BARCELONA – Joint Meeting: ICANN Board & TEG
Thursday, October 25, 2018 – 10:30 to 12:00 CEST
ICANN63 | Barcelona, Spain

CATHY PETERSEN: Cathy Petersen, icann.org.

UNIDENTIFIED MALE: Those who are in the back, if a seat is available still then please grab a seat. Come along and then if there is a [tech] member, then please, I'd like to have your names stated. No, okay. [inaudible]

All right, let's get done our business. We have a maybe handful of agenda both from the ICANN Board side and the TEG side. That's my understanding, so I'll start with our side. David, please.

DAVID CONRAD: If you could put up the agenda slide, please. Okay. So we have, this is sort of the Administrivia portion of the show. If anyone has any agenda items they'd like to add, please let us know now. Otherwise, we will jump right into a debrief from, on the KSK rollover. This is actually a slide deck that I believe Matt has provided on several occasions so some of you may have seen it already. For those who haven't or who are unaware, you might have heard that we rolled the KSK on October 11th and it went mostly without hitch. I was actually, honestly, a little pleasantly surprised at the lack of problems that we had. We, well, we'll get into that.

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.

First slides. Thank you.

So as mentioned, at 16:00 UTC on 11 October, the Root KSK was the zone signing key signed with the new root KSK. KSK-2017 was put into use for the very first time. Next slide, please.

Seventy-two hours after the rollover, there were very few issues that were identified. ICANN itself didn't receive any direct reports at that time. We had seen a small number of issues via observation, areas like Twitter or mailing lists, and other operational forums. The notices that we saw were basically individual system administrators saying, "Oops, forgot something," and just letting other people know they made a boo-boo. There were no reports of significant numbers of issues or of folks impacted.

There were two outages that we were made aware of that may be related to the KSK rollover. One of them, EIR, which is an Irish ISP, if you go to that URL and scroll down into the comments, you will see a large number of very irate people. It looks like EIR went down for about 12 hours and indications are that it probably was related to the KSK. We have attempted to contact them. We have attempted to contact people to contact them. They have been unresponsive so we can't be sure but it does look like that was the stereotypical probably associated with the KSK rollover where they had validation enabled and didn't update the key so everything went bogus on them.

The second one is a little less clear, Consolidated Communications, a Vermont ISP. They also appeared to go down, however, the news report describes a situation where the number of customer complaints – it was an acquisition by a state telephone company of a local ISP – and the number of customer complaints appeared to have gone up over 2,600% so it sounds like they were having problems. It's unclear whether the problems that occurred on October 11th for that ISP were directly related to the KSK, but it was suggested that might be the case so it might be. Next slide, please.

So this is some of the outreach that we had done prior to October 11th. We spoke at over 100 events over the three and a half years that the KSK rollover was actually being pushed. That provides a list of them. There were a lot more. Next slide, please.

We identified a number of minor software issues related to the KSK rollover and reached out to software vendors to get those fixed. That included BIND, Unbound, PowerDNS Recursor, and Knot Resolver. As a result of Warren Kumari scrounging through Github – he suggested we might want to do the same – we found over 2,000 software repositories that had the 2010 KSK embedded in it. 1,400 of those did not also have the – well, no one had the KSK 20178, but 1,400 did not have KSK 2017. We attempted to notify those folks. 638 of those packages had not been updated in six years or were not open to receiving any issue reports. Next slide, please. And one of those, I believe, was the plugin for the Firefox browser that did DNSSEC valuation that was maintained by CZ.NIC. I believe they fixed that and pushed it out.

Other things that we did, we created an automated testbed for the operators to test RFC 5011 automated updates. At the end of the roll, there were almost 1,500 participants on that testbed. We sent a survey, a KSK preparedness survey, to 16,000 networks which there were two contacts per network, so it was about 30,000 e-mail messages. I'm sort of surprised ICANN didn't start showing up in spam filters.

We had a response rate on that survey of about 4%. Actually, we didn't anticipate anyone really responding. This was more of an outreach effort, basically to try to get people to be aware of the fact that the KSK rollover was going to occur. But it triggered 150 operators to ask us questions and some of those questions, frankly, were a bit terrifying. There were also, as a result of that, nine new subscriptions to the 5011 testbed, although by that time, the testbed wasn't actually useful because of the timeframes involved.

We contacted the major public DNS providers and confirmed that they were ready. Actually, I think one of them said, "Oh yeah, we should probably do something about that," which was also a little surprising. Next slide.

So additional outreach communication type efforts, we had, obviously, the dedicated KSK roll webpage and it was available in nine languages. Between July 2017 and October 2018, there were about 60,000 unique page views. The next highest was 10,000 for the Japanese page. Media relations, we published almost 500 articles around the world on the topic. Next slide.

Social media that I'm sure that means a lot to people that do social media. I don't actually know what it all means, whether that's high or low or anything, but I assume our Comms Team did a good job on that. Next slide.

We had 40 resources published, blogs, papers, educational videos, PowerPoint decks, and in June of 2017, we sent a note to telecom regulators and GAC reps for 150 countries telling them about the KSK roll and that was the 2017 KSK roll but then we postponed it so we sent a follow-up letter saying, “Oh, we postponed,” and we would reschedule and told them the reschedule date.

In May, I sent a note to a whole bunch of Internet exchange points around the world to try to get them to notify their members. Next slide. Okay. It’s okay. I hate it when that happens.

But we weren’t the only ones to do the outreach stuff. We, Verisign participated actively in outreach. I believe they contacted 7,000 networks. The CIRA, I know, did some work. A bunch of other folks did some work.

The RIRs also, I believe, notified their members and also, particularly APNIC, provided a lot of data that we used for our outreach efforts, both in terms of, we sent notes to 99.5%, or sorry, validators, people who operated validators had covered 99.5% of the DNSSEC-related validations that we were seeing according to Geoff Huston’s APNIC Google ad thing and so there was quite a bit of outreach and communication effort. And largely, it seemed to have mostly worked except for Ireland.

So upcoming milestones we have, the Q4 ceremony had generated the revoke KSK-2010 thing, bit, set the revoke bit, that's it. And in 11 January of 2019, the root zone will be published with that and that's sort of significant because that'll be, I believe, the new maximum segment – or, sorry – response size for the KSK rollover. So it might cause some other folks to have some issues just because the response size gets big, although we've seen… Amusingly, we just saw someone who had a response size of 26,000 bites so we don't think response size is that big of an issue for most folks.

