HOFFMAN: So you just sort of contradicted yourself. Given what you said at the end, do you still feel like a standby key is a good idea, or you're simply advocating for a discussion of standby keys?

WARREN KUMARI: I think a much larger discussion of standby keys and some testing would be good. I personally think that a standby key would be a very good idea, and I think that we should do it. But I'm moderating that with the fact that every time we have thought something related to DNSSEC, we have been proven to be surprised. So, yeah, I think we should probably have one, but more research needed, more discussion needed.

PAUL HOFFMAN: Thanks. Yoneya-san, who's not a Californian, I can tell.

YOSHIRO YONEYA: Hi. This is Yoshiro Yoneya from JPRS. I support keeping KSK rollover regularly because the regular experience [matures] the operation and the software. The operators want more stable operations so that this kind of event has to be regular. We know how to do it, how to respond to it. So I support regular rollover.

PAUL HOFFMAN: Thanks. And I don't mean to put you on the spot, but do you have a picture of a timing that goes with regular? Or is regular the thing that's most important to you?

YOSHIRO YONEYA: I think regular is important. So timing is, I think, once per two years or three years would be good. Less than five years.

PAUL HOFFMAN: Less than five years, but regular is what you're emphasizing. Okay, great. Oh, please.

UNIDENTIFIED MALE: Not Warren from Google. [inaudible], .dk, Denmark. Two points. I agree to do it regularly. Before coming here and seeing these nice graphs, I would also have said do it more often than five years. I'm a little bit more in doubt to what the regular intervals should be. More research needed, if we at all can figure out what it is. Which we've seen before in deciding when to do the key rollover for the first time, it was very hard to find out what might happen and get in touch with those resolvers. So it's not sure that can actually find out what's going on with these extra queries right now. But at least we should give it a try. I think more regular would be hopefully better.

I also agree with Warren, funny enough, on more research needed for the standby key. In Denmark we actually got rid of our standby keys.

PAUL HOFFMAN: Please say much more about that because I have not heard from anybody who has thought about it enough to actually make a decision, much less that decision. Can you talk for like five minutes on that? That would be really useful.

UNIDENTIFIED MALE: I don't know if it's going to be five.

PAUL HOFFMAN: Okay, can you say a few sentences at least?

UNIDENTIFIED MALE: Yeah, what I was going to say is, what does the standby key prevent or hinder? What risk are you trying to solve by having a standby key? What scenarios do you think you might need them for? What we found out was that our private keys, both the online and the standby key, were on the same infrastructure. So if there was an infrastructure compromise, then we could roll the key. Well, then the standby key is compromised as well. We might as well not have it.

Then, of course, we had to look at do we need to have a separate standby key [and] different infrastructure, and we came to the conclusion not to do that. It would be easier to just generation a new key and then quickly do the resigning. But then again, we're

not the root. It's much easier at the first level, top level than at the root.

PAUL HOFFMAN: You're not the root, but you obviously have thought about it. So please don't denigrate your thinking.

UNIDENTIFIED MALE: No, no. It's easier for doing emergency rollover.

PAUL HOFFMAN: That's true, yeah, below the root. Was anything written up about those decisions, even if it's in Danish?

UNIDENTIFIED MALE: I don't think so.

PAUL HOFFMAN: Okay, could you look and find? Because that would be very useful to the whole community. And if you can do it soonish, like in the next few months. And I'm saying this to you, but I'm saying this to anybody in the room – if your organization has done any research, even just we wrote up a page about this or that, having those be published and a link to them on the mailing list. And again, even if it's in Danish I would spring for the translation since I happen to

have a bunch of Danes in my background and one of them is a professional translator. So I would do that.

UNIDENTIFIED MALE: I don't think it was actually written down [inaudible] in English, but yes.

PAUL HOFFMAN: Or if the people who made that decision would be willing to write a paragraph or two, those kinds of things I think would be very helpful in the bigger discussion. So thanks. I'm avoiding Warren, but I will go to Warren. But really, if you folks have thoughts, please do share them. Ah, Jacques?

JACQUES LATOUR: I beat you to it. In the context of ICANN, how would a standby key work if you have one HSM signing infrastructure?

PAUL HOFFMAN: You tell us. Seriously. This is a very open discussion.

JACQUES LATOUR: I'm asking you.

