

Introduction.....	1
Attendee roll call.....	1
Open Items from Prior Meetings .....	3
Github .....	3
Fellowships.....	3
Speaking Engagements and Work with Other SDOs .....	4
Improve publication/messaging of tech-related efforts .....	4
Technical track at ICANN meetings .....	5
DNSSEC Policy gap in RAA.....	5
Improve DNSSEC metrics.....	5
ICANN should encourage adoption of DANE .....	6
SSAC looking at RPKI status and actions.....	6
Facilitating back-end DS management automation.....	6
Internet of Things Status Update.....	7
Dependencies between Internet Operations and ICANN.....	12
Asynchronous dependencies shifting to synchronous dependencies .....	12
CZDS .....	12
Questions Provided by ICANN Board Members.....	15
Large scale technical evolution. ....	20
Problems with centralization of data and routing .....	22
IPv6 deployment update. ....	24

## Introduction

*[This is an edited transcript of the Technical Experts Group, 11 Feb 2015. If the editing has introduced errors or lack of clarity, that is due to the editor, not the original speakers.]*

### Attendee roll call

>>KAVEH RANJBAR: Kaveh Ranjbar, CIO, RIPE NCC, and K root operator.

>>MARK KOSTERS: So Mark Kusters, CTO of ARIN.

>>HOWARD BENN: Howard Benn representing ETSI, and also I'll pass on the apologies for my colleague, Francisco da Silva, who unfortunately can't make it to this meeting.

>>DANIEL DARDAILLER: And Daniel Dardailler, W3C.

>>SUZANNE WOOLF: Suzanne Woolf, Root Server System Advisory Committee liaison to the ICANN Board, which makes me both a Board member and technical.

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

>>LARS-JOHAN LIMAN: Lars-Johan Liman, operator of i.root-servers.net and also co-chair of the Root Server System Advisory Committee.

>>JIM GALVIN: Jim Galvin from Afilias, representing SSAC. Vice chair.

>>DAN YORK: Dan York, Internet Society.

>>WARREN KUMARI: Warren Kumari, one of the Internet Architecture Board reps to the TEG.

>>PAUL WOUTERS: And Paul Wouters is the other one.

>>WENDY PROFIT: Wendy Profit, ICANN staff, Board support.

>>BRUCE TONKIN: Bruce Tonkin, ICANN Board.

>>DAVID CONRAD: David Conrad, ICANN CTO.

>>STEVE CROCKER: Steve Crocker, ICANN.

>>PATRIK FALTSTROM: Patrik Faltstrom, chair of SSAC and also operator of I-Root.

>>ASHWIN RANGAN: Ashwin Rangan, CIO for ICANN.

>>MARC BLANCHET: Marc Blanchet, Viagenie, IAB member.

>>MARKUS KUMMER: Markus Kummer, ICANN board.

>>CHERINE CHALABY: Cherine Chalaby, ICANN board.

>>RINALIA ABDUL RAHIM: Rinalia Abdul Rahim, ICANN board.

>>KIM DAVIES: Kim Davies, ICANN staff.

>>KUO-WEI WU: Kuo-Wei Wu, ICANN board.

>> (saying name) representing ITU-T standardization sector.

>>JAY DALEY: Jay Daley from .NZ.

>>PATRICK JONES: Patrick Jones, ICANN staff.

>>MARGIE MILAM: Margie Milam, ICANN staff.

>>FRANCISCO ARIAS: Francisco Arias, ICANN staff.

>>DAVID CONRAD: I went through the transcript of the previous session and tried to collect sort of a synopsis of the various suggestions that were made in the section that we were talk about how to improve ICANN's technical stature. This list is sort of what I came up with and I wanted to give sort of a status update of where we are in the various items. This -- you know, I understand that there was -- you know, these were just suggestions that were provided by TEG members and there was no commitment made by ICANN to actually do these. But I did want to provide information to the TEG members that we actually are taking this stuff into consideration and are trying to do things necessary to improve ICANN's technical stature.

### Open Items from Prior Meetings

#### Github

So Jay Daley raised the question of increased use of GitHub, and during that session it was pointed out that we actually do use GitHub. There -- since LA, there was a little reconfiguration of GitHub and we're in the process of moving some things around. The repository on GitHub will actually include all of the sources for stuff that we use that are appropriately licensed. If there's anything that's non-proprietary that we can make available, the intent is to use that repository for that purpose.

