SAN JUAN – Joint Meeting: ICANN Board & TEG
Wednesday, March 14, 2018 – 17:00 to 18:30 AST
ICANN61 | San Juan, Puerto Rico

CATHY PETERSEN: Hello everyone. Those in the TEG and the Board and the BTC, please feel free to sit up here at the head table, thank you.

RAM MOHAN: This is Ram, I'll invite board members to come up to the table as well, this is TEG and the Board.

ADIEL AKPLOGAN: Welcome everyone, welcome TEG members and Board members to this session, I will be chairing as David is not available now. So, we have three main topics on the agenda for today's session; we have DNS privacy, the current states of work on DNS privacy, we have DNS capture analysis, what are the monitoring and measurement that has been done on DNS data, what are the types of analysis done, and this will be done by Matt and Mauricio from ICANN perspective, ICANN org., and the last formal presentation by Jay Daley, it will be on domain name and website classification.

We don't have [inaudible] from now, but if there is any topic you have, write them down and we can take them and discuss them on time.

So, the TEG and the Board met at two ICANN meetings, for last year we had a meeting during the policy forum, so twice a year roughly to discuss some of the topics that are of interest for the TEG, but also for the Board.

Before we start, I would like to thank all of the presenters because they have volunteered with a very short period of time to work on this presentation and be ready to deliver them, so thank you very much for that effort. So with that -- yes, Kaveh?

KAVEH RANJBAR: May I suggest a very quick AOB, that's about future of TEG meetings and Board technical committee, thank you.

ADIEL AKPLOGAN: Yes, thank you. So, one OAB, TEG and Board Technical Community; there has been some discussion about how to structure this going forward. So, we will start right away with Tim Wicinski on his presentation on DNS privacy. Tim?

TIM WICINSKI: Great, thank you, Adiel. Tim Wicinski, some of you know me, I'm the co-chair of the DNS privacy working group for the ITF, as well as one of the co-chairs for DNS operations and also, I've been appointed as one of their technical liaisons. To give you a little overview; to talk about the problem space, why DNS privacy is important, the current state of our technical standards, the implementation status, because that's always kind of key, and a big thing for me is operational deployment. In my day job I work as an infrastructure architect at Sales Force, so we're a very large enterprise company, and we're doing lots of crazy stuff with DNS, so as an operation thing I really care about operational deployment of things. And also, future directions; where we think the ITF is going with this, where we think the world is going with this.

So, as we know, DNS is 30 years old, there's a lot of information that gets leaked, for example, the fully qualified domain name of every lookup gets sent to the root name servers. Some requests, of course, expose too much, much like people put their names, our company gives me my laptop name; I have no choice over that, and of course that gets leaked out to the Internet. And then there's things like EDNF client sub-net that all the CDN networks and folks use to give us better geo-location. That leaks data, and there's privacy concerns about that, so on the privacy side the folks are very concerned about that.

So, from it ITF side, a little history; July 2013 is what everyone called the summer of Snowden, and that's when ITF published 7258 where we talked about pervasive monitoring, it's an attack on Internet and organizations. But, under the radar a little bit was in 2016 the EU approved the GDPR specifications, and they go live in May of this year, so compliance for these is a big deal.

And at my workplace, there is a large number of people working very hard on this, because one thing we do store is laptop company customer information, and so we're scrubbing everything from logs to databases to everything, so this is not a trivial problem. Now, what's interesting is I don't think that GDPR folks have really understand DNS yet, so I don't think they've gotten to us yet; they're still worried about websites and things like that. So, that's probably on the good side.

So technically, on the technical standards, DNS security, the DNSSEC stuff was published in March of 1999, and really to try to sign DNS zones and validate sign zones, but it's always lacked a must-have, it's always lacked that killer thing that makes us want to do it.

One of the things that I think is big is the DNS authentication of named entities work that's been going on in the ITF, where basically we embed certificate keys into sign zones to validate with browsers and servers authenticated certificates, but there

hasn't been a lot of rollout of DNSSEC. ICANN is the big driver of this, with the gTLDs. The US government is requiring it as well as Germany and Netherlands; we have a giant requiRament at my employer to do DNSSEC because of the US government requiRaments.

And we're running to sort of these interesting large scaling problems that one of my coworkers presented at the last DNS or just a couple of days ago, and not everybody is doing DNSSEC, a lot of the Cloud providers just don't care, and the DNS vendors, they support it, but what they tell us is, "We support what our customers want, and our customers haven't really asked for it." And outside of the infrastructure folks, the only folks that are really trying to do any sort of enterprise adoption is Cloud Flair, and we're probably going to be the second one, and it's highly scary to us actually.

DNSSEC validation, this is basically how you validate your lookups, and this is all done at your DNS resolvers, and it's a very low percentage of the user population. The Googles do it, the DNS on the 8's as I call it, or now the quad-9 folks who now DNS on the 9's, they both do DNSSEC validation and they're probably the few folks of any sort of large-scale that do it.

About a month ago, Jeff Houston of APNIC, who all know and love, wrote a nice post about peak DNSSEC, he's feeling looking

at their data that the validation numbers are starting to drop, and he's wondering if we've reached peak DNSSEC and things are going to start to drive.

And honestly, if you look at the business constituency at ICANN, they just avoid the problem completely because behind everyone's firewalls there's so much dirty laundry, it's an embarrassment to get caught in it, and we're just as guilty of it, but you could probably go to any company, look behind their firewall, and they're using root zones they create themselves so they don't worry about leak at all; if we create this root zone that's not in the root TLD, that means that any kind of leakage that happens, because they just assume it will, people can't actually work backwards and find us, so that's not an uncommon thing, I've heard that from multiple large-sized companies and it's kind of frightening.

And with zone signing, it's really limited to what I consider the Internet infrastructure companies, the TLD's, the ICANN approach drives it's at the root. Some of the larger vendors, folks like Google, but even Google doesn't sign their zones, and it's very interesting, they have a very hand wavy idea; they do validation, but they won't sign their zones, it's like they like to wave their hands about it.

These are sort of going, we've gone in a different direction, there we go, apologies, I used the wrong button.

Other work that's gone on in privacy, of course, this stuff is non standards, DNS curve, DNS crypt, the DNC, the .onion stuff, anybody of course in ICANN world Ramembers .onion, that was defined as a used name inside the IDF made solely because they wanted to get an L certificate, and that was the only way they could get it done, and honestly there's like several other of these type of torrent type things that want to do the same thing, that we're basically trying to ignore in the ITF at this moment.

On the standard side, we've done a few things. The query name minimization; stop trying to send the full name to the root name servers, which actually helps us in the GDPR situation, so we were kind of ahead of the curve there. And that's slowly getting rolled out, we're starting to see the tool being supported.

The next one was the DNS over TLS using port A53, instead of 53, and it's basically TCP based, and it does use TLS but we kind of walked around that authentication light; can we trust that the TL cert that we're seeing from this sort of customer? And the [inaudible] working group is basically focusing on this problem.