In 22 March of 2019, that will be the first root zone since 2010 that will not have the KSK-2010 in it and then in Q3 of 2019, we will remove the 2010 KSK from the HSM on the East Coast and in Q4, we will get rid of the KSK-2010 for history. There were still some t-shirts with the old KSK on it and they may become [your] items. Next slide, please.

UNIDENTIFIED FEMALE: That's it.

DAVID CONRAD: Then that is it. So any questions, comments? I actually just want to say thanks to everyone for the efforts that you all had taken with regards to the KSK roll. I know a number of you had done outreach and tried to get people to be aware of it, so thank you. Now questions. Sorry. Yes, sir.

UNIDENTIFIED MALE: Do you expect to have another KSK rollover at some point in the future, and if yes, you think it's going to cost as much in resource and effort to warn people?

DAVID CONRAD: So I was thinking maybe next week. It seems like a good time. No, so yes. We're going to have another rollover. We are, the current DPS states that we will be rolling again in five years. There has been some discussion within the community of changing the periodicity of the rollover and what we're planning on doing – well, actually I should say what I believe IANA is going to be doing is sending out a note saying, "Here's our suggested time for the next rollover. As a straw man proposal requesting a public comment phase to get input on what the community thinks the right periodicity should be for the KSK rollover."

I should also say that OCTO, the Office of the CTO of ICANN, is handing over responsibility for the KSK roll to the IANA functions team. We, OCTO, took over the KSK roll, primarily because IANA had a few other things to do with the transition and all that sort of stuff. But it's actually IANA's job and we did it simply because they didn't have the resources. We are now transitioning it over to them, so future KSK rolls be done by IANA. It's possible that if we're looking at things like an algorithm role or something like that, that my team, OCTO, may get involved again. But at this stage, when we're just doing the standard roll of just updating the key, that'll be done by IANA.

UNIDENTIFIED MALE: Thank you very much, David, [inaudible]. This is reviewing the KSK rollover, and the KSK rollover cannot be happen, only the effort by the ICANN OCTO but it involved a tremendous amount of the corporation from the community. In that time, David presented not only the technical side but the engagement and communication side. If you had anything, you found from that kind of activity, we are more than happy to be shared and then that's from us. I'm asking you for that, Steve. No, Fred, sorry.

FRED BAKER: Well, so do you have any idea what interval the IANA will be recommending?

DAVID CONRAD: Not at this time. I know there's been discussion of one year, two years, three years. I heard yesterday that someone thinks ten years is the appropriate time. So it wasn't IANA. It was someone from the community. So my guess is that IANA will probably propose a straw man, maybe two or three years, and the whole intent of that is to initiate the discussion within the community to come up with a consensus.

UNIDENTIFIED MALE: I think there is no particular procedure or process to consider and set forth such kind of policy so that we need to create something appropriate for that, right?

DAVID CONRAD: So yeah, my impression is that the plan is to use the public comment mechanism as the process.

UNIDENTIFIED MALE: So that’s mainly to change the DPS for … Okay, thank you.

DAVID CONRAD: Exactly.

LARS-JOHAN LIMAN: I would just like to officially thank you very much for actually doing this and for all the tremendous work you put into this. This was a feat. It was something extra. I now [inaudible] of these things, once when the root zone was first signed and now when KSK was rolled. Both of these things went very smooth and I would like to thank you.

And as a comment to procedures, please don’t overcomplicate this. I think David’s proposal is perfectly in line. We have ways to handle questions that are not part of an exact process and we should use them. Thank you.

DAVID CONRAD: Well, thank you for your thanks but it was definitely a team effort and with the help of Verisign, and APNIC and the RIRs, and everybody else, this was definitely a community effort so thanks to everyone.

And on the process, my assumption is that there would be a public comment thing and then that would generate a recommendation by IANA to ICANN Board who would probably then throw it to SSAC and RSSAC and everybody else, appropriate, maybe RZERC, I don't know, for their input and advice and then the Board would make the final decision. But that's just my guess at this stage. We haven't worked out those details.

LARS-JOHAN LIMAN: That was intended as a "you" in the plural, very much including yourself and I think process, perfect.

UNIDENTIFIED MALE: Thank you. Thank you, David, and for the record, thanks for the rollover to you and the team.

About the public comment, and as you mentioned, this is a process that should be talked about, but I think that's actually an interesting case because this has some technical aspects to it, the frequency of rolling the key. This is just thinking out loud. I was thinking, going to the whole community, to start at least that question might cause, as you said, there are people who might think ten years or 15 years and there are people who might think one year.

I think it's very good to lead the community or provide with some technical facts coming from, first of all, IANA as you said they will do, but maybe SSAC, RSAC, RZERC or other committees who might have technical comments and say, "Okay, this is the considerations from them," or consultations with them and this is what they think. Now rest of the community, based on these facts, what do you think? Because if you start with the whole community at once, there might be a lot of time spent around suggestions which have no technical merit or have technical issues so they are basically moot.

DAVID CONRAD: So would it make sense for the letter, or the draft proposal that IANA puts out, to include sort of a background? Here's our proposal, here's why this proposal, this is what we did in the past with, perhaps, a pointer off to the post-mortem that we're going to be doing on the KSK 2017 roll just as background material so that people who do want to respond to the public comment, of course, they will read all the documentation that we provide because that's what they always do. But do you think that would be a workable approach?

UNIDENTIFIED MALE: Definitely, and then maybe some industry practices because this is unique but in general for public key management practices, things like that. Thank you.

UNIDENTIFIED MALE: Thank you. Any questions or comments in this regard? If not, we will to move ahead.

DAVID CONRAD: So this is the agenda slide. No, this is the next topic. Actually, if you could go back a little bit, it'll explain. So this is a presentation that was given to the Board Technical Committee meeting in January of this year and this is some of the materials that led up to the Board resolution that was passed in Genval – I don't remember the number – called Root Server Strategy.

For a little background, this topic came up to the Board. When was the first, where's Ram? Oh, there you are. When was the first time we raised this topic? It wasn't L.A., was it?

RAM MOHAN: No, I believe it was almost two years ago.

DAVID CONRAD: Right. Okay. So basically, the issue that arose was some concern that based on projections of increase in denial of service traffic that were being experienced by some of the registries, there was some concern that if those types of attacks were targeted at the Root Server System, that there was a potential risk that the Root Server System itself could be brought to its knees. So that's a little background and this is a set of slides that was provided to the Board Technical Committee by OCTO, trying to provide some food for thought on that discussion. Next slide, please.

Okay. So I'm sure most of you or all of you know all of this stuff, that the Denial of Service Attack is actually really real. There have been now multiple attacks greater than one terabit per second. The Dyn attack, I believe, was at 1.2. There was an attack against Github at 1.7 and if you actually graph that out, the graph that results looks scarily exponential. And the reality is that there's no reason to believe that the state of art today wouldn't allow for even larger attacks. The number of large-scale botnets seem to be increasing. The number of elements within those botnets also seems to be increasing.