It is currently up. I don't know what the status is with regards to whether it was public or private. It's part of the configuration that I mentioned. Kim, do you know offhand if it's public right now or is it still private?

>>KIM DAVIES: Everything that's available at that URL is public. It's more a case of we have a lot of private projects. We're going through the process of working out which ones can made be public.

#### Fellowships

>>DAVID CONRAD:

Moving on to the next one, so Jim Galvin had suggested that we increase the use of technical fellowships. I am actually actively discussing with my management the idea of creating secondments and sabbaticals where people who are interested in working on ICANN-related technical matters could have like office space and maybe even a laptop -- hopefully not Windows -- that would enable --

Not to say anything bad about Windows. Don't -- that's just this morning -- the stuff this morning was sort of entertaining.

And I'm actually -- we're right now in the middle of the budget stuff. You know, coming up with the budgets. And I have -- I'm going to be meeting with my management to try to ensure that that is a -- something that we can move forward with. Increased engagement in technical groups from both Warren and Daniel.

#### **Speaking Engagements and Work with Other SDOs**

So since the LA meeting, which I would like to remind people was about 70 working days -- seven zero working days -- ago, we have had a non-trivial amount of engagement in technical groups. I spoke at PTC on two sessions. They weren't really all that technical. It was about Internet governance.

I spoke at the Future Internet Conference in Nanjing, China. We gave a presentation on NANOG. Ed Lewis, who is now my evil mignon, actually gave a presentation at NANOG on the new gTLD program.

And at the Honolulu IETF, we had I believe ten people from ICANN participate. In addition -- just to embarrass him, Terry, could you stand up for a second? I would like to introduce one of the new Internet area ADs for the IETF, Terry Manderson.

>>DAVID CONRAD: So we are trying to be engaged. I probably missed a number of engagements, particularly in the areas of SSR that the SSR team is involved in, but we are trying to ramp up our engagement and will continue to do so.

A popular one, improve publication/messaging of tech-related efforts from is Dan, Jim, Martin, and Warren. So I am actually -- yeah.

>>JONNE SOININEN: Just one point that I wanted to say still about the engagement. I mean, like we also have now -- from the Board, we have Suzanne Woolf, who is the RSSAC liaison who has been selected for IAB as well. So we are cross-populating those organizations as well. And she's very happy to be in the limelight, I guess.

>>DAVID CONRAD: I wasn't aware of that. Congratulations, Suzanne.

>>DAVID CONRAD: Should also mention you're still co-chair of DNSOP, right? So you're really bored. You need more to do.

#### **Improve publication/messaging of tech-related efforts**

I'm actively working with ICANN's communications team in relation to trying to improve the messaging and content related to the technological efforts we're undertaking. Since it has been only 70 working days since the last meeting and there were these minor details like the Christmas holidays and, you know, other various holidays, and, you know, preparing for this particular venue, I haven't

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

gotten as far as I would have liked on that, but it is an area that I am actively trying to improve.

#### Technical track at ICANN meetings

Filiz and Patrik suggested a technical track at ICANN meetings to try to provide a venue for the people who are interested in technology to meet and participate together and hide from all the lawyers. I didn't say that. This -- I actually raised this with the meetings team back in early January, and as you are aware, we are restructuring our meetings somewhat, and it was agreed that we will be working to create a technical track for the Buenos Aires meeting. Not fully sure what that means at this point. We're going to be working with the meetings team to develop that. If people have specific suggestions, please let me know and I will work to try to ensure that that -- that those suggestions can be met, as much as feasible.

#### DNSSEC Policy gap in RAA

Moving right along, the -- we fixed the DNSSEC policy gap that was identified at the LA meeting. Not to put you on the spot, Francisco, but do you have any information about that?

>>FRANCISCO ARIAS: So Jim and I talked briefly about this a couple of times, so Jim, help me here. I'm trying to remember where we left this one.

I think we didn't have enough clarity still on what the issue was, did we?

>>JIM GALVIN: No. At the time -- I mean, the issue is that there's just a -- there's a slight lack of synchronicity between what the RA requires of registries and what the RAA requires of registrars, and there's a question of mandate and how you make things happen and whether you're obligated or not to provide services.