PAUL HOFFMAN: If you have a design that you think would work, great. If you just have an idea how it would work, IANA has plenty of good folks in the crypto group who could decide that. They may decide it can't work. But it's likely that if the community wants a standby key that they would come back with a proposed or possibly multiple proposed ways of doing it and ask the community. But if anyone here is, as you were, using standby keys, your design might be useful. How many people here are from ccTLDs or gTLDs? And how many people, of those hands up, have a standby key, a standby KSK?

WARREN KUMARI: Define that. Published?

PAUL HOFFMAN: Actually, thank you, Warren. Warren asking define that. Is it published or not published. Of any sort? How many have a standby key? Okay, I was hoping for a few more hands and then I would bug you to write those things down. I will bug you to write the thing down that you just raised your hand about Warren. But having these kinds of things written down – for example, one of the things that we had talked about was algorithm rollover. Some of the ccTLDs have done algorithm rollovers, and fortunately they've been writing up the results. So having these things written down will be of great value to this conversation.

Okay. So, Warren, I think I have avoided you long enough a second time. Please.

WARREN KUMARI: First off if it's okay, I'll quickly respond to Jacques. It kind of depends on what you think the purpose of the standby key is. If the purpose of the standby key is so that you have another one ready in case the primary is compromised, then you will need a very different design for building up more HSMs, etc., than if you simply want another key so that it's easier to do key rolls. We could, if we had two keys, it looks as though 5011 might make key rolls easier if you always have two published and then you just move out one and promote the other instead of introducing and removing.

JACQUES LATOUR: Okay, so my understanding of a standby key was something you have [a DS for] that's public.

WARREN KUMARI: But if it's public, then there's a long time for resolvers to learn it and put it into their thing. Revoking becomes less of an event. So it's sort of potentially always being in key roll state versus swapping it in as an emergency. So this might require a whiteboard and more discussions [inaudible].

PAUL HOFFMAN: Better than a whiteboard, a mailing list.

WARREN KUMARI: Yep. But what I was originally putting my hand up for was to respond to your question about how often or how regular is regular. At one point, there had been a number of people discussing annually or possibly even faster. Some of that I think was before we had seen the most recent set of data. But I think that the most recent set of data is something that we should figure out what the cause is, track it down, and squish the set of particular issues. And then the next time we roll, there will presumably be some other set of issues but hopefully a lot smaller. So although it causes some pain, there is some advantage to having a very regular set of rolling like annually or something similar in that it forces you to actually clean up the hygiene and make things better. But obviously, this requires more discussion, mailing list, etc.

PAUL HOFFMAN: You said what other people thought. What do you think?

WARREN KUMARI: I knew you were going to ask that, and I was hoping you weren't going to.

PAUL HOFFMAN: That's what I care about more.

WARREN KUMARI: I personally was thinking annually or even at one point I was thinking six months, which I realized was [inaudible].

PAUL HOFFMAN: What do you think now?

WARREN KUMARI: What I think now is that I cannot give that answer until we've looked a bit more to solve this particular thing. Once this particular thing has been figured out and solved, I think it will inform us a lot better as to, is this risk structural? Is this risk one weird implementation? So once we know that answer to what's causing this particular issue, we should be able to make a better decision.

PAUL HOFFMAN: I'm sorry. Just because no one else is raising their hand, I'm going to keep poking at you. You said once this issue is fixed. The line at 4,000, if that stayed constant for two years, is that fixed or not?

WARREN KUMARI: What I'm calling "this issue" is the one after the revoke. The other one, I'm like that's weird and annoying. But the unhappiness is the revoked and the continuing to climb revoked.

PAUL HOFFMAN: Okay, thank you.

WARREN KUMARI: And there are some people discussing it.

PAUL HOFFMAN: Yeah, thank you. I see a hand.

UNIDENTIFIED MALE: [inaudible] from [.oe registry]. Looking at the graphs, I thought it's maybe not clever from cryptographical point of view but what if it was rolled back to the old KSK key. So maybe the situation would normalize. I don't know if that – now since it's revoked it's too late.

PAUL HOFFMAN: Yes, and that was the part that – once you revoke a key even for five seconds, you can't go back because those people will never trust it. And quite frankly at the time, no one asked us to go back. We were watching and such like that.