And in [inaudible], our focus was always step-wide solutions; we could not boil this ocean. DNS is so embedded in our

infrastructure that to try to start over or to build something totally was just doomed to failure. We focused on the stop to resolver, and that's the most technical; that's your browser or your laptop talking to your recursive server in-house, and that does reveal the most information. To us, the harder problem is the recursive server talking to the authoritative servers.

Now, the TLD's of course are one set of authoritative servers, most of the companies' domain servers are another set, and we feel there's a non-technical solution involved here, this is like layer 9 type stuff. We don't think it's going to be easy to get all of this without something else happening. And we're also trying to track implementation and usage, because one of the big things is; we want to see the stuff being deployed, we want to see people using it, we want to see if it's actually useful.

And so, there is several different DNS over TLS clients, and there's actually some trustworthy DNS over TLS recursive servers, folks like unbound and folks like that have actually done some work, and there's actually been some interesting work in the mobile space. The Android folks have literally actually have code in the system that's committed but not pushed out that basically supports this DNS over TLS, and I've seen an example of this on the iPhone side as well, but it's not coming from Apple.

And there's several different clients and forwarders, and the big one in the server side has been Stubby, which is a variation on the GetDNS as well as Knot and Unbound, there's been great work there. As I said about the mobile stuff, it's committed but not released, and I've seen some demos of that and the charts there, the DNSprivacy.org actually give you some great detail if you're really curious.

So, my big thing is operational deployment; I want to see stuff, and we want to see stuff deployed, I want to see folk using it, and I believe this is a the thing, the side of ITF standard that no one uses, most people in the ITF find that fairly abhorrent because we want to put out standards that actually get deployed and get used, and if they're not then that's just a 0 to me. It's the user awareness, it's a hard issue to drive on the user awareness side, because there's a bit of a complication there, but I do feel mobile is going to drive this, much like mobile is driving everything on the Internet right now. All the traffic is mobile, everything is going mobile.

And we do see, at least I see, the tangible benefit for the business constituency that they may not see yet, which is this sort of shared Internet server infrastructure. There's so many companies that are in AWS, they're in Google Cloud, they're in Rackspace, they're in OVH, they're in these sort of spaces and

they're sharing the same servers, the resolvers, they're seeing the shared network infrastructure having something like this is definitely something that they haven't started thinking about, but the security people are starting to think about it, and so that's going to be a place where this becomes a bigger sort of lever for folks. And the folks I've seen deploying the DNS over TLS stuff is the Quad-9 folks, and they're really doing deployment at scale, and that concerns me, because I definitely—if it's not going to really take off then we've got to do something about it.

Future directions; GDPR is happening, lots of area of course on this stuff, the client sub-net, the DNS logs, the specific transparency process that's going on. There is a little bit of leakage there. One thing that's also going on in the ITF is the DNS over HTTP process, which is attempting to do DNS over port 443, and actually I feel that that will probably be more successful than DNS over TLS because everybody does stuff over the web, right? Everybody does, and there's already people in Asia-PAC that's doing that, because it gets around all the firewalls; everybody has to talk over port 443, everybody has to talk securely, here's how we're going to solve that problem.

So, I feel that on the privacy side, that's going to be the place, or in terms of what's going to make the bigger impact in the future?

It's going to be DNS over a web portal, over a web socket. And so, we want—I mentioned that DNS over HTTPS, that solves middle boxes, that solves China, and actually I've seen actually active code, long talks with folks about that. There are people looking at DNS over Quick of course, which is the Google deployment of web, basically web traffic.

And, we're starting to work on the resolver to authoritative portion on the DNS over TLS, but I think most people are scared that we're going to have to approach ICANN and talk about the root servers, and that's one question, but there's so many more authoritative servers. We have to talk to—how do we talk to .com? How do we talk to any of the TLD's, right?

And, I do believe that they want to see, and any TLD operator I think would come back to us and say, "I want to see deployment, operational numbers, I want to see how this thing performs. We just can't turn this on if it's going to destroy what we have, so that's why I feel having good operational numbers will actually help deploying the DNS over TLS. Though, if you're going to pick a bet on which horse, I feel it's going to be DNS over HTTPS. I do feel that the privacy thing is, that's not fully understood yet, and can we get a better integrated in decline operating systems?

Can mobiles and laptops support the DNS over TLS by default? Can we build on, I think what really works, and I think we all know this; is stuff that just turns on and works; it's already there, we don't think about it. I get laptop software from my corporate IT people, they deploy a bunch of stuff, I just run it, right? I don't have to think about it, it's just -- and so, I think there's lot of people, that's what they get. And, that's—it has to, I feel that's going to be the driver of this sort of stuff. But, until then, it's definitely going to be piecemeal, or it's going to be slow going.

And, you have limited scope in this. I think you can help on the TLD and gTLD side, but it's hard; you need bigger carrots, right? You can't just force people to do stuff in this world. And, we have…how do we figure out the best way to convince people this is the right thing to do? And I think one way is to show deployment at scale and try to avoid these traps. But, I just look at IPv6 deployment; I mean, our company, I work for a large enterprise company and we're still not deploying IPv6 because it is cheaper for us to just go onto the market and buy large quantities of IPv4 blocks because we've got a giant checkbook, right?

People don't like to hear that, and we're bidding against AWS, bidding against those kinds of players. IPv6 is probably another 20 or 30 years in the process, so these are slow things moving.

It's frustrating for me, because I think we should move faster, but something has got to help drive that, and I think mobile clients are going to drive the bulk of that. If we can get the big mobile platforms deploying DNS over TLS, by default out of the box, that's a large user base; that's a pretty massive user base actually. So, that's where we are in that, and I hope I sort of covered all the bits. If there's any sort of questions from folks, if I miss stuff…

ADIEL AKPLOGAN: Thank you very much, Tim. Very interesting. Oh, and I actually like the—

TIM WICINSKI: Question? Okay.

ADIEL AKPLOGAN: Hi Cherine. Go ahead.

CHERINE CHALABY: Thank you very much. Could you go back a couple of slides? Keep going…keep going, yeah, this one. So, at the very bottom you say, "a root server is only part of the solution." Could you elaborate again a little bit more on that statement?

TIM WICINSKI: Sure. In looking at resolvers, talking to authoritative servers, of course the top of the chain is the root servers, right? But they talk to the TLD and the TLD's talk to the second-level servers, so we can work on the resolver to authoritative, we can sort of interact with various people that run authoritative servers.

I feel that we would want to basically work up the chain by first starting with folks that run domain servers, and then talk to the TLD's, because I think as you collect more and more operational data to show; this works at scale, this isn't' going to break your infrastructure, that's going to give the TLD folks more reassurance, "Oh, I can deploy this and it's not going to break."

And then, I think the root operators, I haven't talk to many of them, but I think they would be in the same boat; they'd want to see, "How would this work if we turned this on? What sort of impact would be on our infrastructure? How would we support it, that sort of thing?"