And the way that we left it the last time that we had talked was you were -- I believe the action was to you, I think. We haven't followed up since then. You were going to go back and talk to -- to some legal folks on your side to get an evaluation of what you thought about that and then we were going to talk again. But I confess, I didn't even bother to follow up with you either on where you were going on with that, so -- but it's only been 70 days? Isn't that what you said, Dave? So --

>>DAVID CONRAD: Okay. Thank you. We will follow up and try to nail that one down.

#### Improve DNSSEC metrics

Improve DNSSEC metrics to aid deployment from Dan. So one of the initiatives that we're undertaking is what -- something called the Internet health indicators, and

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

that would -- some of those indicators will be related to DNSSEC. The statistics will be published, so that should help address that one.

Encouraging DANE from Patrik Faltstrom.

#### **ICANN should encourage adoption of DANE**

So we are in the process of implementing DANE. We're anticipating that we'll actually have deployed DANE within our infrastructure by the Buenos Aires meeting.

We're also working to pivot our activities in encouraging DNSSEC to focus on validation, to provide an infrastructure that DANE could actually take advantage of.

The effort there is still underway. We would look for any other suggestions on how the community -- particularly the technical experts -- would think that we could help encourage the deployment of DANE, and I see Dan wanting to comment, surprisingly.

>>DAN YORK: Right. And I know we declared this a DNSSEC-free meeting zone so I will try to just be brief and say I would be very curious to talk a little more about the latter part about that, encouraging validation because it is a topic that came up in the DNSSEC workshop that took up the part earlier, so let's definitely talk more on that.

>>DAVID CONRAD: There was a suggestion from Wilfried Woeber and Dan regarding having SSAC look at RPKI, and I will defer to the chair of SSAC about what the status of that is.

#### **SSAC looking at RPKI status and actions**

>>PATRIK FALTSTROM: Yeah. That's me. So regarding the RPKI, we have not started to look at RPKI yet, but it's on the list. If -- one of the things that we are looking at, together with the other SO and ACs inside ICANN is to look at priorities. So if it is the -- so the more other parties want us to look at RPKI, the higher it will become on our priority list, but at the moment it is not prioritized high enough to be a work item.

>>DAVID CONRAD: Okay. Thank you very much.

#### **Facilitating back-end DS management automation.**

And then the last one was related to facilitating back-end DS management automation.

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

Jim had provided that, but I've heard it from a number of places. I'm aware that Olafur Gudmundsson, who is now at Cloudflare, is working to -- I guess he's developing an Internet draft on this and the technology at Cloudflare. We at ICANN are still sort of looking at how we can facilitate this. So any -- did you want --

>> I'll just say there's a much larger discussion that we just had that we do need to involve you at some point.

## Internet of Things Status Update

>>HOWARD BENN: Thank you. First of all, apologies to the Samsung advertising here. I was going to transcribe these onto the ETSI template. But after many, many years of traveling for the first time ever, I forgot my power adapter and ended up with a laptop with no battery in it. Fortunately, the shops in Singapore are very good and I managed to buy one.

So this really is just a bit of background about what's going on in the Internet of things. Things are changing very, very quickly. In fact, since I last created these slides, there's been some quite significant changes. So move on to the next slide.

The numbers are large. Not quite sure whether I believe the 50 stroke 80 billion connected devices in 2020. Still sounds like quite a lot to me. But those are some of the predictions that are coming out at the moment. Maybe it will be only 20 billion. It's going to be a very large number. And we can see this happening all the time now. So, increasingly, there's an increasing number of devices getting connected to the Internet.

The moment we're up to -- well, this slide, again, was 7.2 billion connections. That, of course, is way out of date now. Because it was 70 days ago. We're up to 7.38 billion mobile connections at the moment.

The next number I got from Wikipedia -- you guys should know much better than me whether it's correct -- 270 million names. I think what we call those unique domain names. Whatever that number is, it could increase significantly. I'll try and explain why. So move on to the next slide.

5G is the next-generation that we're working on right now. And this is the kind of thing that you normally see when you talk about 5G. It's all about smartphones, higher bit rates, getting gigabits per second. This is kind of the exciting front end. But an awful lot of work is going on on the more boring aspect. So lower data rate, wider area coverage for connecting devices to the Internet over cellular connections.

This is kind of what's happening out there right now. I think you could have chosen any device. There are increasing number of devices that you can buy at relatively low cost that will connect via your home Wi-Fi to the Internet. And some of them have names.

I think it's highly unlikely, as we move forward when we start talking about large numbers of devices, that the devices themselves will be directly addressable with a domain name.