Let me go back to the graph and I'll extrapolate just a little bit here. One of the interesting things about this second line at 4,000 is no one can explain, at least to me well enough, to say why something should be doing that for more than 48 hours after the second red bar so very, very soon and still actually be a validating resolver that is answering questions for anybody. This is an indication that I cannot get the DNSKEY set that I want, and therefore it should be serve failing all of its queries. So is this an indication that there are a ker-billion resolvers out there that have nobody asking them queries? Which is where I would go. Or is it something even more mysterious?

Actually, Eric and then Jaap. I thought you had your hand up.

UNIDENTIFIED MALE: Tim.

PAUL HOFFMAN: Tim. I'm sorry.

[TIM]: I was confused because you said Eric and that's not me.

PAUL HOFFMAN: Yeah, well, that's be I used the wrong name.

[TIM]: The notes I have been taking so far, do we want to see another roll? I'm interested to see another roll in the not too distant future partly because of the research scientist in me that wants to see what that graph looks like one more time. Because if I had to make a theory right now, I think that big peak is hardcoded old key. But I wouldn't think that it would get higher with another roll, but we don't know.

And for the standby key, for another process I run that's not DNS related I keep a standby key in a separate HSM that we roll through every six months. So that we have two ready to go that are published and they're accessible by the consumer of it. So that we use one for six months, roll to the other key, refresh the old key, and then put that into the safe for the next six months. So that we always have two ready to go.

PAUL HOFFMAN: Okay, thank you. Jaap?

JAAP AKKERHUIS: Jaap Akkerhuis, NLNet Labs. One theory [inaudible] came up with why this happened is that there are [inaudible] resolvers out there which actually uses the [resolve] library. And so they are not really doing any validations. And so just because they're just blindly using the library, the library still is trying to do validation [inaudible]. Now the [inaudible] change it will be always be invalid. So they're trying to look up the keys again. So that could be possible, but the only way to find out [is track] who is doing it and trying to find out [why].

PAUL HOFFMAN: Well, if you're correct, then that curve should drop radically once we take the key out, which will be on 22 March.

JAAP AKKERHUIS: But we'll see.

PAUL HOFFMAN: Yes, exactly right.

JAAP AKKERHUIS: But that's a theory. And a lot of people using the BIND libraries and Unbound libraries in new and unexpected ways, and they might be causing this behavior [inaudible].

PAUL HOFFMAN: Great. And one of the things that we've found during the – so some of you might be aware that we were going to roll the key on 11 October 2017. And I won't go into a lot of detail, although it's in the report. There were a bunch of anomalous results that we started seeing, so we postponed for a year. And during that year, because we were seeing some numbers we didn't expect, we discovered that there were some applications using some libraries incorrectly through good research that people in the community did. And we were actually able to get in touch with those software developers and they were like, "Oh, whoopsie!" and fixed it. And when they pushed those changes out in their software, we could actually see a drop. So that may happen again. All of these things are possible. Don't know the timelines.

Other thoughts? Oh, yes, please.

UNIDENTIFIED FEMALE: This isn't really helpful to planning the rollover, but it's common when you have a few clients that are querying too often to rate limit those clients. Has anybody given any thought to whether or not that would be a good idea in this case to have some algorithm to rate them at the top? It looks like Warren has a comment on that.

WARREN KUMARI: Yes, that's definitely an interesting question. A number of the root servers currently do stuff like RRL, which should we think be limiting them to lower numbers than we're getting assuming that a bad response causes them to re-query, that it is actually a feedback loop based system and not just them spewing queries. It seems as though maybe they're not respecting that or it's not working right.

One obvious thing to try is to watch the rate of queries and then stop answering to one client and see if that stops them asking. That will then show if a bad response is triggering another query or if them just asking as fast as they can is what they're doing. So that will be an interesting experiment, and I believe somebody is doing it soon or as we speak. But, yeah, possibly rate limiting and also if they are well-behaved hosts, you could potentially make them slow down by just answering certain queries slower.

PAUL HOFFMAN: But we don't know. Other thoughts on any – here, I'll go back one to here. Anything from here or anything else? Yes, please?

JACQUES LATOUR: In my view, the data that we're seeing is either bad code, botnet, home gateways that are firmware that's messed up somewhere. I don't think it's big problem because if everything [was]

[inaudible] [fail], people would complain and they fixed it. They're using something else. So the user impacts of this I don't think people have been dead for a week or two weeks or whatever.