So, I sort of feel that it's basically, we're going to have to work up from the domain owners to the TLD's to the root server folks, and you could give good guidance, much like with the gTLD process and stuff, but you can't probably make everybody do certain things, right? I know with the gTLDs you're able to get

them to say, “Oh, we can do DNSSEC,” so maybe for the next round we can put something in there about DNS over TLS, or something of that nature, but that doesn’t affect all the ccTLDs, that doesn’t affect some of the legacy stuff. So, I think the way you do that is you show that it’s not going to break their infrastructure.

CHERINE CHALABY: So, just a quick follow up, and when you say, “we,” who?

TIM WICINSKI: I mean ITF; I mean when we’re sort of building these tools and folks like me would go and we’d talk to folks and say, “This is a good thing to run,” and I go and I talk to Matt Larson and I say, “Matt, I think ICANN needs to do something about this,” and wisely he would probably say, “Let’s make sure we don’t break anything,” that would be the first thing he would say probably.

And when I have talked with folks, and we’ve talked with some of the vendors, the DNS vendors, these are third-party folks, some of their employees are here in the round table here, they wisely say, “We would like to do this, we’re not against it, we just want to make sure that we don’t break our own infrastructure to our existing customers by deploying something like this.” And, it’s a very valid question, and yeah, I think the root servers have

always been setting a good example, but it doesn't always mean that people always listen, right?

ADIEL AKPLOGAN: Thank you, any other questions? Just so that the audience, I mean, this is between the Board and TEG, but if you have questions, feel free to go to the ICANN and ask as well, because it's open. Question?

DANIEL DARDAILLER: Daniel Dardailler from W3C. You mentioned DNS over HTTPS, and that Raminds me of one of the comments that Steve Crocker made earlier on the former director's lease, which is that we have an issue to actually update the zillions of DNS software on the client's side, and so for various reasons, but one of the features of DNS over HTTP if it's integrated in a web page is that the update is automatic; whenever you reload the webpage, you get new code that is going to do the DNS business, so is that something that people have been looking at as well as being more open, the HTTP port being open through firewall, there's sort of another feature of easy to update web code.

TIM WICINSKI: I know they're approaching it on; can we get the protocol worked out on talking DNS over HTTP, most HTTPS. Then, I think they want to go; where are we going to go from there, sort of thing. The point was, they're trying not to get ahead of themselves by doing everything, they're; let's get the protocol, let's get that sorted out so it's fairly standard, it's defined, it's published, people start using it then I think you're going to see that application side where things like that happen.

ADIEL AKPLOGAN: Thank you very much, and one comment I may have at the end; you compare this to the slow uptake of IPv6, and the fact that the market is allowing people to go and get IPv4, but the question may be, because this touches on privacy and touches on people's use of the internet and the protection of part of their privacy, maybe that can be a drive that we can see differently from IPv6, so this is more at the infrastructural level.

JONNE SOININEN: Yeah, Jonne Soininen, I'm the ITF Liaison to the ICANN Board. I've been a little bit following the IPv6 space over the years, and you're right, there is here some kind of possible driver that IPv6 didn't have; it didn't have a privacy—well, you could say that it

has some support for privacy, more than IPv4, but it's still not the killer, so to say, killer application for that.

This has more possibilities in that sense, however, I can see that these both have the same inherent problem, which is that there's a long tail of different software, long tail of different players, and that there needs to be a certain kind of, what I would call, industry or market coordination; you have different vendors on the laptop or handset on the host side, and then on the resolvers side, and then you have something else in the kind of like when you go up on the root, and this has to be coordinated that this is to actually happen.

But, in IPv6, it would be wrong if I wouldn't correct; I think actually the uptake has been quite good, but it's very, very uneven. In some places, it might be extremely good. In some places it might be just absolutely not there, and that makes the reason the clear coordination or kind of alignment over the whole space, it kind of like makes the really moving to that really difficult, and it actually drives people who are doing end-user services, that the IPv6 is the lowest common denominator. I'm not sure that here you couldn't—here you can use the old and the new more in parallel than in IPv4, so this doesn't necessarily have all the problems that V6 has.

ADIEL AKPLOGAN: Good. Thank you very much, and a space to watch. So, we will move to the second topic on the agenda, which is about the DNS capture and analysis. This will be a presentation in two parts, first part will be by Mauricio, who will tell us a little bit about DNS Stats; a tool that is used and maintained by ICANN. And the second part will be by Matt Larson who will also talk about some of the actual measuRament and analysis that is happening within ICANN org. and Engineering Team.

MAURICIO VERGARA: Thank you very much, my name is Mauricio Vergara. I work for the ICANN DNS Engineering Team, which is in charge of all the operations on the ICANN manage root server, and also the ICANN domain name portfolio. I'm going to tell you about a ltitel bit of history of what we have been doing in the past 4, maybe 5 years, and what has led us to create this whole suite of software called DNS Stats.

As many of you might know, there's a lot of operators in the DNS world that is using right now a software called DSC, which is used to present what is the traffic reflected on DNS on their servers. At the moment, when we wanted to start using DSC, we were using DSC, we noticed that there was some issues in terms of how we were presenting the information, so we decided that it was needed to create a new way to present this information,

so alongside with the help of Sinodun, we created a tool called hedgehog which is different to DSC, it was already using a database as a back end, and it was dealing with problems for instance for what happens when you have more than 100 servers, and it's a little bit difficult to plot.

One of the main concerns that we had at the moment was that we didn't want to do a fork of service that wasn't used by anyone else, so we decided to go on the open source route, so everybody else can use that.

Currently, the software is still in use, and is running on version 2.4. You can see life on the side, stats.dns.icann.org all the services that DNS Engineering is currently serving. As you can see over here, this is a small screen shot of what the presenter of Hedgehog is doing right now, which it can look like traffic that is getting the root server managed by ICANN on any given day, divided on five different regions. This is pretty normal for us to use, and it looks like the cardiogram of what is looking on the traffic that we are seeing.

So, the second problem that we noticed when we were using this was that we needed to change the way that we were gathering the files inside of the DNS servers. Because of that, there was created this second portion on the servers itself, which was to create a collector that it will replace all the DSC

collectors that was being used at the moment. In the beginning, we noticed that the DSC connector didn't have some information that we found useful like DCP receipts, or ICMP messages.

It's my understanding that the DSC maintainer currently is working on implementing these features, but 4 or 5 years ago when we were deciding to adopt those, it was a requiRament for us. Because of this, we noticed they need to have a new format to capture all the information that we were having on that server.

That is why we started to develop a new format, which is called the compact DNS, or also known as C-DNS, which is an efficient file format to help us transfer and see what is on the DNS Traffic. All the specific agents on this format has been taught into ITF in the forms of a draft, which has had a very positive reaction, and it's currently on the version 06. We would like to have that to the next stages in the near future.