Back in January, there were 955 root server instances and if you do the math and assume the impossibility of equal distribution, that's one gigabit per instance. I know for a fact that many of the L-single, the single unit for the L-server system would not be able to handle a gigabit into the networks because the way those instances are deployed, they rely on the network provision of relatively small sites that aren't really provisioned for that much traffic.

So the assertion there, that the root system at this point in time is not safe against a very large scale, Denial of Service from a widely distributed set of sources. Next slide, please.

Why are we in this situation? Well, one reason is the just abysmal security in IOT devices. They're easily compromised and commandeered into botnets. They're relatively powerful now and connected to relatively high bandwidth connections. For example, the IP cameras that exist need to feed that video out and the Mirai botnet took advantage of that.

In addition, the low project margins mean for poor attention from manufacturers to non-essential areas such as security and manageability so you end up having easily compromised machines with no way of updating them without replacing them. In many cases, it's unlikely that those devices will be patched. Even if they are updatable, very few end users will have the interest or understanding of how to update them unless they update themselves and that implies a service, so that people would have to pay for which sort of implies that people won't pay for it so you get back to the same place of having devices that are unpatched.

Spoofed addresses remains an issue. They make reflection attacks easy and DNS servers are an ideal vector for reflection attacks. BCP 38 which would deal with spoofing is still relatively low, although I gather there has been some increased attention to BCP 38 as probably a result of the manners work that ISOC is doing.

And there's this problem associated with deploying BCP 38 is it suffers [inaudible] of externalities. The ISPs that are deploying don't get a whole lot of benefit for doing it and potentially significant costs. Next slide, please.

There we go. So the implications of a massive Denial of Service attack. So obviously, Anycast instances that are hanging on off thin pipes are going to fail first and if the failure is sufficient time, then it's going to cause the traffic that was directed to that instance to go to the next announcement that they can see, which can cause potentially a cascade failure and if the other ones goes down, then It may cause the traffic to swing back and you get this interesting little ping pong effect.

It's possible for Denial of Service mitigation and scrubbing services to be applied, but if the attack is intelligent, what's the difference between a good query and a bad query?

Volumetric attacks can saturate incoming bandwidth on the path to the Root Server. This can limit the bandwidth hitting the root server instances, but it has the same effect for clients depending on that incoming bandwidth. If you're flooding a root server instance with video, the capacity of the link is filled up and your DNS queries probably aren't going to get through.

The other issue that actually isn't mentioned on the slide is that you don't actually need to take out all of the instances, of all of the root servers. Depending on what your intent is, if you're just trying to get people's attention, then you can take out the important bits, whatever the important bits actually mean.

The Dyn attack did not take out all of Dyn's instances across their network. It only took out ones in interesting places like New York and Washington. This wasn't probably intentional. It's just the with things worked because that's where most of the traffic went and since most of the traffic then is now being blocked from getting access to the service, it got the attention of the press and the companies that we're reliant upon that service.

One of the questions that had been raised is, well, can TTLs, long TTLs address this? And it potentially can but the question then becomes how long can an attack be sustained, particularly, if it is an attack that's spread across millions of machines around the planet that are doing relatively low transmission rates. The nice thing about how easy it is to have botnets these days is you can just keep making more.

So worst-case scenario, massive botnets sending valid queries for existing domain names. If you have a billion machines, that's not outside of the realm of possibility with the wonderful world of IOT then, and you have a whole bunch of names to query against, then you just cycle through those names and send these queries, and the whole point of it is if you do, a simple volumetric attack, it's hard to imagine that not actually succeeding. Next slide please.

So looking at how to address this from the icann.org perspective, our view is that expanding the L-root is necessary but not sufficient. I, personally, believe that it's an architectural issue within the DNS itself. The idea of a single point in which all traffic must flow, that it doesn't usually scale very well and yes, we have mitigated that somewhat by using Anycast, but still there is a logical single point that has to be addressed.

Each Anycast instance acts as a routing catchment and Anycast localizes attack flows into those catchments. So as long as the interface is up and announcing things, then it'll absorb traffic but when it goes down, then the announcement will stop and traffic flows into other instances, and lots, of course, it's static0-ly nailed up. But that just means that you're redirecting people into a black hole.

Expanding root server system capacity to stay ahead of Denial of Service capacity is, in our view – and when I say "our", OCTO – an unwinnable race. It's always going to be cheaper to add more attack capacity than is it for organizations, particularly, organizations who are providing the service, not as part of their business but as sort of an additional capability, particularly to the Internet at-large than they'll be able to cope. And in fact, I would make the argument that the increasing Denial of Service capacity is faster than anyone can cope these days.

icann.org must make a strategic decision on how much to invest in L-root. There is need for sufficient capacity for legitimacy and responsibility, but it's not an easy answer to determine how much that sufficient capacity actually is.

So Steve Crocker coined the term "Hyper local" and this is a known idea. It's been around for quite some time, simply the root zone is small enough to replicate it into the individual cursor resolvers and that has a lot of interesting benefits, things like it removes the resolver's dependence on the root system. It improves performance, reduces the amount of junk traffic going to the root servers. It's actually documented in RFC 7706. One approach to do this is documented in 7706 and it's implementable today. In fact, a number of people have actually implemented it.

It does have a disadvantage, a number of disadvantages. One is that we lose visibility at the root servers to see what queries are being sent. There also have been concerns about the ability, the likelihood of people misconfiguring things. One of the experiences with the KSK rollover was that there were a lot of folks who didn't configure KSK 2017 when they needed to and there's always the cases of renumbered root server addresses still getting queries after more than a decade, sometimes more than two decades.

So it's not a perfect solution, but it does have some additional advantages and it's also something that anyone can deploy if they so choose.

If we do go down that route, clearly, we would probably need a different way of distributing the root zone. Right now, the provisioning system allocates out to a small number, a relatively fixed small number, of clients and if a zillion recursor resolvers are trying to fetch the root zone, then we'll probably need to have a completely different distribution structure.

One of the options we suggested to the Board was that we could explore a high performance and resilient distribution network to allow for widespread adoption of the hyperlocal concept. So that's something that was part of the resolution and that was done in general. Next slide, please.

Another useful thing. NSEC aggressive use, it can actually, as I'm sure you're all aware, it provides a way for authoritative servers to say "no names exist between two labels" and this actually has some, it can remove an entire class of Denial of Service attacks, the ones that do random string queries into a domain name server.