PAUL HOFFMAN: Or four months.

JACQUES LATOUR: Or four months or whatever. So it's background noise on the Internet. There's going to be background noise for – I can gauge depending on how you put your hand if you agree or not. I think it's not a big issue. There is information there, but I don't think it's super important with the rollover because the stuff that's serving people that people need to use, the system, the APIs, all the production stuff is working today because they would have fixed it when it broke, if it did break.

So these are interesting data points, Internet data. I'm sure if we look at other data sets, there is interesting stuff that we don't know why it's happening. It's not just this. So there's background noise on .ca. We get a lot of spike once an hour for five minutes. I have no clue where it's coming from, what it is. I'm looking into it, but it's not breaking. If it broke, people would complain and it would have been fixed. So I wouldn't worry too much about that.

What I would like is I'm in favor of doing this quick, the next one. I think six months or something like that. Because we need to be good at this. If we're going to use DNSSEC, we need to make sure that we understand it and we're not afraid of it. And then after that, we should a year or two rotation. Once we understand that we're comfortable with it and we're confident in it, then we roll on a regular basis. But I wouldn't worry too much about the queries because, like I said, it could be [Andrei] that wrote a script to mess everybody up, right?

PAUL HOFFMAN: Okay, thank you. I'm going to ignore you again, Warren. Really, because we want a plurality of opinions or voices or questions. Russ?

RUSS MUNDY: Thinking about one of the particular presentations we had just shortly before lunch, Duane Wessels did a review of some of the traffic that he'd seen and Wes Hardaker gave us a view of what he was seeing at B-root, the analysis itself was not exactly identical but they were similar. And especially in Duane's there was somewhat of a concern that the rate of growth was continuing to go upward after the revocation. That may or may not be really related to the revocation but the incident of the timing makes it at least very questionable. If in fact this changes after the 22nd

when the key is removed, then it's probably something that as long as the overall rate doesn't continue to increase needs research to figure out what it is but at least in my view then we need to start looking at how soon we can turn this into a regular event. You've asked a bunch of times different people what's their opinion on what is the regularity.

PAUL HOFFMAN: I've been a little bit polite. I said if you have an opinion. I wasn't trying to force it. But if you do, please.

RUSS MUNDY: Yeah, I do, and that is we should strive to be able to do a KSK rollover in he root once a year because that is far enough apart that things can be done to correct things in between but it's close enough together that people get used to a cadence of the change. So I think that we should strive for that but not until we figure out what's going on with some of these data indications that we're getting at the moment. So figure that out, then look to try to get to a year. Thanks.

PAUL HOFFMAN: Great. Thank you. I'm going to do something different now. Jaap and [Andrei], you both work for companies who distribute resolver code. Russ just said once a year. Since you have

customers who rely on you, customers who use your software who don't rely on you and such like that, how do you feel about once a year?

[ANDREI]: Actually, I have an even stronger opinion about the thing that we do [open source router] who is doing DNS validation. It's much more complicated. Those guys who use [our] resolvers, they usually know what they are doing so they're fine. But doing like end user customer [inaudible] it's much harder. And honestly, we fixed the [met kit], the bootstrapping stuff and the router on the last moment when the rollover was running. So every year would be – as a technician I feel it's perfect. I support [inaudible], but as a vendor of a [inaudible] device I'm a little bit scared of that.

PAUL HOFFMAN: And I was asking you as the latter. And Jaap?

JAAP AKKERHUIS: Well, we all have to different opinions when we should do this often and not really defined often that much. I mean, I guess we should first wait until the dust finally settles. But once a year is fine with us when [inaudible] you think that's – I mean when things go smooth, especially for people getting used to it and all that stuff. We update our software more often than once a year.

PAUL HOFFMAN: Okay. Vicky?

VICKY RISK: I basically have the same comment. I mean, we update our software quite a bit more often than once a year. We have a new stable version every year. But as we saw from the [BIND] [inaudible] version presented in one of the earlier talks, plenty of people are running five-, six-, seven-year-old code out there. I don't have a solution for that. If I did, I would bank it right away.