This tool is also an open-source tool, and it's under Mozilla Public License, open-source license. The DNS Stats Compactor, it's comprised of basically two programs, one of which is called Key Compactor, which is very similar to what is currently used on several applications for networking called TCP Temp. The main function of this compactor is to reap the traffic from one or

more interface, or even another bigger file, another TCP file created and generate the C-DNS, the format that we have been working for.

And everything that is not contained into that C-DNS format, we can still store it in the form of pickups, so it can be analyzed for future revision. This compactor has on the other hand another tool that is called the inspector, which can be used to do the opposite instead of going from pickup to C-DNS, we can reconstruct the traffic generated on the C-DNS to generate a pickup file which is kind of the language that most of the research or analysis people is doing nowadays.

In the past few weeks, there has been a new version of this compactor, which has added solomnization and the output of Inspector. This has helped us to be prepared for all the incoming changes that have been discussed on the mailing list of DNS operations and also some RSAC about anonymization of data. We are trying to push this new version of CDS to also be deployed more widely and adopted by other researchers. Just as a reminder; the compact DNS currently uses about 30% of what is the regular pickup file, which could be really, really useful in terms of bandwidth when you are retransmitting this to the researcher's side.

In the future, we would like that the presenter is having a new version, that is going to be the version 3 which code name will be Wombat, and we will replace PostgreSQL as a back end with a cluster of ClickHouse and ZooKeeper. Instead of doing the particular graphs that we were showing before, we will start using Grafana to plot the data, and you will have also the ability to create your own plots, or even export the pickup files on request.

Everything on the presenter's side is currently in development. I'm going to show you later one screenshot of that. And, on the compactor's side, we are trying to get more detailed captures, so we can help all the researchers to do better analysis through time, and we are already in talks with the people from DNS org., so we can start doing that on a regular basis, and only using the C-DNS format, so if they want, they can create new back end tools to operate directly from C-DNS, or they can transform that into the pickup files and use what tools have been used until now.

Also, as I was telling you before, the anonymization schema has been working with several parties, and is currently released about a week ago, so we are really happy to start seeing this working in the near future.

Finally, I would like to show you how the new presenter, how it will look in the future, this is one of the internal test developments that we are using right now, and how we want to go in the direction of this. One of the principal reasons that we are using this new format is to be able to propagate a lot of information without losing as much as possible, and we have been using this to frame all the OCTO researchers, so they can do analysis, for instance the KSK rollover and stuff like that.

So, we invite to everybody to start using these tools, and if you have any questions I will be more than happy to answer here or in the rest of the week.

ADIEL AKPLOGAN: Thank you, Mauricio, but we will move with Matt's presentation, and we will take questions just after. Matt?

MATT LARSON: Thank you. So, Jay had suggested the theme of DNS Capture and Analysis, and so here is another presentation on that theme. It's very short, I just wanted to go over what the office of the CTO Research Team does in this area at a very high level.

We have access to several different kinds of DNS related data. At the moment we have traffic from root servers; B, D, F, and L. B,

D, and F are in PCAT format, although we're phasing out access to that data in 2018 and we're going to be focusing exclusively on root data which as Mauricio said is in the CDS format.

Actually, thanks to Roy Aarons on the team, we post-process root sever data into to a very clever plain text format, and it turns out if you're willing to give up just a little bit of the detail, you can wind up with a file full of, or a directory full of files with names of some of the parameters in the queries in them, and other parameters in the files themselves, and you can use the familiar and beloved UNIX text processing tools like GREP, SED, AWK and Friends, and do an awful lot of analysis very easily and quickly. So, SQL is for losers.

The rzkeychange plugin that I mentioned if you were here for my presentation about the KSK Roll, that's a plugin for DNS CAP that allows us to get statistics now from 12 root servers. These are high level statistics about counts of packets and queries processed as well as the RFC8145 trust anchor reports. We also have a resolver test lab and the relationship to DNS capture is that that lets us capture DNS traffic in a controlled environment. We've used that to research different aspects of the root KSK Roll for example to see how resolvers behave related to DNSSEC and the root key, and we can capture that traffic. And then, not directly related to DNS capture data, but we do have historic

root and TLD zone files which are very handy to have lying around.

And here's just a very short high-level list of the sort of projects that we use DNS data for; Roy did an analysis of Corp.home mail and compared statistics to the inter-aisle report from a few years ago. If you've been here in Puerto Rico for any length of time, you've probably been unable to avoid my talking about the KSK Roll, and the role that RFC8145 data has played in that.

You also perhaps heard Alain Durand and Christian Huitema talk about the ITHI project; the Identifier Technology Health Initiative, and the metrics that they're developing, some of them are based on DNS traffic, also using a DNS CAP plugin actually to do the particular measurement that they are doing. In the past, it was about a year ago, we've done joint resolver behavior research with APNIC, merging what we see in root server data with what they see from their Google-based ad measurement. We have statistics on DNS deployment that WHOIS tracks, different usage of operational parameters, and then as I mentioned a moment ago, we've also studied the behavior of validating resolvers in the context of our test lab. So, that's all I had, I just wanted to give a brief overview and description of what we're doing.

ADIEL AKPLOGAN: Great, thank you Matt, thank you Mauricio again. Questions on DNS Measurement Statistic and Tools that are in use today? Really, no questions? Oh, there is one.

MARTIN SUTTON: Well, we know there's a lot of interest in also how we can better secure the internet, make it more reliable, do people find their way to these statistics? Do people talk of OCTO also outside parties, can ICANN use that to improve data, how does that work? I know it's not a purely technical question, but I hope you allow it.

MATT LARSON: I'm sorry, I need to hear it again please, I guess I maybe didn't understand the question?

MARTIN SUTTON: Basically, you're on top of a lot of movement, and you measure that, and you can generate tons of data that could be used for analysis and discovering vulnerabilities and things you can do something about. So, I can see that outsiders may be interested in that data to do something with, or is it only for your own use?

MATT LARSON: Oh no, it's definitely not only for our own use, although I will confess that there is some of these efforts that we could do a better job on publicity and publishing the data, although on the other hand, there's some things, as I said, you've probably been able to avoid my talking about some of this root KSK Roll data, so no, our intent is definitely to make the results of our research available to the community to help however we can.

MARTIN SUTTON: But it's not happening at the moment yet?

MATT LARSON: I guess it depends—some things yes, and we have things in the works as well that aren't yet published, but we definitely have published some things for sure, and as I say, we've done a lot of research on the RFC 8145 data that I've talked about here extensively.

MAURICIO VERGARA: Just to clarify, are we talking about the raw data or the outcome of the research and analysis done?

MARTIN SUTTON: My deeper intent is; I try to get a feel for the value of what we're generating here. And, as I said, I'm not technically skilled but I can see we are on an enormous wealth of data and it would be great if that would be used to its best use.