Of course, the bad guys would simply just use a different Denial of Service attack, but it would be an incremental improvement, so that was also a suggestion that was made. Next slide, please.

Whoa, connection lost. I guess they're taking down the network. No.

UNIDENTIFIED MALE: DDoS.

DAVID CONRAD: DDoS, yes.

UNIDENTIFIED MALE: Do you have the printed material after this?

DAVID CONRAD: No, I think that was the last slide.

UNIDENTIFIED MALE: Then that's good to ask the people for comments, questions.

DAVID CONRAD: Actually, yeah. We can go to questions, comments, screams of outrage. Yes?

UNIDENTIFIED MALE: Actually, I want to make one comment. It's that from a privacy point of view, this is actually really excellent that you're not seeing all these queries and that TLDs do not see all these queries. So in the modern way of don't collect what you don't need, it is actually a really good thing to have.

DAVID CONRAD: Right. That is an additional benefit of hyperlocal is it removes one tier of queries going to the root servers and with the increased deployment of things like minimization and other privacy enhancing techniques, it's our view, within OCTO at least, that the use of root servers as a vantage point over the long-term is going to be challenging and that we need to look at other alternatives in order to obtain sort of the telemetric data that we need to actually understand how the DNS works.

That's not to say that we won't continue to use the root where we can because reality is that it's going to take a zillion decades to actually get to a point that the root servers aren't seen, at least a very good sample, of what the DNS is actually doing but it is something that we need to take account of moving forward.

UNIDENTIFIED MALE: One thing I didn't understand from your presentation and reading 7706 which was a quite short document – thank you, Warren – is the expected size of the uses of this distribution network. From what I got of 7706, I got a feeling that it could be in the tens of thousands but from your discussion of a high capacity distribution network, I was thinking, "That's what you need to sell a billion [inaudible]."

Obviously, these things start small and then grow, but what kind of sizes are you thinking of?

DAVID CONRAD: So there are, the last time I looked, something like 10 million open recursor resolvers out there and those are, in the view of some, actually misconfigured. So it's not the default to be open so the people actually have to go out of their way to make them relevant. So the actual number of resolvers on the Internet is probably higher, maybe. Who knows? Then the question is are those actually really resolvers or are they just some sort of malware that's pretending to be a resolver?

There's someone behind me who probably wants to make a comment.

WARREN KUMARI: I am actually part of the [TGI] because I sat up here earlier.

Yes, the RFC does say that this is designed specifically for people with high latencies and kind of implies that maybe it won't be widely deployed. To some extent, that was something that just needed to be said to get the document through.

[UNIDENTIFIED MALE]: That never happened to anyone before.

WARREN KUMARI: Yeah. So my personal view, just me, is that it makes sense for a lot of people to do this. I don't see a reason not to really, but that's just my view.

DAVID CONRAD: Right. So just let me finish with Harold's question. The assumption that I guess our team was making was that it would start small and build up and we wanted to develop a distribution network, or not necessarily develop, but coordinate or work with or suggest or work with the community to identify a distribution mechanism that would scale to millions because ultimately, it could go that direction, right?

One of the things that we've been playing with is deploying the root zone into https so that it can be put over standard content distribution networks and we actually have that stood up on the transfer servers that we operate. So ultimately, the assumption is that we have to assume that long-term, we're looking at millions.

UNIDENTIFIED MALE: [inaudible], do you need to make an immediate comment? Okay, then I'd like to have Alyssa for the immediate comment.

ALYSSA COOPER: Alyssa Cooper, IETF Chair. I just wanted to let people know that the RSSAC liaison has flagged this to the IAB as something to look at a little bit more closely, so as you say, ICANN has only limited options but it's not only ICANN that is able to potentially make some progress here, so I think it's probably something that the IAB is going to be looking at and talking about and trying to figure out if there's further protocol development work that could be useful in the IETF or anything else the IAB can do to look at this problem from a more holistic perspective. So we should stay in touch about that, but it's not like everyone is expecting ICANN to solve this problem by itself.

UNIDENTIFIED MALE: So the first thing I would do would be to measure how much BIND install can handle of [inaudible] requests.

DAVID CONRAD: Just in response to Alyssa, yeah, one of the topics that was discussed on a couple of occasions is what can ICANN Org actually do? And we can deploy clusters. We can deploy L out onto a cloud. We can throw more instances around the world. But ultimately, it is our view – OCTO's view – that this has to be a systemic improvement. It's not just, a single organization cannot, should not be assumed to be able to take on all of this issue and that was actually part of the rationale for the actual Genval resolution.

UNIDENTIFIED MALE: Thank you. Josh, thank you for your patience.

[JOSH]: Thank you. There's been some speculation about what the Quad8, Quad9, Quad1 servers, whether they are actually practicing hyper-locality, whether they have an embedded root and keep it up. What do we know about that for sure?

DAVID CONRAD: For sure? Nothing.

[JOSH]: What do we suspect?

DAVID CONRAD: So my suspicion is that some, that probably all of them have the ability to run without the roots. They probably have mechanisms that allow for mirroring of the root zone, but may not actually turn it on unless they need it. But tI must stress that that is a guess.

UNIDENTIFIED MALE \: Thank you.

LARS-JOHAN LIMAN: If there is actually an option that we could look into and I haven't really discussed this or thought it through, we already have a network of roughly 1,000 nodes that already have the root zone on it. It's two clicks away from providing zone transfer. That's the existing server network.

So that could be an option to look into, what we can do, providing a zone that way and the roles of technology for providing that with a signed zone transfer, and I'm not speaking about [inaudible] because that's a symmetric key which doesn't work well in this case, so you would have to go to a [inaudible] key instead. So there are technology already in place to do that, as you know, and explaining to the wider audience.

DAVID CONRAD: So there was some discussion of that internally within OCTO and someone, probably Roy, because he thinks that sort of way, is that, raised the point that if you're relying on the root servers as a distribution network for a mechanism that you use in the event that the root service is not available, then you've sort of created a problem there. But my suspicion is that, and there doesn't have to be, a single solution to this, right? That, in fact, it's probably best not to have a single solution. You don't have defense and depth. You want to actually have multiple ways of providing root service and how that root service is obtained. As long as you have multiple ways, then you don't have to worry about one of those ways getting taken out.

LARS-JOHAN LIMAN: Diversity is good.

UNIDENTIFIED MALE: Thank you, Jay, and [inaudible].

UNIDENTIFIED MALE: Thank you. Forgive my ignorance, David, but has anyone done any calculations on the impact on the existing root servers of multiple resolvers doing IXFRs off them and they know the number resolvers that might do it and those sort of things, and if so, will that be published at any point?