Here's a little bit about what the industry is doing. And this is where things have changed slightly since the L.A. meeting.

So, if we have a look at standards and what's happening -- so I've chosen smart home because it's the one that most people understand. It's easiest to explain. But these cover a whole variety of different scenarios, everything from smart city, smart energy, everything is getting smart these days.

So, if you look at within the home at the moment, the way the standards are being written is you end up with three separate domains. So you end up with devices. The devices then connect to some form of gateway within the home.

And there are a whole raft of standards out there for connecting those devices to the gateway whether it's Z-Wave or ZigBee, Wi-Fi, you name it all, different sorts of ways of doing that.

What the industry is doing -- and OneM2M has spent a long time standardizing is making that a very simple, easy process.

So the onboarding is one of the trickiest things. You enter your DIY shop. You buy a device, whether it be a new light switch that you fit in your house or a new light bulb. You can just plug it in, press a button. And hey, presto, it all gets connected.

The other side of it is on the cloud-based services and the cloud-based platforms. And this is where I think there's some difference in between two of the larger consortiums that's just been pulled together. So we've got OIC, the Open Interconnect Consortium, and the Allseen Alliance. And what these two groups are trying to do is to pull together these different standards to make easy, operable systems.

Open Interconnect Consortium haven't opened up specs yet. But they're going down the line of having more intelligence in the device that sits in the home. Whereas, the Allseen Alliance are very much along the lines of getting everything in the service platform in the cloud, all the provisioning directly through the cloud.



Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

Now, this does have an impact on ICANN. Because, if you have a look at what is going to get a domain name, nobody is quite sure what the final outcome is going to be. But it will either be that service platform -- and that will be a Google or a Facebook or who knows who's going to come along and create the next great Internet thing.

And, therefore, it could be business as usual for ICANN, really. It will just be kind of the same old ways of doing business.

If we move to the other scenario where the device in the home gets a domain name, I think things could be significantly different. Because the number of devices will grow amazingly quickly. So what you may see is a greater demand on domain names. We may see greater pressure on some of the DNS. Difficult to know exactly what's going to happen. But I thought it would be good at least to start the conversation now to say from our world, this is the thing that we should be looking at. Thank you.

>>DAVID CONRAD: Thank you. Question related to the addressing. Presumably, these devices are currently both -- they're dual stack, v4, v6.

>>HOWARD BENN: Yeah. So the gateway will definitely be dual stack. It's still unclear exactly where the devices are going to go. So the devices at the moment tend to be using proprietary radio standards where they don't actually get an IP address. Some devices do. Some high-end devices. Nest, for instance, you'll end up with an IP address in the device. Increasingly, everything will go to v6.

>>DAVID CONRAD: Yeah. I personally actually deployed the Iris system, which -- there's a hardware store in the U.S. called Lowes. And they actually made some sort of arrangement with Iris. And it uses Z-Wave and ZigBee. And those do not use IP addresses? Is that correct?

>>HOWARD BENN: Yeah. What you do is, depending on which standard you end up implementing, whether it be NQTT or something like that, you do end up with a name. But it doesn't directly then translate to an IP address for the device. But it does within the gateway in the home.

>>KUO-WEI WU: It's very interesting. As you mentioned about the domain name in IoT, when you're talking about the domain name structure, it's very different to here these days. And also, in that area, I don't know how the governance structure began to develop.

>>HOWARD BENN: Yes. One of the questions I often think about is why do we have these two separate splits in architecture? And a lot of it actually doesn't come down to technology but things like governance. Privacy and security are the two big

ones. So most people don't want to have their front door opened by a Facebook page. On the other hand, everyone is well aware that they don't want their box in their house to be hacked, because they can't keep up with the latest security. Interesting space.

>>KUUO-WEI WU: This is Kuo-Wei Wu again. As I know, particularly, if the people using RFID to build kind of the domain name structures, is the current service in the VeriSign?

>>HOWARD BENN: I don't know the answer to that one.

>>KUUO-WEI WU: No, when the IoT, the technology somehow people using RFID -- RFID, yeah. And somehow, actually, the VeriSign came to Taiwan and said they're running a server. So, if the Taiwan wants to run the national DNS, they're more than happy to connect that.

You know? So that's -- actually, go back to the original question. At that area, the name structure governance actually is still empty. It's kind of -- well, I don't know what would be happen when that is really going on in the future.