WARREN KUMARI: Warren Kumari, Google. So yeah, the RFC8145 data that Matt's group publishes has actually been really useful; that's something that a lot of people have referred to, it's data that isn't really available anywhere else, and so I think that data in particular has shown a lot of things which people weren't really expecting, and is providing very valuable input to a lot of other decisions, so that data, I can't speak to the others, I'm not sure as much.

ADIEL AKPLOGAN: Okay, so let's move to the next presentation, Jay?

JAY DALEY: Thank you. My name is Jay Daley, and I'm going to talk about Domain Classification. This will have virtually no technology in it, you'll be glad to know.

So, classification in a nutshell; you start with a standard industry classification, there are a number of those, most of them

internationally are aligned, such as NACE in Europe, or ISIC, the international one, but the US one is a generation behind and quite different.

You then classify domain names by their website content, which I'll explain a bit more; there are two methods to doing that, manually, where a human visits a site and classifies it, and a well-trained human, well, a cleric with a few days experience even, can achieve between 500 and 1000 a day to classify domains. And the other method is machine learning. This is where you have a WebCrawler that grabs the text from the site and uses a trained neuro network to classify it. The output is a single primary industry classification for that domain name based on the website. Sometimes people have multiple secondary classifications as well, particularly if you're using a neuro network and it's giving you some idea of probability.

So, just to give you an idea of what I mean by "classification," this is an extract from ISIC, ISIC level A is Agriculture, Forestry and Fishing, and underneath that 01 as you see is crop and animal production, hunting and related service activities. I bet nobody's ever said that at ICANN before. And then, further on down, growing of non-perennial crops, and then down to growing cereals except rice, leguminous crops and oil seeds, so

this is what an industry classification is like, and it has got a bit of depth to it.

Now the benefits of classifying your domains are; firstly, that national statistics bodies use classifications to size national industries, so the try to measure the number of countries and organizations in a particular classification, the number of employees, and the turnover, or the total turnover in industry.

So, you can now begin to establish the economic value of the domain name industry within a country. You can look at the market penetration by companies in turnovers and the value of the industry served by domain names. And then, at the registrar level as well, you can help a registrar understand if they are specializing in particular verticals in particular industries, and the registrar can direct their sales and advertising.

So, classification is something that's just emerged in the last couple of years, maybe three years at the most, and is now being used by some of the more advanced companies and registries around.

So, there were three world leaders that I'm aware of, there may be more, .nz which is attempting to classify every domain name in their register using a mix of manual and machine learning. CENTR which has a working group of EU's, largest registries and

registrars, which has created a new specific classification standard for domain names, which is an important piece of work. And then, a commercial service that quite often comes to ICANN meetings data provider that classifies domains among many other data points.

I'm just going to give you a bit of data of each of those; so, .nz has 700,000 domain names in their registry, they have a large cluster of a system called Hadoop where they customize distributed WebCrawler so that it can go out and get the pages, all 700,000 domain names. They brought in a set of students and then they classified 100,000 plus domain names and then used tested multiple machine learning models to find one to classify the rest of it.

The accuracy that they get varies by the antic code because of the number of domain names that might fit into each code varies, and so there might not be enough data to classify it properly. .nz are rolling out a commercial product where they combine the classification with traffic measurement, and then a registrant can then compare their traffic against other sites in the same industry, and they can understand if their investment on say marketing or website improvements delivers a relative gain in traffic compared to others in the same industry, which nobody else can tell you.

Now, CENTR, we have Andreas over here just sitting in the front row who's one of the people this is the brain child of. They have a registry/registrar data group, which is this group of your largest European ccTLDs and registrars, and they're working on producing frameworks, classification and tools, so they have quite a big ambition. And, the intent is to help the industry better understand the market of domain names, and it's supported by CENTR.

So, one of the most important pieces of output is the domain industry taxonomy. So, this is a classification of industries and sub-industries matched up to the European NACE codes. Now, the reason they've had to redo this is there are lots of new types of businesses that particularly appear on the internet, such as different types of auction sites for example, that aren't necessarily property captured it in the standard industry codes, and they have a website about this, and this again allows relative market penetrations in a country and cross-country comparisons which is very useful for CENTR members.

And the final one to talk about is; data provider. Now data provider is a commercial service, and so they do large-scale global data calling. They try to grab 30-50 pages per website which is much deeper than most and collect 150 data attributes.

And then, which includes industry classification and they build a trust score.

They also attempt to understand the language, and the country, so you understand the country using a different set of identifiers, such as the ccTLD or the language and the like. And the interesting point is that their identification of the country of the website based on the content, is generally different from and more accurate than that identified from the WHOIS, which is subject to very variable registrant-based data.

And as you can see, their type of customers are, both within the industry, and without and the use cases by companies to date include; understanding particular types of companies, understanding the digital footprint of a website, and various forms of market intelligence.

So, that's it for me. Any questions?

ADIEL AKPLOGAN: Thank you very much, Jay. Questions for Jay?

DANIEL DARDAILLER: Daniel Dardailler, W3C. In relation to the categories that exist, already the standards that are applicable and I was wondering if it's on a finer level, on a granularity level to define the difference

between ICANN and IETF and W3C for instance, or is there just in this category like a tech organization, you see what I mean? What is the quality of the level for the technical community kind of business, just as an example.

JAY DALEY: So, in the example of our industry so to speak, no, there would not be the resolution required. In more established industries, there is the depth of resolution. One of the problems though is classifying to that full depth is quite difficult, and sometimes people pull back and only classify say at 2 or 3 levels instead of at the 4th level to do that with. And, it's some of those sort of sub-classifications, that fourth level, that people are working on for the next version of these standards, so that we can then get a better resolution there.

ADIEL AKPLOGAN: Any other questions? No. Okay, so we will get into the AOB. We have a few requests for AOB, currently three, one from Kaveh on the linkage between the BTC open-session and the TEG, one from Ram, a feedback question, and I have added one as well on a future topic for this meeting. There has been some discussion about 5G on the mailing list and look at how we can frame that as well. So, we'll start with Kaveh, if you can take us.

KAVEH RANJBAR: So, let me provide a bit of background information first. The TEG has a sort of very informally, after some history we have developed basically to have two meetings per year, so normally on the Policy meeting, which is the meeting in the middle of the year. The ICANN meeting, there is no TEG meeting, it's the first meeting and the last, the AGM meeting of the year. So, the next TEG meeting would be in Barcelona as it has been established.

Also, in Abu Dhabi, the board started a new committee, it's called the Board Technical Committee. The Board Technical Committee has three distinct responsibilities, if you look at the charter, to summarize them; one is to look into internal IT practices in [inaudible] and ICANN so the IT projects, and this is mostly fiduciary responsibilities. The second part is to look into the technical request or the technical interaction with the constituencies. Mostly I assume the interactions with RSSAC and SSAC, but if there are other technical discussions between board and other constituencies, BTC will be the one who will channel those discussions.