DAVID CONRAD: That's a study that we've been talking about doing, just haven't had the resources to actually get into it. I don't know if other folks have actually done it. Oh, I sense someone next to me again.

WARREN KUMARI: So, yes, we kind of did. It turns out that very much all resolvers have very much all of the root zone very much all the time, right? Because there are things that cause people to go off and do lookups. This means that they go along and they do a bunch of separate queries to do that. Getting the root zone in one blob is the same amount of data. Yes, it is transferred over TCP instead, but it's a much smaller number of sort of separate operations.

So it looks according to the analysis we did a while back, it looks like the load is likely to be probably lower than the current set of lookups but there is definitely some fudge factor there. We did sort of a small simulation and tested it, and it seems to be lower as well by doing the lookups.

UNIDENTIFIED MALE: Okay, so the rationale for having a different distribution network is not load. It is risk. Is that right? Because we're now introducing a new production service effectively on the root servers as someone has just nicely written in the chat [inaudible] to say.

WARREN KUMARI: So one of the original kind of ideas with hyperlocal root is, and also adding additional signatures like A-zone digest to the bottom of the zone, is that it doesn't, it no longer really matters where you get the zone from. Fetching it from the current root servers, it's easy. It works. Fetching it, fire HTTP from something, that might work too. Fetching it, fire –

UNIDENTIFIED MALE: FTP, from FTP.Internet.net.

WARREN KUMARI: Yep, that, a bunch of different distribution things. What's the one that I'm thinking of? All I can think of is block chain at the moment and that's not –

UNIDENTIFIED MALE: Before we get too [inaudible] with the technology, I'm just thinking from the Board perspective, I think the Board needs to understand the reasoning why they would need a distribution network because this is an extremely large amount of money and that it's not actually volume-based. That's not the problem. It is a separate set of problems and those are still to be potentially articulated. Okay. Thank you.

LARS-JOHAN LIMAN: Do you need to have it immediate or you can wait for them?

UNIDENTIFIED FEMALE: This is on a little bit of a different topic because I just wanted to make everybody aware that the SSAC does have a current work item that's been socialized this week regarding IOT devices that describe the risks and the opportunities that pertain to the DNS ecosystem and so while you do still want to consider looking at solutions to the potential risks of these growing DDoS traffic to the root DNS servers, there's also the consideration of being proactive with mitigating the risks at the root of the cause. And yes, the pun is intended.

UNIDENTIFIED MALE: Thank you. Lars-Johan?

LARS-JOHAN LIMAN: Thank you. So I agree with the risk perspective because that's an important point and that, very much, with your [inaudible] regarding diversity. There is also some other things that start to tie into this and that is that we will probably see other traffic flows regarding DNS service in the future, thinking primarily of DNS over HTTP which will probably start in web browsers, but I cannot forecast here that it will probably migrate into other products as well.

That will entirely change the traffic patterns around DNS service and the roles of change the importance of root servers and lower that importance on the current network. And it adds diversity so it may not be an entirely bad thing. That's all.

UNIDENTIFIED MALE: Thank you. I think you … Then who is that? Sorry. I have my own queue here, remote.

CATHY PETERSEN: I'm reading a question from Dan York. I agree with research that RFC8198 can help here. Have you been in touch with resolver vendors about getting RFC8198 implemented?

DAVID CONRAD: Yeah, so we're trying to work with all the, at least the open source resolver vendors. We've also had contacts with non-open source resolver vendors in a bunch of different areas including NSEC Aggressive Use and I believe 7706 and doing zone transfers via HTPs and a couple of other things, basically just seeing what their interests are in deploying these new technologies. So one of the, if I remember correctly, one of the items in the strategy was to look at providing additional support in those areas.

UNIDENTIFIED MALE: Thank you. Warren, Brad, and [inaudible].

WARREN KUMARI: Kind of following up, actually from Dan's question, which made a nice lead-in. So actually, a large number of the current open-source resolvers do aggressive NSEC, BIND and Unbound at least, them well. And so that, I think, helps prevent DoSes on the root.

Another thing which is happening and is deployed in a number of open recursives – sorry, not actually open recursives – open source implementations is serve stale which actually means that if a root, or actually, if any DNS authoritative server is unavailable and there is still information on the cache, even if it's expired, it can still be used as sort of an ancillary last resort.

We hope, possibly, that this sort of thing makes attacking things like the root and authoritative servers less interesting because although it might cause a small bit of delay, it doesn't actually harm, create harm to the users. This means that hopefully attackers have less incentive to do attacks.

UNIDENTIFIED MALE: Thank you. Brad?

BRAD VERD: Really quick, looking for clarification on a few things. But you stated a couple times and I've seen this a couple times as far as things that in the L-root strategy or the IMRS strategy that add more instances and whatnot. I just want to make sure that I understand when instances are added by any root server operator, they're out there announcing the IP space that's owned by that organization. In this case, though, I just want to make sure. Is it OCTO's interpretation that the hyperlocal possibility or mitigation tactic, let's say, is in the scope of the IMRS deployment?

DAVID CONRAD: No. At least our view is that hyperlocal is something that's entirely in the realm of the resolver operator, so it's not something that is related to L-root per se. It's a mechanism by which resolver operators can pull down the root zone and begin to serve that information locally as opposed to referring queries out to the root servers.

UNIDENTIFIED MALE: Thank you. Let me closet the queue by now, then Lars-Johan is the first, the last comment for this section.

LARS-JOHAN LIMAN: Thank you, just I forgot a comment to Jay regarding capacity capitulation for zone transfers. There is some prior [inaudible] in that area but done individually because anyone who runs an Anycast network runs into this problem with distributing zone transfer so if anyone wants to go down that route to do any calculations to look at this problem, I think you should reach out to [inaudible] Anycast operators for some input. Thanks.

UNIDENTIFIED MALE: Thank you very much for the very active and substantial discussion. So let's go next agenda. Let me project it. So Paul?

PAUL WOUTERS: Thanks. I think this is the bigger deck. This is not the five slide deck version. So I'm just wondering if you can bring up the smaller one.

Okay so meanwhile, I'll just start talking for a little bit. So with DNSSEC, we've added a lot of security of the data to the DNS hierarchy but there's still one issue that a lot of opponents of DNSSEC always have and that is that the keys at the top of the hierarchy are really, really powerful and they can do a lot of bad things to their children and sometimes these children/parental relationships are not entirely voluntary. There's strength in the hierarchy because there's only a few people that can make statements about the child, but it does give the specific parent a lot of power.