I agree it's good to do it on a regular basis. It was a lot of overhead on us and I suspect the other open source publishers as well this last time because we had spent so much time on testing, coming up with new tests, fixing corner cases, implementing multiple different telemetry options for tracking the resolver versions in the key tags. It should be much easier the next time around.

I would like to see us understand what's going on with some of these weird behaviors because to me there must be some bugs in there somewhere.

PAUL HOFFMAN: Okay, and these were the three software vendors I knew. So there might be other ones in the room who I either don't see or don't know. Does anyone – okay, so please?

STEVEN CARR: Steven Carr from Infoblox. We're kind of in a half position in that we leverage BIND as our main DNS engine, however we don't actually support 5011. Which is, yes, we know that's a big issue.

PAUL HOFFMAN: And as of now we know it's a bigger issue.

STEVEN CARR: It's a bigger issue and so selfishly I'd like it to be rolling on a year basis which gives our developers a kick to formally support 5011 and put it into the product so that the users don't have the configuration overhead to roll the keys manually. However, when it comes to actually signing a DNS zone, so leaving the root to one side, our recommendation for your own zones is on a yearly basis for the key signing key. That's the standard figure that we have in our UI is a year. So unless you go in and change it, it will force you to roll the key signing key every year. Which we made that decision pretty much on the same basis that other people have suggested, that it needs to be on a regular occurrence that people don't forget how to do it and just that general cadence of keeping

things moving. Because if it's more than a year, you're going to get into a situation where people may move jobs, people may come in and out of the industry. Somebody gets presented with the system, something happens, and they're not familiar with what is supposed to happen because it hasn't happened for the previous three or four years. So something that is regular, yearly basis, it keeps things present and keeps it fresh in people's minds.

PAUL HOFFMAN: Great. Thanks. Vittorio, can you speak to PowerDNS or not? And no is an okay answer. I thought you were more on the policy side these days.

VITTORIO BERTOLA: You should really discuss this with [PowerDNS].

PAUL HOFFMAN: I'm sorry?

VITTORIO BERTOLA: You should really discuss with the PowerDNS people that I [inaudible] because I don't know.

PAUL HOFFMAN: Okay, well, please bug them to get involved on the mailing list. Again, I was just trying to pick out any of the software developers in the room.

We have just a few more minutes, and I'd really like to avoid having Warren speak again. Even though I love him dearly, I would like to hear new voices. So somebody who has not said anything who has heard this discussion who might have an opinion, especially folks out back away. Oh, please. Thank you.

AKIRA KATO: This is Akira Kato from WIDE Project. I think the algorithm change is very interesting and contributes to reduced packet size if you choose a proper algorithm. But in some cases, the [data] algorithm may need to be supporting everybody. So it is essential to announce that the [inaudible] [rollover] will take year of, for example, 2025 or something like that and then encourage the people to develop and get prepared. That is the announcement might be necessary, I think.

PAUL HOFFMAN: Great. Thank you.

AKIRA KATO: And the other [inaudible] information announcement is okay?

PAUL HOFFMAN: I'm sorry? Oh, please.

AKIRA KATO: Okay, so that earthquake is about 100 kilometers and 50 kilometers [deep] and the magnitude was 5.2. You don't have to worry about the tsunami at all.

PAUL HOFFMAN: Right.

AKIRA KATO: But if you are near to the shoreline and feel a big earthquake, it is essential to climb up the building of five floors or above to survive. Thank you very much.

PAUL HOFFMAN: As someone who has a friend who actually lived through a tsunami, I heartily recommend following that advice. But again, the Japanese do it different than we do in California. And I'm going to look into that. Actually, I find that interesting. A little bit less interesting than the KSK at the moment. Please?

FREDERICO NEVES: A comment regarding the algorithm rollover.

PAUL HOFFMAN: Yeah, so name [and affiliation]?

FREDERICO NEVES: Federico from NIC.br. From the history that we have only doing our first roll, we took a lot of years for doing that. I predict that probably doing an algorithm rollover it would probably take more than what we took to do a regular roll. So perhaps we should not spend time with this now, even that I think that we should do, because probably this packet size will not be a problem in ten years. We will probably have different transports that will probably overcome this. Kind of sad comment.

PAUL HOFFMAN: Great. Thank you. So since we are now out of time, I would like to thank you for the contributions. I again will exhort you to – let me put up that URL again – to join the mailing list even if you just want to watch. But particularly if you said something today to repeat it.