And finally, there is the work with OCTO and looking into future or related developments and research in the DNS area, which is actually very much aligned with what TECH does. We also tried when forming the BTC, we also tried to keep it as open as

possible. The first item that I mentioned, the fiduciary part, because it's mostly report of internal ICANN projects and IT related stuff, sometimes security related to the organization. At the moment, we decided maybe it's easier to keep it closed, but we are every time evaluating the meeting, so the aim is if possible to open it up.

The second part, we discussed that, the plan is to also open that up to second port, which is basically the technical interaction with the constituencies, most of the information is already open. We are also, we are aiming to open the deliberations.

The third part, which is looking at the future of DNS and the research and basically evolving technologies has been always open from the beginning, which is every meeting, so there are only three meetings. And when we try to come up with the agenda, it was one meeting before Abu Dhabi which we just did a test run, that we had a meeting, and there was one in LA during the LA workshop.

Trying to come up with agenda, we actually thought there is a lot of overlap between what we get as agenda items and what's being discussed in TEG. So, the idea was, or the proposal at the PTC open meeting on Saturday if I'm not wrong, earlier this week anyways, was to possibly combine these two meetings, maybe having a bit more time, but we will have the open part of

PTC meeting combined with the TEG meeting. So basically, we will have one larger meeting which will be attended by the board technical committee members, plus any other party interested in technical issues in ICANN.

Just a proposal coming from the PTC, there was support there in that meeting. If that's the case, then we will need to work out the mechanics of that, for example, how to call for agenda items and the rest of things. But first is; does everybody agree, and is there any objection that we explore the possibilities for that, how to formalize that? And then we will need to look into procedures on how to move forward. So, is there an objection, any comments, please?

UNKNOWN SPEAKER: No objection, I think that is a marvelous idea. I think it would be very useful for the Board to set their expectations at how technical or not they want the presentations, and what they want to get from it. And then the community can the respond with appropriate ideas at that level.

KAVEH RANJBAR: Point taken, thank you. Any other comments about that? Okay, Adiel, please?

ADIEL AKPLOGAN: Just to add as well that we discussed that in case it is that, the public parts may move to the week of the meeting so that everybody is able to attend.

KAVEH RANJBAR: Oh, definitely, I think if that happens, because Adiel is right; normally we have the BTC meetings during the workshop part of the meeting, which is earlier than that, or before the actual ICANN meeting, we will definitely move it to somewhere during the week, for example, in these large meetings. So, we will do that. Okay, hearing no objections, what I will do—

ADIEL AKPLOGAN: There is a question here.

UNKNOWN SPEAKER: Just a quick question; you said the BTC is meeting at every ICANN face to face, but the TEG is not, the TEG is only meeting, so you'll have a meeting with all of us, no matter what, which means that if we share agenda, more action items in this new group we will have one meeting off.

KAVEH RANJBAR: No, so basically, I should have explained a bit more. The BT as I mentioned, BTC has three main purposes, so we plan to basically in the first ICANN meeting of the year have all the three sections, which the last section is, which is shared with TEG, the second one we will only have the first two, which are right now closed. We are aiming towards opening them up. As soon as that happens we will discuss that and how to involve them. And the TEG meeting, the last meeting of the year, we will again have three sections, the last section will be shared with TEG. So, I don't think that will cause any issue.

Okay, so based on this, what I will do, I will work with the OCTO office and communications with the TEG, we will receive emails, we will try to come up with a process here, just to be practical I will try to do it within BTC. We will circulate a proposal on how to move forward, basically on procedures, how to collect agenda and all of that. And we will try to get rough consensus within TEG. If everybody agrees, then we will make that operational and we have enough time until Barcelona. Thank you, everyone.

ADIEL AKPLOGAN: Thank you, Kaveh. Second point on the AOB, Ram, can you?

RAM MOHAN: Thank you, this is Ram Mohan. I wanted to take a little bit of time here and place the Board members who are here at this session on the spot and ask for the Board members to provide some feedback as to the value of this kind of a session, and what they'd like to see more, better, different from the TEG.

I think it's an opportunity not just to only provide feedback, but to have an interactive dialog right here, we have the time, so that's what I'd like to suggest. And what I'd like to do if you guys are okay with it, including Board members sitting not at the table, is that I'm going to poll you. So, I'll start with you Becky, and ask what feedback you have?

BECKY BURR: No, me. I was paying attention. And, I'm a little embarrassed to admit that notwithstanding the fact that a lot of this was about privacy, significantly, DNS privacy standards discussion was a little over my head. I am interested in it, and I'm not afraid of technology, but I'm not an engineer, so that's my one observation. I would like to know a lot more about it.

RAM MOHAN: I see two comments, one from there and then one from here.

TIM WACINSKI: Tim Wacinski, so as the new person, and actually being my first TEG and given that assignment. I was cautious about how I was supposed to present the data, present the topic and stuff. And so, I tried not to go too much, but I realize also, because it's a techno audience, sometimes I went towards the techno side, so I will be glad to sit down and talk to you about anything you want in any sort of—so that is my failing for not quite understanding the room as well, so a learning experience. Thank you.

WARREN KUMARI: So, Warren Kumari, Google. About 4 or maybe 5 TEG meetings ago, I can go through and have a look and find the actual link, I think I presented on a much more introductory thing for DNS privacy. I believe that they are all recorded, so I can see if I can find the thing and sort of a more friendly introduction for her.

RAM MOHAN: Thanks, anything else, Becky?

BECKY BURR: No, I find these meetings very interesting, I learn something. I didn't meant to suggest that you should dumb it down, because I have the impression that the purpose is actually technical, but I do get something valuable out of it every time.

RAM MOHAN: Thanks, Matthew?

MATTHEW SHEARS: This is my first time attending a TEG session, so I very much appreciate what you presented. I very much echo what Becky said, but just to add, I guess what's missing for me is the implication; what the impact is on ICANN's operations, and maybe it doesn't relate specifically, but is there a way of drawing this out for those of us who are in the audience who come, put the pieces together, that would be incredibly helpful, thanks.

RAM MOHAN: Cherine, to respond to that, right? No, you're just in the queue. I'm going to come to everyone, but does anybody from the TEG want to respond to Matthew's comments about requesting for implications, don't just talk about the tech but talk a little bit about implications of it, Jay?

JAY DALEY: Yeah, I don't think we've ever had a set of guidelines for speakers that set out what they're attempting to achieve.

RAM MOHAN: Okay, sounds like that's a useful thing to do. Jonne?

JONNE SOININEN: Yeah, as a response to that, I think that Jay first of all is right; we have actually never thought about it this way, that we would give guidelines on what we do, and I think this is a good start for that, but maybe we should go beyond this round and actually provide some guidelines from this, and also for ourselves as kind of a 'what we actually expect from this.

RAM MOHAN: Thank you, Jonne. Let me switch to the other side of the room and come to you, Cherine for feedback?

CHERINE CHALABY: So, I look forward to these meetings because I always learn something new. There's something about the format that is not working for me. If you look at this, there's a 60-minute presentation and only 15 minutes of dialog, and to me that's a pity because if we are here face to face, we should have a dialog. Presentation could be sent by emails, so if you look at the slide, right?