So there's two attacks that DNSSEC does not help against if you cannot trust your parent. So the first attack is where a parent, so for instance the .org or the root zone, could pretend that there's no NS record or DS record and immediately put some data in there that's [deep sign]. So in the example that might come up soon on the screen, I had an entry at mailarchive.ietf.org and it would be directly served by the .org TLD, so it would completely bypass the ietf.org zone. So there's not currently much we can do against this problem.

So the second attack that's there is, of course, that a parent can just replace an entire zone. So instead of having an NS record and a DS record, it will just replace the zone entirely with itself. That you can see now in the slide.

So how do we address these issues? So next slide please.

So one of the drafts that is currently circulating early in the IETF that has not yet been adopted, so it's very experimental, first idea, is that we add a new flag to the DNS key, and that flag basically says, "I am a good parent. I will never skip my children, so if you see that happening then somebody took control of my private keys and you should never trust us."

The good thing about encoding this as a bit in the DNS key is that if you sent this to the parent in the form of a DS fingerprint, that bit is part of the calculation. So if, at some point, the parent decides to change that, for instance, in a targeted attack to someone specific, then you would actually no longer match the DS record and it would also get flagged as bad use of the key and it would get rejected by a validating resolver.

So this is a really nice way of making sure that a parent cannot skip a child and answer very deeply in the tree with malicious answers. Next slide.

The second thing, what it does is that it provides us to have a method of doing what we call DNSSEC transparency. Currently, if you want to see what certain keys do to see if you can trust them, like let's say has the root key ever signed www.noads.ca, then you have to do a lot of logging of data and it turns out that whenever you look at this, you have to log all of the DNS data to be sure that a DNS key is not maliciously used. And so that's obviously impossible because we invented the DNS, and particularly, because it's hierarchical and we don't' have to store all of this at one place. So an audit of that entire tree is also impossible.

But if we have this flag, then actually, we only need to log a very few limited things. We need to log what the DNSKEY is doing and then we can quickly flag all the violations of it. So it's actually a really good way to start DNSSEC transparency. Next slide.

This was first proposed a couple of months ago and it's been briefly presented at IDS and there were two things that actually came to mind. So the first was actually, "Well, we just did a root key rollover. If we want to add this DNS key flag, we'll have to add another root key rollover for this flag," because, as I said, the DS record of the root key that you would pre-load would change and so this would basically be a key rollover.

But we don't actually need to do a key rollover because if you're at any point in the hierarchy, if you say, "I promise not to skip my child zone," then you're also implicitly saying, "I expect my parent not to skip me because then you would skip that limitation that I'm putting in place."

So okay. So we thought about for the next version, we will put this in a drafts and making this actually explicit, so if you set this flag, you don't expect to be skipped by your parent. And the good thing about that is now we actually don't need anything in the root key to be changed because it is implicit. So if a TLD would say, "I will not skip ay children." It means, really, that the root key can also not skip that TLD. And so these are a little bit of subtle interactions that are now happening, so that's why I thought this would be a good time to at least bring this up at the ICANN Board so people are aware that there might be some technical things coming up that would actually put limitations on what you can put in the root zone later on, and once this gets codified into software, you will never, ever be able to remove this again, so we better, if we're doing this, we better do it right at the first go.

And there's another minor thing is that a lot of the data we actually want to protect in the DNSSEC. It relates to public key, so other than DNS Key and the signatures, we are looking at things like TLSA which is like putting TLS certificates inside the DNS. And we really want to protect those two, but those are usually encoded with the prefix. As you can see in the example here, _443._TLSA.example.com, it would be really annoying if we would have to have delegations for those because example.com would not be allowed to skip _TLSA.example.com and then you would have to add a lot of zones.

So since all of them start with an underscoring prefix is actually really easy. If we could just add and say, "This no skip rule, actually, there's an exemption for underscore domains," and since these normally aren't allowed at the traditional second-level domain registrations, we don't think that that would be a problem. Okay, any questions? Thank you very much. Lars-johan?

LARS-JOHAN LIMAN: I would first like to note, and I'm looking at you for confirmation. These are not problems that kind of came with DNSSEC? These are old systemic DNS problems that were not solved by the NSEC and looking at proposing?

DAVID CONRAD: Correct.

LARS-JOHAN LIMAN: Yeah, just to make sure that no one believes that the problems have their roots in DNSSEC. The second thing is I carefully note the exemption for underscore labels but don't you also run into problems with [inaudible]?

UNIDENTIFIED MALE: Yes, I mention. That's right. So let me just clarify the station and date, so sometimes a TLD will have a clue for one domain and that's actually used for multiple other domains and after that domain vanishes, then technically, the clue should have been removed but then hundreds of domains might go down because they lose their name server. And so often, TLDs will then often disclose and they will sign it themselves or they will keep it in their own zone and since it's in their own zone, it has to be signed. And in this case, that would indeed violate of not deep signing anything. So in this case, that kind of sign glue would then get ignored. So that is an issue that is mentioned in the security considerations off the draft and maybe that needs a little bit of careful thinking, but that is a side effect.

LARS-JOHAN LIMAN: That's not what I was referring to… Actually, I was referring to the NS records at the apex of the zone. If I have example.com and my name servers are ns1.example.com and ns2.example.com, I want to have the A-records in the zone.

UNIDENTIFIED MALE: Those are in your zone so those are not a problem. We've had this discussion three days ago. I'll take it offline with you.

UNIDENTIFIED MALE: Could you repeat the name of the draft?

PAUL WOUTERS: It is draft.wouters.powerbind because it has not yet been adopted by a DNS subworking group.

UNIDENTIFIED MALE: Okay. Any other questions or comments? Thank you very much. So okay, that's the old agenda which is prepared. Any other ones? If you have, anyone has input or questions or comments, I am very happy to take it.

Okay, P. Wouters. That's why I couldn't find it.

WARREN KUMARI: Sorry, Warren Kumari again. I have sort of a general question for the board. Are these useful or what could we do to make these sort of meetings with you more useful? Is the level of the discussion appropriate or less technical, more technical? Basically, how could we make these better and more useful [inaudible] for you?

UNIDENTIFIED MALE: Okay. Thank you very much for that question. If anyone from the Board have comment. Okay, Cherine?

CHERINE CHALABY: So I think this is a really very good question and I remember Steve Crocker who set up this group together, putting together the people from the [inaudible], I think the IETF, RSAC, SSAC, now the Board Technical Committee and OCTO together. And if you look at the discussion today, it was a mixture of possibly a good discussion around the KSK rollover because that was rolled over and everybody got engaged and then there were a couple of good presentations on DDoS and DNSSEC.