>>DAN YORK: So, David, it's Dan York. Going back to your question about IPv6 and the piece there, there are a number of these lower level IoT, you know, networks and connections that do support IPv6 and work in different ways. You mentioned ZigBee. And they did recently come out with -- I think it was last year in 2014, it was a 920 IP or something. There was a new specification they had that does support IPv6 as part of that. So it varies -- within the program that I'm part of with things -- we've certainly been talking to folks around that. I don't know the full status of where they are, but --

>>PATRIK FALTSTROM: Yeah. So what happens is EPC Global that run the naming system did a procurement process around 2004. And VeriSign did win that procurement to run the root under ONSEPC.com domain name. I looked for the web page for the specification. And that web page -- the link to the EPC web page, that link was dead. So you cannot read the specification on the original space. On the other hand, the name servers are still up and running at VeriSign. But we need to reach out to EPC Global if it is the case that you want an authoritative response to that question.

>>DAN YORK: David, Dan York. Just to follow up on the two, I meant to say the work that ZigBee is doing as well is based on work coming out of the IETF. There's other groups that are also involved with that. There's a whole internet of things directorate within the IETF that is focused around standards in this space.

>>DANIEL DARDAILLER: Daniel Dardailler, W3C. I'd like to come back to the number I saw, 50 billion or so. Today, when you get a web page, you actually generate hundreds of connections, not just one get HTTP web page made of multiple pieces. And I would expect that, with the Internet thing and a layer above that we call the web thing, we'll see the same scenario. That is, whenever you get data on a thing, then you would like to add context and you query dozens of other resources to actually do something with this data. And for the web thing, in addition to the real thing or the Internet thing, you also had virtual thing, like people. Someone, you know, is not part of the web thing, not part of the Internet of things. But the other thing the Internet evaluation of the number has to be re-evaluated in terms of the connection, not just the object itself. But the number of TCP connections that are going to be generated is going to be huge because of the added logic of the platform above it.

>>WARREN KUMARI: Warren Kumari. And Dan stole part of my answer. I was just going to point out that the IETF is working on 6tisch and 6LoWPAN, which are specifically designed to do v6 over these sorts of things. You know glossy networks, low power networks. And that seems to be moving along fairly quickly.

>>JAY DALEY: It seems clear to me that all of these devices in the house are going to need a name of one sort or another. The light in my bathroom will be called "bathroom light" rather than any obscure series of hexadecimal digits. The problem, of course, is whether or not that will be a domain name. And I wonder whether or not there is a role for anybody around this table to be promoting the idea that it should be a domain name rather than some other form of naming system. Because we know that domain names will scale, will work, are secure, et cetera, et cetera rather than creating a new naming system for this.

>>WARREN KUMARI: I guess I can respond to Jay quickly. Actually, the home net working group was working on a set of stuff to allow naming of that sort of stuff. And the previous TEG appointee from the IAB was one of the big editors of that, Daniel. And the current draft has expired. I think it was replaced with an IETF version. It was originally an individual. And I can look sometime and get you more info. But there is definitely work to make it so that at least within a boundary, you'll be able to have things with same names. There is, of course, some difficulty then, if you want these to be global names and where exactly you draw the borders and also where in the name space that fits. This might end up being sort of a special use type name, which is a topic we might get into later.

>> Can I just add -- I hope this isn't illegal at this meeting. But I was talking about marketing. I was talking about somebody going out to promote the cause of DNS rather than the technology in this particular case.

## Dependencies between Internet Operations and ICANN

>>JAY DALEY: Thanks. I managed to explain the slides last time, but we didn't get a chance to talk about it. So that's why I'm doing this again.

The first slide, this is talking specifically about ICANN and a change in the dependency that I see taking place between Internet operations and ICANN.

### Asynchronous dependencies shifting to synchronous dependencies

Where there were asynchronous dependencies before, I think we're now shifting to synchronous dependencies. Where we had a decentralized we're moving to centralized dependency in some cases. And where we saw in-sourced ones, i.e., ones that are just controlled by small groups of technical operators that we understood moving to a mixed source dependency where some things are now being contracted out to third parties.

So here are some examples. The centralized zone data service has centralized authorization and centralized data access. That's new compared to the system that was in place previously.

The Expert Working Group on WHOIS is recommending centralized authorization, centralized data access, and I believe centralized break in and theft as well. That was a joke.