DNS Privacy, 20-minute presentation, only 5 minutes, very little time for any dialog. DNS Capture; 20-minute presentation, only

5 minutes for dialog. And Domain Name, the same thing. So, it's all presentation, presentation, presentation, and very little interaction, interaction, interaction. And I'd like to reverse that if possible, it's a suggestion, that's all.

RAM MOHAN: Thank you. Kaveh? So, this is in response to Cherine's request?

KAVEH RANJBAR: Yes, so I agree to basically all of the comments. I think the comments are fair. What I suggest to do is I will take it up as the BCT chair, again, from the Board perspective I will compose a proposal, send to the TEG, try to see if there are objections or there are proposal comments, and then…

RAM MOHAN: Thanks, Kaveh. Anybody from the TEG who wants to respond to Cherine's comment? Tim?

TIM WACINSKI: So, I agree with Cherine. I love interaction. When I give my presentations for my day job, they're very back and forth, right in the middle of stuff, so I do enjoy that as well and I agree with you completely on that.

RAM MOHAN: Thank you, Tim. Patrick?

PATRICK FALTSRAM MOHAN: Thank you very much, Patrick Fältström from SSAC. I think the proposal from you Cherine is really, really good. That it should be 5-minute presentation, 20-minute dialog, so we should turn it around. Secondly, I also support what Kaveh suggested, and I think this whole TEG, TLD sort of thing should be driven by the, now, given the structure, I think the Board Technical Committee, and the Technical Liaison Group. The rest of us should be used as resources that can help in the dialog, thank you.

RAM MOHAN: Thank you, Warren?

WARREN KUMARI: Warren, I think that would be a great idea, but I guess I should also mention that in previous TEG meetings we've had a number of presentations and then no dialog at all. Not because of time I guess, just because towards the end of a meeting everybody is tired, and/or maybe the presentations weren't very well formed.

So, we'll do a ten-minute presentation, so then possibly working more on making sure that the presentations are actually going to be of interest to the board and are going to be providing useful information and not a little soapbox for us to stand on and talk.

RAM MOHAN: Thanks, Warren.

CHERINE CHALABY: Quick response; I think the presentations should be, and this is a suggestion, right? In the form of framing an issue, and then inviting people to discuss and debate the issue, that would be my suggestion.

RAM MOHAN: Thanks, Cherine. George?

GEORGE SADOWSKY: Thank you. Well, this is an exercise in awareness for me. I think it's very important, but it's also an exercise in humility when it's driven home to me that I don't know what I don't know. And, in that sense, it's very valuable because when we discuss issues like this and we don't have the value of such expertise

surrounding us, at least I sometimes fall into this trap of thinking I understand the issue and therefore I have my solution and it's the right solution, and understanding the depth and the variety of the kinds of things that are being talked about by this group make me understand better what I can contribute and what I can't in addressing specific issues.

RAM MOHAN: Thanks, George. Any response from TEG members towards George's comments? Alright, then we come to you, Martin?

MARTIN SUTTON: Yes, thank you. I've been to some of them and I must say, you always take something away. And I'm 100% with the suggestions that are made. For me, it would help if the purpose of the—explained their—comes forward. Why would we talk about it, it says Board and Tech, it doesn't say BTC and Tech, right?

So, if you can think about what it strategically means for us, why we should be informed, that would help me to listen better. I do think it's good to give a little bit of the tech info and go a little bit deep there to tickle me also, just over the Board where I don't get it anymore.

But indeed, the other element is; so, what does it mean, what can .org do, how can we make a difference? It would be great if there is a question like, a request, or a suggestion of action as well that would be discussable and then last but not least, I think we are more than three minutes in the presentation, but still keep at least half o the time for dialog. If it doesn't happen, okay, let's get to drinks earlier.

RAM MOHAN: Thanks, Martin. Sarah, and Jay?

SARAH DEUTSCH: So, as a newcomer, I would just advise that this be drafted as if everyone at this table from the board knows nothing, because at least there's some people, we know tech, but only as explained to us in commoner's English, so trying to keep things basic, and also not only understanding the implications but maybe boiling that down further, like what keeps you up at night, what's most worrying about some of these topics so we know what we should be concerned about.

RAM MOHAN: Thanks, Sarah. Jay? And is there anyone else from TEG who wants to respond? Okay, Jay.

JAY DALEY: In addition, I think to the types of presentations that have an action at the end of them or a risk or something like that, I think general environmental technical ones of other things happening in our industry are important. That's why I brought the web classification one here, there's nothing for you to do, but I'd be very surprised if many of you know anything that is taking place, and it is a big and growing part of our industry, and we'll have implications for it.

And there are still things for all of us to understand in those implications, but it's good that you're an action-oriented board, but I think it doesn't always have to be the response the presentations we give.

RAM MOHAN: Thank you, Tim?

TIM WICINSKI: I think Sarah's made—I think all the Board members have made some great comments. I just feel the next time I get tapped to do one, I'll actually do a much better presentation or you all, that's all I can say. So, thanks.

RAM MOHAN: Thanks, Tim. Lousewies?

LOUSEWIES VAN DER LAAN: Thank you. I think actually most of the questions have been put. What I was still thinking is, I was wondering about the convergence of views within the tech. My experience, I've been only to one ITF, it's that you put five engineers in a room, they're going to have twenty different opinions, that's how I was reassured that ITF couldn't be captured.

So, I was wondering here if everybody here is quite harmonious in their views, are we getting the kind of diversity of views or competing analysis of technical situations, or do you guys all just agree, or are you so brilliant that everybody works together, so I'd like to get a feeling for whether the composition is reflective of the full scope of technical expertise there is on the planet.

RAM MOHAN: I'm positive we'll have diverse views on that question. Warren, who else?

WARREN KUMARI: So, I think that for these particular things, what we're discussing is generally, everybody is talking about a different specifics sort

of things, and also, we're generally speaking at a high enough level, or not down in the details enough that we get along perfectly well. With that said, I think it would be useful if there were different sets of views brought along.

But also, the things that we generally talk about here are not that controversial, it's usually once it's been said that everybody agrees that this is not a bad idea, then it gets spoken about here. So, it would be more fun though to have sort of a cage match.

RAM MOHAN: Anyone else from the TEG on this? Oh, Tim?

TIM WICINSKI: Yes, I tried to speak from my operational sort of point of view in sort of deploying large-scale infrastructure, and especially about the privacy bits. But also, from the ITF point of view, when I pointed out that DNS over HTTPS is probably going to be the bigger winner, that sort of hurts as well because I think the other ideas are more technically elegant. But, it's like, I look at it and say, "That's going to actually probably end up winning in the long run." I've tried to basically balance both of those, and also, I was really thinking operationally how we deploy this server stuff, right?