I think we ought to keep on revisiting the agenda. And I think it is important that I find meetings that are issue-based where maybe we get together to try and advance the issue is much more valuable than presentation-based. That would be my suggestion. I really like that meeting. I think it's a good gathering but I think it is important to keep on assessing how to make it more relevant and more effective for all of us. And I think an issue-based one where we bring an issue to the table and really have a deep discussion around it will be, to my own mind, as valuable as just an informative meeting, even more value for me. So I don't know if everybody has a different view, but that's a suggestion.

UNIDENTIFIED MALE: Thank you, Cherine. Nigel.

NIGEL ROBERTS: Nigel Roberts, newly-minted Board member. I can only speak for myself. I have a little bit of a technical history, though not as much as people in this room. I loved every minute of it. Thank you very much.

UNIDENTIFIED MALE: Thank you very much. I'd like to make comments to Warren's question. Yes, that's a very good question to put and then I'm actually echoing what chorine said. First of all, the presentation and the information which wax exchanged in this meeting is brilliant. Then we definitely need to have the other technical part of the Internet ecosystem to get together and cooperate and then we need to fulfill our own mandate at ICANN and ICANN Board so that's really appreciated. Then maybe your question is representing maybe to set up the [inaudible] a little bit, lacking some focus or some proper working mode or something.

In that time, the Board Technical Committee take roll of, make it clarified then let me work on that. I will need your input for that afterward. Thank you very much for your question. Very good. Ram?

RAM MOHAN: Warren, this is Ram. You and I have had some offline conversations on this as well. During my time on the board, I came in to the TEG meetings with a lot of anticipation for dialogue and to get to a deeper understanding of a specific topic, or perhaps, a couple of specific topics. And as time has gone by, it seems like they are, these meetings and the presentation have been somewhat random and it's not clear whether there is actually a good level of absorption and/or understanding an/or utility long-term from the Board side.

So, I think it's a noble idea whose implementation still has a long way to go, if it should still stay alive.

UNIDENTIFIED MALE: So, what it says on the webpage for the purpose of the TEG, which is a subset is "The purpose of the TLG is to connect the ICANN Board with appropriate sources of technical advice on specific matters pertaining to ICANN activities." What I think might make some of this more useful is if the Board knew, and maybe it does that they can reach out and be like, "We would like to talk more about this," or "We would like some more information or input on this sort of thing," and maybe that already is known, but I'm not quite sure how that information would flow or if Board members are comfortable. I know stuff about blargh, and it would be useful for me to have [technical] advice or information on that.

It's also a difficult thing sometimes to stand up and say, "I don't' know much about this and I probably should. Can we get some advice and guidance on it?" Possibly even just having sort of like a somewhat closed mailing list or the awfullest thing where Board members can be like, "somebody was mentioning about this. Can I get some technical input?" I don't know. Maybe that's best handled within ICANN now that I think of it, so David might punch me.

UNIDENTIFIED MALE: Thank you, Ram, for the immediate response.

RAM MOHAN: Thanks. Just very briefly, one, I think that is a very good response. I actually think that the Board as a whole, that individual Board members may have value in getting deep into the details and understanding the various specific pieces. But if this group really wants to have, in my opinion, an effective medium into the Board, then I thick the presentations have to shift from information about the tech to information about the risks and the pros and the cons of going a certain path and the relevance for, and the reason why the Board ought to be thinking about this a certain way.

If you look at, one of the topics that has come up recently, for example, is [do]. What's the value of the Board to learn all the details of [do]? The real value ought to be for this group to come back and say, "You may be hearing about [do], you may be hearing about [blocktree] and whatever. These may be areas where there may be a positive impact or a negative impact and then the Board, I think, can look at that and say, "Okay, how does that fit into its mission, its remit and then what can it do about it?"

Right now, the presentations are from where I used to sit, nice to have but I could have picked that up from watching these at an IETF or a DNS OARC presentation.

UNIDENTIFIED MALE: Thank you very much. I'd like to close the queue. We have the four comments on the queue. [inaudible] first.

UNIDENTIFIED FEMALE: Ram just stated exactly what I was going to comment on is that what the Board really needs to understand overall is the risks and one of the things I think would be a fit, again, this is my very first tie sitting as the SSAC liaison to the Board, but as I look at Harold being the liaison to the IAB, Kaveh liaison to the RSAC, myself LIAISON for the SSAC.

I look at, looking at the work that we can all do in the respective communities that are part of the multi-stakeholder environment and as we try to address the technical issues and understand where the risks are to ICANN's mission, I think having those articulated in some way will be quite useful in these meetings.

UNIDENTIFIED MALE: Thank you. David?

DAVID CONRAD: I'm actually going to cede my time to [inaudible] because I think he was going to say what I was going to say.

UNIDENTIFIED MALE: Yes. Thank you, David. I think this is something that we have already start looking at on how we create interaction between the TLD, the TED, the BTC so that this is addressed.

I think as and the [inaudible] are just, the BTC will have a role in defining the agenda. I think on the other hand, the TEG can also suggest topics, but it's the way that the topic is framed that will make it useful for the Board by kind of [inaudible], "We are raising this because we think this my have impact on ICANN mission or need for the review from the organization." So those two aspects are going to be taken into consideration in what we would propose and probably what would be discussed within the BTC and also within the TEG to try to document that and have it as a way of working with. The BTC just is a new element of the Board so it's trying to find it sway in all of this to make it current. So thanks.

UNIDENTIFIED MALE: Thank you. Thank you very much, [inaudible]. Then I put Cherine for the last of the queue and ask you for the conclude this meeting before that. Daniel, please.

DANIEL DARDAILLER: Yes, Daniel Dardailler, W3C. I want to point out to another benefit of this presentation, this kind of presentation and meeting for people like me and probably people from ITU and [inaudible] that are not, like people from ITF and RSAC, very expert in all that. We learn a lot by hearing all these technical discussions but that relate to what we understand about whatever the key signing and DNS and we can relate that information into our community. And when people ask me "what about this KSK rollover?" I understand it better because I hear David talking about it simply.

UNIDENTIFIED MALE: Thank you. [inaudible]?

UNIDENTIFIED MALE: So now recently departed from the Board, looking at this from the Board perspective and now looking at this from a different perspective, I think that what should be done is that, and I think that is actually reflected in almost all the comments here is that there should be a long, hard look on what is the actual purpose of this and what should actually come out of this.

I think that we didn't do a good job recently to actually do that and I hope the new Board is better than we are and I'm sure that they are much better than we were. And it looks at what does this actually, what should this actually achieve and adjust accordingly.