And I would also like to exhort you if you are in an organization that has thought about this, even if you don't have a strong opinion and the organization has even a weak opinion but an official one, by all means write it up and send it to the mailing list. That would be very, very valuable.

For those of you in the room who are researchers who have done any even informal research or are making wild guesses about what those graphs mean or what the graph is going to look like on March 23rd, or 24th given the 48-hour TTL, that would be really useful.

And by the way, NIC.br did do a nice writeup about algorithm rolls and such like that. That has been very useful. Actually, I think the CZ.NIC folks did as well. So things like that as compared to just, oh, I have this opinion or I've thought about this. Things like we've looked at this and we have research – very, very valuable.

We'll collect some of that separately. So again, it's IANA who is going to make the decision on moving forward and such, but ICANN org will in fact collect URLs for statements like that and put them on a webpage such like that so that other people can find it easily.

But we look forward to more input from you, especially on the list, and thank you.

RUSS MUNDY: Thank you very much, Paul and Matt, for joining us. You're very welcome to stay. The last portion of the day is an open Q&A session. We haven't done one of these for a while. Originally the idea of the program committee was that we would ask for inputs

and comments and questions on both the program today and the program that you would like to see at the next meeting. Because we always want to get more input from the community about what is useful to the community with respect to DNSSEC or any DNSSEC related topics.

So with this, I'd like to ask for folks to express their views with respect to what we've had today and what you'd like to see in the next workshop. Anybody. Not all at once. Ah, Yoshiro.

YOSHIRO YONEYA: This is a question to the audience. Can I speak Japanese? Because there are several Japanese here. [speaks Japanese]

UNIDENTIFIED MALE: I'm first time to join such an international event. So, for example, today this discussion topic is very professional topic. So I think it depends on the contents of the topics. Some very difficult topics if Japanese translation available, many Japanese can attend this kind of event I think. Thank you.

RUSS MUNDY: So as a quick follow-up to that, I would like to ask if any of the people in the room – I've seen some people making use of our interpreters and listening – do any of those folks have comments

with respect to how helpful is that to you for anyone that was listening to our interpreters today. Somebody that was listening. I don't remember who it was. I saw them earlier. Well, I'm not seeing anyone raise their hand. But I do want to very much thank our interpreters. They've done a marvelous job. So we can take more on that if anybody has any responses. Yes, Vittorio, go ahead.

VITTORIO BERTOLA: My only comment would be that I understand why we are doing French and Spanish, but maybe it would be good to have also the local language at least when we are in [town] to have the local language which is not French, Spanish, or English.

RUSS MUNDY: Okay. Any more inputs there? So other comments and questions from the participants? I know we have, like I said, several of the presenters here in the room yet if anyone has questions they thought of after the presentations that they want to ask. I see where a number of them are sitting, so feel free to ask. No? Okay.

So with respect to the next workshop, comments from the audience about what was the most – let me ask this. What did you think was the most useful to you in terms of the presentations

today? And what would you like to see more of next time? Somebody start the discussion.

UNIDENTIFIED MALE: More of a general comment, and don't take it as a critique to the ICANN staff organizing this whole event. It must be tremendously hard to get everything organized with all the many sessions. This morning I had a security, an anti-abuse, and a DNSSEC workshop going on at the same time. So try to look at anything security related to not overlap with DNSSEC. I know it's not possible to do that every time with so many sessions, but just a general to give it an effort at least. Thanks.

RUSS MUNDY: Thank you. Yeah, so a lot of conflicts for the technical folks today particularly. Okay, we will note that and provide the feedback to ICANN and hopefully we can get it improved but we never know. But it's good to have the feedback. It wasn't just one or two of us that noted it. Other people have too. Good.

What else do folks have in terms of thoughts for today? John?

[JOHN]: I enjoyed what we might call the Warren and Vittorio smackdown. No, they both sort of – I mean, historically most of the talks have

been fairly low level and technical. Here's what we did with DNSSEC in our registry. Here's how we dealt with this particular key roll problem. Here's how we debugged these resolvers. And having a slightly higher level, looking at higher level issues of, what problem are we trying to solve and are we taking the right approach? I mean, I wouldn't want all of that because you run out of interesting things to say. But I would look forward to having other people do thought pieces like that in the future.