Who cares, why do we care, that kind of stuff. And from the enterprise world, we're just a different beast then a lot of folks in this room, right? We run big infrastructure, but we're essentially a sales company, right? And that's what people think of us as, and we fly under the radar and we build very large infrastructures, solve these very complicated problems, and the security stuff is very big for us, so we look at this a lot, so I tried to really sort of balance all those, but I didn't go into explain some of that as much as I should have, but yes, I definitely like that sort of point of view too.

RAM MOHAN: Thanks, Jonne?

JONNE SOININEN: Yeah, one thing that I can say about the people that are selected for instance from the ITF to come and represent here, they are selected in the way that they can give you a balanced view, and I think that for instance Tim tried to a little bit, well, it isn't quite a cage fight, but like you said, you tried to show that they are different solutions to the same problem, and some of them seem more likely than others. So, it's also the people that are coming here are also people that can represent these in a very balanced way.

RAM MOHAN: Thanks. Jonne, did you want to chime in on feedback?

JONNE SOININEN: Yeah, so first of all, the presentations here always have very good quality and very good, there's no question about that. It seems clearly that we, as a Board, haven't maybe payed enough attention to give actually guidance on what we want to see, and that is a little bit kind of like something that I find that is good that is coming out of this discussion, but what I see that this is not as actual feedback to the presenters, the presentations were very good.

This is feedback to the board that we actually can give, that we can actually give the right expectation, and as Jay said, we've never had a style sheet or a guide to go on what is the right level of doing that? And it's clearly an area I would see this as an action to the Board to come back and be more proactive on solicitation these and asking what is the level that we want.

RAM MOHAN: Any response on that? Okay, let me switch—oh, Cherine.

CHERINE CHALABY: Hang on a second; point of order, you went around and selected all the Board people and asked for their feedback, you're not going to get out of providing yours, are you?

RAM MOHAN: No, I'm making sure everybody else gets their words in. Cherine?

CHERINE CHALABY: I just wanted to make sure that no one took my comment as saying the presentations were not of good quality, because they were very good presentations. I am just feeding back what we hear in also what we hear in also our meetings with other constituency and people are saying these face to face meetings don't often happen, so there are maybe two or three times a year.

So, let's use them for a real dialog rather than a presentation, so it's just feedback that maybe a change of direction in those meetings would help us all to dialog more, that was the point I was trying to make. So, please, make sure there is no criticism of the presentations who were excellent quality, thank you.

RAM MOHAN: Thank you, Avri?

AVRI DORIA: I wasn't sure I had any feedback. I think some of them went by a little fast, and there was more I wanted to know; I wanted to see some of the tools and how they worked and where they were being used. I got sort of lost on the classification trying to understand where that was going and there wasn't really enough time for me because that was really quite a new look at things, and so that one sort of left me.

I would have liked to have known more as to why this one is having more success than that one, and this one, but you know, I thought it was fine. At some point I thought, I really did want to see more of the things in action, and things being used, and how they fit, especially on like the tools. It's like, okay, the tools are things that I didn't know about all these tools, and so, seeing them sort of in action and how they were all used would have been kind of fun.

RAM MOHAN: Thanks, Avri. Kaveh?

KAVEH RANJBAR: So, I'm loving it. So basically, I see a lot of big opportunity here, and very good feedback, so I think based on what I received, yes,

this should be used more as an opportunity for dialog, so I'm already trying thinking of some possible models. Definitely we will need more time, but yeah, maybe short presentations or no presentation, we'd focus on the strategy points for the board, so when requesting for agendas maybe we should highlight them, and a lot of time for discussion back and forth, that's is what I am envisioning.

RAM MOHAN: Thank you Kaveh. Lito?

RAFAEL LITO IBARRA: Thank you. I agree with most of what has been said. I would like to highlight a couple of things; the first one on the agenda I think is very good that it would be prepared in a joint effort by TEG, OCTO, BTC and also I agree that presentations should be shorter maybe, and a little less technical and with possible implications or consequences of the certain technology could have in the eco system and the eco system that we will give more value to the technology we are seeing.

And lastly, I would like to also include a follow up on previous subjects that we have seen. So, for instance, one of the last times we saw block chain, maybe we can follow up on what is

the current status of that technology, and so it is with DNS privacy now, and so on. We can follow up on the next meeting.

RAM MOHAN: Thank you Lito. Warren?

WARREN KUMARI: So, one thing I did sort of want to Ramind people is that the purpose of the TEG and TLG are supposed to be as a technical resource for the Board to ask questions of. I mean, Board members have wide range of backgrounds and things that they specialize and have knowledge in, and we exist, to a fair extent, so that if Board members have a set of questions or things they'd like to know more about, we're a resource that you can ask us questions.

I think everybody here is deeply passionate about the stuff they care about, and so they will be more than happy to talk your ear off, so if any Board members want to know more about anything, feel free to ask the TEG, TLG, whoever, and we'll be happy to talk.

RAM MOHAN: Thank you Warren. I guess that brings, I guess Warren, that brings it to my comments or my feedback. At the last time that

the TEG met with the Board, after the meeting was over, I sat down with some of the TEG members and I said, "Do we really need—what's the value of this meeting?" Because there's a set of things where there are folks saying things that just fly over everybody's heads, and then there is another set of things where it seems like it's just somebody's passion or somebody's interest, and you just come and present it, right? And not much information on what the relevance is.

But you already heard other folks say that. I have some thoughts on what would make it really interesting for me, as a board member in these sessions. What would make me make sure that I not only schedule this, but I want to be here and other Board members to be here. I'd love to have it where there are really four segments to the meeting; one is an explain segment, you know, and it's just pick a topic, and just explain it, explain it in a way that allows for the board members to have some, a better level of understanding.

I'd love to have a segment that is next-gen; things that are coming that you ought to know about, and that part I don't expect to get deeply tech, but that part is, this is what's happening next, and this is what's on the horizon, that would be interesting. It would be really interesting to also have I think a live segment. Last time where there was a live piece that was

shown, it was very, very fascinating to watch something happen, and we know that demos usually don't work when they are live, but it's actually a useful thing even if it's a cam demo, right?

But it makes it come alive, it makes it real. And the last segment I think it would be useful to have a segment of what should we, as a Board, worry about? What are the risks on the horizon that we ought to be thinking of? And, it doesn't assume that it's not already being thought of, so the intent there is not about, are you so dumb that you can't even figure that out? But it's really in the area of; we think these are things that are risks on the horizon.

So those four things in there I think would make for an extremely valuable interaction, thanks.

ADIEL AKPLOGAN: Thanks very much, Ram, and just in time, we just reached the time allocated to this session. I think the feedback session was very, very useful and for the OCTO as well because that will help actually work with BTC to ship from a future meeting. I will just withdraw my last point of the agenda, which was talking about the next meeting and talking about the 5G which was discussed on the email briefly. But we'll package that into preparing the next Board and TEG meeting. Thank you all for your

contribution and your valuable input, thank you to all the presenters. I think we have reached the end of this meeting, thank you.

[END OF TRANSCRIPTION]