And then, finally, bootstrapping for RDAP, I should say, is now talking about the realtime use of a centralized infrastructure. So we will have web pages doing a JavaScript call to a JSON file that's on an IANA distributed CDN at approximately 50 to 100 million lookups per day, I would imagine, given that it's WHOIS lookups.

So my real question is are we sleepwalking? Do we know this is happening? Do we want this to happen? And where are the guiding principles behind it? I'm not necessarily saying this is a bad thing or a good thing, because this is a change in dependency that, to me, is close to an architectural change. And we're not actually planning this.

### CZDS

>>DAVID CONRAD: On the CZDS, in a previous life, I had some involvement in the development of the pilot prototype implementation of CZDS. And this was around 2011, 2012. In that version of CZDS, it basically acted as a password storer. The access to the data was still held by the registries. And the registries actually assigned the user ID passwords.

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

And the sole purpose of CZDS was to store the -- basically, the collection of user ID passwords to facilitate the end users who wanted to obtain the centralized zone data service from having to manage all the user ID passwords themselves.

We did allow for the data to be stored on the CZDS servers, but that wasn't a requirement.

I will gleefully admit I have not followed that particular part of ICANN. But Francisco undoubtedly has. So --

>>FRANCISCO ARIAS: Francisco Arias, ICANN. So CZDS allows two options. One, in which, as you said, David, the registry can still be the one receiving the database, hosting the database -- sorry, the zone files -- and offering access to the users, or the centralized version in which the users register in CZDS and obtain the zone file directly from there. So they are the two options in the system.

>> My experience is that most people use the latter. In fact, I'm not actually aware of anybody who uses the former. But we don't poll that many zones.

>>DAVID CONRAD: Right. I think at least back in the day, when I was involved, the -- no one actually stored the zone data on the CZDS server. But with the new gTLD program I presume large number have decided it is much just simpler to store the data in the CZDS since it was a requirement to make that data available. I just wanted to clarify that.

>>DAN YORK: Dan York. Jay, are you really asking us are we making a conscious choice around this? Is that a better way to say it as far as these changes are happening? Are we doing this consciously? Or are we just letting it happen?

>>JAY DALEY: Yes. I'm saying I don't think we are doing this consciously. If anybody thinks we are, please let me know. If we could go forward to my questions, again, please. Thank you.

Most importantly, I think we have guiding principles to discuss here. One of the main principles the Internet has been built on is compartmentalization of failure. And we are actively going against that in some cases here. It may or may not be a good thing, but I think that needs some form of conversation. And I think it is actually a rather important top-level strategic item for ICANN, if this is not a simply something that's technology. Because it's a risk.

>> Yeah, I will say that I personally -- I'm not speaking for ICANN -- personally uncomfortable with the WEIRDS approach for the bootstrapping stuff.

I understand the reasons. Marc Blanchet actually conveniently explained to me the reasons and rationales. I'm hoping that, in the evolution of RDAP, at some future point, as the underlying infrastructure becomes more capable of dealing with, say, something other than A and text queries in JavaScript, that we will actually be able to do something more, shall we say, intelligent with the bootstrapping protocols. But my impression -- and actually Mark may want to comment here -- is that the bootstrapping selection that was made was sort of a -- sort of the compromise -- I don't necessarily want to say least worst out outcome, but the outcome that was sort of the only one that was able to gain sort of somewhat consensus. Do you want to comment?

>> Well, that's mostly it. The working group looked at different possibilities, essentially, two categories where IANA was registry based and the other was in DNS. And, you know, we spent, you know, a few meetings on this. And the final decision was, as you said, the least worst. There were no perfect -- you know, given the constraints that -- and requirements that you know, we set at the beginning about what would be the use.

>>JAY DALEY: Right. And the implication of the choice by the IETF is a non-trivial resource requirement on ICANN in order to -- you know. The bootstrapping file is mostly static. So it does fit easily into CDN. And ICANN, of course, can make use of CDNs. But it does imply a certain set of operational considerations that will undoubtedly -- Kim, if you want to talk to this at all -- will probably have some implications on processes within the IANA.

>>KIM DAVIES: You know, I think on the bootstrapping mechanism itself, it's been designed in such a way that it's minimal change, I think, to our operational practices. As has been noted, I haven't heard that specific projection in terms of number of queries. But we are mindful that it may have a significant load to it. So we're rolling it out on a CDN.