And with that, it's like Ram said. If it's presentation stuff, can it be gotten somewhere else or being part of the normal agenda of the ICANN meeting, for instance, the tech day or so on, then the question of what is the actual role of the TEG should be looked at very hard.

UNIDENTIFIED MALE: Thank you very much for the viable suggestion. Yes, [Adiel] said that the BTC is just one year old and then TEG is a new set-up for us to work better, then that's for the Board's responsibility to have it better. Then Cherine.

CHERINE CHALABY: Thank you. We were just talking about the purpose of the TLG versus the TEG, and it's interesting. The TLG is clearly written, the purpose of the TLG is to connect the ICANN Board – not ICANN – with appropriate sources of technical advice on specific matters pertinent to ICANN activities and the TLG and the TEG within that says the TEG, which I suspect is this meeting is focused on forward. So I am confused, to be honest, is this a TLG or TEG meeting? Which one is it?

DAVID CONRAD: so the TLG is a subset of the TEG, so if the TEG is meeting, it's including the TLG. So this is the TEG and BTC joint meeting. So it's all of the above.

CHERINE CHALABY: Okay. So I'm thinking, "Okay, so the Board is connecting, therefore, with this group even though it's meting between each other." Yes?

DAVID CONRAD: yeah.

CHERINE CAHABY: So I'm thinking, "Okay, so this is beneficial very much on an individual basis. There's no doubt." Nigel, he loved the presentation. I loved the presentation. But I don't know what to take to the ICANN Board from here. I'd love to be able to take something because if the terms of reference, this is the board connecting so I would like to be able to take something to the Board and say, "The Board met and here is X, Y and zed." Or am I getting this wrong?

DAVID CONRAD: No. Well, so one of the challenges with the TEG, the Technical Experts Group, has been trying to figure out a way to make it the most beneficial to the Board as we can and we've tried a bunch of different approaches back when Steve was Chair and then subsequently. None of them have fit quite right. None of them, at least in my view, have been quite beneficial in providing the kind of information the Board needs on technical-related matters.

So one of the things that's actually that I've tasked Adiel with is to come up with a proposal that will work through the BTC to figure out a way to improve the TEG/TLG and the interactions with the Board, and we're going to be working with that. We'll have something done by Kobe, which is the next scheduled TEG meeting, at least, I hope. No pressure, Adiel.

So obviously, I would be interested in any input. Board members would have to figure out how they think this technical channel can best serve their needs, but we will come up with a draft proposal and work through the BTC to provide that to the Board.

CHERINE CHALABY: Can I continue a little bit? Yes, because this is very, very helpful. So I think the Board Technical Committee represents the Board on technical matters. So I think, rather than just the word "Board", I think it's the board Technical Committee. Now the Board Technical Committee has a list of priorities that it needs to recommend to the Board.

For example, we receive recommendation on RSSAC 37 and 38. We have a discussion with SSAC on [inaudible] named collisions and things. There are other things. A suggestion, perhaps, is one of these priorities, not everything, but one of these priorities ought to, for example, be part of this discussion so that the Board Technical Committee seeks the input of the wider technical community and then we'll be very helpful because when the Board Technical Committee comes back to the Board and says, "Not just us and OCTO thought about that, but we also discussed with the wider technical community, with the TEG, and here is the input from them," that would be very helpful.

DAVID CONRAD: So one of the decisions, I believe, we made or Kaveh made – he ran away – was that the public BTC session was going to be with the TEG, so the intent was that this, the BTC, has public and nonpublic sessions and that the public sessions would actually be with the TEG so that's one step in the direction that you're suggesting and we can look at additional ways of sort of facilitating that communication, the technical communication to the community related to the Board through these sorts of sessions.

UNIDENTIFIED MALE: Thank you very much, David. I actually –

CHERINE CHALABY: Sorry, I'm not going to go over this until I understand it a bit more. Sorry. So you're saying there's a private session and a public session? And you say we have different agendas for each?

DAVID CONRAD: Yes.

CHERINE CHALABY: And a different purpose?

DAVID CONRAD: Well, they're all related to trying to – well, I don't want to speak for [inaudible] here, but the charter of the BTC is to be the technical input into the Board. There are some times where that is a private session where there's discussions that the Board may want to hold in private related to technical issues, and then there's stuff that doesn't need to be private which is made public.

CHERINE CHALABY: Okay, thank you [inaudible].

UNIDENTIFIED MALE: Okay, thank you very much. Cherine, do you have any last words?

CHERINE CHALABY: This is really a worthwhile habit. This is fantastic forum where we all get together so I'm not being critical. I just want to make it so much more valuable to the Board that we can really take something back to the Board. That's my only issue, but this is an excellent gathering and excellent information, excellent presentation. Thank you.

UNIDENTIFIED MALE: Thank you. Warren?

WARREN KUMARI: And I'm just making this up on the fly, and I think it's David's group so hopefully he's okay with this. Possibly what might be useful is the person from [ETSY] whose name I've completely forgotten said that it was useful [inaudible] – W3C. Wow, I can't even get the groups right. – is that it was useful for sharing information between the groups.

Possibly would it be useful for us as the TEG to meet at some point before this meeting to do these sorts of presentations, which are more information-sharing kind of between groups. Not official, just, "Hey, these are the sorts of things we're working on." And then save the board type meetings for these other risks of things that might happen. As I'm saying this, I'm realizing that might actually be more of an IAB thing than something but I don't know.

UNIDENTIFIED MALE: Thank you very much, Warren. Just to reply to that, I think we should consider that all the TEG meetings are of that kind which is mostly about input from the DNS root zone expert. At some point in the past, we've been discussing items that are more in the area of [inaudible] like DOI and things lie that, so it's not generalized but I'm happy with having private meeting with you guys.

UNIDENTIFIED MALE: Thank you. Jay?

JAY DALEY: Yeah. If I could just make my annual plea to the other technical people around the table. Please think about the strategic issues hat the Board needs to know about. Please don't get into the depths of the technology. We do this too often and it's all fun and we love it, but we're really here to help the Board understand the future.

UNIDENTIFIED MALE: Thank you very, Jay, and let me conclude. Thank you very much for the various input and then the – I am actually, as the incoming chair, in the conversation with David and others, how we can get better. All input along that context, I really appreciate the input from everyone. Thank you very much. It's a really good meeting and the meeting adjourned. See you next. Thank you.

UNIDENTIFIED MALE: Thank you.

UNIDENTIFIED FEMALE: Thank you, everyone. All the presentation materials have been uploaded to the public schedule. The recording for this session and transcripts will also be uploaded to the public schedule shortly. Thank you.

[END OF TRANSCRIPTION]