UNIDENTIFIED MALE: So having a discussion around – a lot of it is in the browser. There's a lot of stuff moving there, so having a discussion in Marrakech or the next one on the browser, DoH, and all the, not the policy around it but the politics I guess too.

UNIDENTIFIED MALE: Also, I don't know to what extent people even know where resolution is happening. We know browsers do it and we hear that Roku does it and stuff. But if people know academics who have actually studied this or could be persuaded to study this, that would be really interesting.

RUSS MUNDY: Okay, good. What else? Vicky?

VICKY RISK: I really enjoyed seeing all the data presented. Besides the fact that I thought that Duane's visualizations were really cool, and I'm saddened that probably people who see the files presented online won't be able to see those. I like it that we were presented with a lot of data which hasn't yet been fully analyzed. It's kind of more engaging when you know there's a problem that isn't solved yet rather than just being presented with the answer. So I thought that was pretty powerful.

Usually in this session we have a lot of people sharing operational experience, and we had some of that. I was particularly interested in the .au presentation. Any time somebody is trying to transfer signed zones, it's always a problem. So those were two things that I really appreciated.

I did like the higher level presentation about, who does the end user trust? We have really never in this environment that I've seen talked about usability or end user research, and that's probably something we've neglected. And we maybe haven't taken advantage of the fact that there are so many policy people here.

RUSS MUNDY: Okay, good. Thank you. You said you enjoyed seeing the data, the heavy aspects even though it wasn't finished. Would you think

more of that, particularly the finished pieces, would be a valuable thing to include next time?

VICKY RISK: Yes. It has given me something to do. I want to take this back and get some more people to look at it and see if perhaps ISC operates the F-root and maybe we can look into some of our own data. Certainly, I'd like to investigate a little further whether or not there's some of our bugs that perhaps account for some of this tail of traffic. I'd like to answer the question that Warren raised which is why isn't RRL dampening this. So it has given me some things to think about and some things to follow up on.

RUSS MUNDY: Okay, great. What other things do folks want to have the program committee pursue for next time? Anything additional that we haven't talked about yet already?

UNIDENTIFIED MALE: We were thinking of adding more security to the workshop, less of the standard roundtable that we do unless it's very exceptional. But what kind of security topics you'd like to see? Because we're putting the [inaudible] together basically. Presentation, I mean.

VICKY RISK: One thing that was already mentioned is trying to de-duplicate the agenda for this day with the other security topics. Even if that means maybe broadening this beyond DNSSEC, I don't know if that's a possibility.

BARRY LEIBA: Maybe this has been done before; maybe there's a reason not to. But can we reach out to some major operators that have not implemented DNSSEC and have them present something about why not?

RUSS MUNDY: It has been a while, but we've done that a few years ago. It has been quite a while. And especially given the passage of time and now that the first KSK roll is done, that excuse has been removed. So I think – go ahead.

UNIDENTIFIED MALE: So we can do it in Canada, in Montreal ICANN, because pretty much all the Canadian ISPs don't validate. So then it's not far for them to travel.

UNIDENTIFIED MALE: Perfect.

UNIDENTIFIED MALE: We've been talking about this also with the rest of the ecosystem. The DNS isn't just the protocol and it isn't just the resolvers. There's an awful lot in the food chain on the authoritative side – registries, registrars. And we've been talking about it elsewhere in ICANN. That also seems to be something that's a little timely. But it's the security of the DNS as opposed to DNSSEC.

RUSS MUNDY: Okay, thank you. Yeah, in fact, that is very much in line with discussions that the program committee has had that we're looking at trying to have topics beyond, if you will, core DNSSEC as part of the DNSSEC workshop. Maybe eventually someday it might have a different name. For a while it will stay DNSSEC workshop, but we'll certainly look at spreading out and getting additional topics.

This then brings me to the next point that I wanted to make with respect to this. The program committee, although we assemble it and sometimes some of us do present on it, the material comes from the community. So we very, very much would like to have inputs and responses to when we send out our call for papers. What are things that you would like to see here? What would you want to present? What would you want to talk about that would be of interest to the community? Because this is really very much

intended to be a community sharing, community outreaching session.

So please look for the call for participation in the next, it will be within a month because the next meeting is not that far away. It's the end of June. June 23 or something like that. Anyway, it's the end of June for when the next meeting is, so please be thinking about it.

UNIDENTIFIED MALE: Where do we look for the call for participation. Where is the call for participation posted?

RUSS MUNDY: We send it to a number of lists. The dnssec-coord list is one of the main ones. Jacques, you send it to a couple of the ccTLD lists?

JACQUES LATOUR: Yes. The ICANN gTLD tech list, CENTR, NANOG, and we have Yoshiro pushes to a couple of others.

YOSHIRO YONEYA: Yeah, DNS operations.

JACQUES LATOUR: But it's one of those e-mails that people delete.

RUSS MUNDY: Well, if folks do have suggestions as to what would be good additional mailing lists to use, please send them to the address for this meeting which is – I think it's in the slides somewhere. Anyway, yeah, so let us know what you'd like to have us add in terms of mail lists for the call for participation if there are some more good ones to send it to. Because it's easy enough for us to do, and we want to get as broad a participation as we can.

So any last thoughts on the program today, the program for next time? Jacques?

JACQUES LATOUR: I had one line item here. It's decoupling domain registration and DNSSEC registration at the registrar. So taking the DNSSEC registration from the registrar, taking that out of the picture and only do it out of [band] with the DNS operator and CDS/CDNSKEY. Because if the domain is compromised, then the bad actor can remove the DS record and un-sign the zone. Then it takes all the benefit of this away. From experience, most of our registrars are not – speaking as a ccTLD – most of our registrars are not interested in doing DNSSEC. The DNS operator portion of our registrar is but not the registrar part. So I think we could have – is that an interesting topic to have for next session?

RUSS MUNDY: Okay, so we should see that in our call for participation. Okay, well, we will take all of these on board as the program committee. We want to just do the quick run-through. This set of slides is one that we have used as the closing for the workshop for quite a while. And frankly, one of the reasons they're here is not so much for the people in the room to see them and hear what's on them as to have them on the website. So when people say, "What do I do? How do I go about doing this?" this is a summary slide that's aimed at where your functionality in the DNS world might be and what you can do to make it happen, make DNSSEC improve your security and make that happen.

So starting with ICANN TLDs, that's the typical set of things to do. Sign your TLD if it's not. We're doing really well in terms of percentage signed. Did Dan have a number this time, Jacques? Do you remember? I think it's between 85-90% now I think for signed TLDs, something like that. So that's really excellent. But for those that aren't, please keep going.

And make sure you accept the DS records. If you're a CC, turn your country green. And work with the registrars because many times the registrars have proven to be the most challenging to get involved in the DNSSEC world.

Today what we saw was statistics, statistics, statistics. And we can learn so much from the statistics. And we can actually end up seeing statistics that we may not understand but at least we have them and can work on it and try to figure out what is going on.

Any zone operator, same thing. Sign your zone. Do verify that your registrars support DNSSEC too. So work with your community for where you are. And again, help with any statistics that you can provide. OARC does a lot of statistics collection. There's a lot of researchers that collect it and are doing analysis. And it is a very open and sharing community for this sort of thing so we can learn more.

Now zone operators, if you're an enterprise, I am happy to say that the company I work for does run a signed DNSSEC zone as an enterprise. So parsons.com is a signed zone. If you're a straight operator, offer it to your customers. Consumers, ask for DNSSEC. If you're an ISP or a service provider, validation, validation, validation. You probably operate your own zone, so sign your zones yourself.

Things anybody and everybody can do. Be a DNSSEC user in as many ways as you can think of. Share what you've learned, what you've done with the community. Send in your thoughts for whether it's the KSK rollover, whether it's the next workshop.

People want to know and want to hear what other people are doing and hearing.

This is the chance for myself and Jacques and Yoshiro as the folks from the program committee that are here to particularly thank all of our presenters and participants today. Thank you, [guys]. And just as important, our sponsors for lunch. Thanks, everybody.

If you want, again, the reminder of where the sponsorship for the organization comes from, that's where it is. The SSAC and the Internet Society Deploy360 Programme.

Websites, everybody has to have websites where they can find things, and there they are. So dnssec-coord is actively going on, once a month calls. And we hope to see everybody at ICANN 65 in Morocco. Thank you, everybody.

[END OF TRANSCRIPTION]