In fact, while we're going through that process, evolving the way we deliver all the IANA registries to be delivered via a CDN as part of that exercise as well.

I think the bootstrap draft, if I'm not mistaken, implies that clients should cash the bootstrap file itself. We're looking at having a sort of long-lived expiry on that file, if we can, to try to minimize the load on us.

I think, from an IANA staff member perspective, I think this is one of those things where we saw it as a healthy debate in the IETF that came with a resolution. Our responsibility in the context of IETF protocol registries is to implement the IANA considerations as they have been settled within that community. And that's what we're seeking to do in this particular instance.

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

>> I'm curious to know if there's any risk analysis done that would help understand how much security needs to go around the integrity of this file as it's distributed.

>> I don't know if -- I think I'm probably being fairly confident in saying I don't think a formal risk analysis has been done. The IETF may -- the WEIRDS working group may have sort of examined that in an informal fashion. I don't actually know. Mark, do you have --

>> Not precisely.

>>DAVID CONRAD: But I guess one of my call it an assumption is that the current model for bootstrapping is a -- almost a temporary solution that, in the longer term, as the underlying infrastructure becomes more capable of doing more unusual things with, for example, DNS queries, that we would be able to, eventually, maybe standardize on an evolutionary change that would allow for an on-demand fetch of the -- via DNS of the necessary bootstrapping information. But that's just my hope, I guess.

## Questions Provided by ICANN Board Members

And now we're going back to the section of the meeting where these are questions or -- yeah, questions -- provided by ICANN board members.

And we -- they were looking -- these are topics of interest. There were four that were provided. Sort of summarizing them, they were just

- increased government involvement,
- large-scale technical evolution,
- IPv6 deployment and
- engagement improvement.

So on increased government involvement, the question here was what impact might increased government involvement be having on DNS operators and the Internet as a whole directly or indirectly relevant to ICANN's mission?

And I sort of allocated about 10 minutes to this. If anyone has any thoughts or comments about this, jump in. Yes, Warren.

>>WARREN KUMARI: One aspect of increased government involvement is the realization that some of the government involvement, involves things like pervasive monitoring. Snowden, et cetera.

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

This has made a number of people somewhat nervous. The IETF has the view that pervasive monitoring is an attack. So it is attempting to design workarounds for this or mitigations for this.

One of those is trying to actually encrypt the DNS. You leak huge amounts of information just through your DNS lookups.

Currently, the IETF is trying to encrypt most traffic. But that doesn't really help if people can figure out what you're doing just by seeing where you're going.

Obviously, encrypting the DNS is a large undertaking. The initial part is likely to be encrypting from the stub resolver to the recursor, which will have some implications, potentially, for ICANN.

But there may be a second phase of this, which will be encrypting from the recursive to the authoritative service. And that will have much more impact, I would assume, for ICANN.

This is all work that's happening in the deprived working group, which just spun up recently. We had our first meeting in -- Hawaii, last place we were at.

>>DAVID CONRAD: So, with regards to the implications of ICANN of the stub to recursive part, what are you sort of imagining there?

>>WARREN KUMARI: Well, there are some recursive operators who show up here who are going to have to -- if they want to support this, do some fairly significant increase in capacity, stuff like that. We're assuming that this is going to be over TLS and TCP connections. That would probably be the main thing for ICANN folk. Obviously, the second part, encrypting from recursives to authoritatives will be a much larger impact. Although we haven't figured out what we're going to be doing, so it may be something that actually involves encrypting from the stub all the way to the authoritative.

>>DAVID CONRAD: And "we" here is?

>>WARREN KUMARI: Sorry. "We" here is folk from the deprived working group, which in full disclosure, I'm one of the cochairs of.

>>DAVID CONRAD: Right. And there's also the Q&A minimization effort. Do you want to talk a little bit about that?

>>WARREN KUMARI: Sure. Yes. The Q&A minimization effort is also to try to remove some of the sort of information disclosure. Currently, if you want to resolve and you have an empty cache, www.example.com, you contact the root server and



Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

you send the entire query name. You actually send to the root server www.example.com. It refers you to dot com. You go to the dot com servers. And, once again, you send the entire string. Www.example.com. It sends you off to example.com.

There's no real reason that you have to tell the root servers the full domain you're looking up. There's no reason you have to tell the dot com servers the full domain you're looking up. It seems to only be an optimization and only a very slight one at that. And so there is work underway.

It's actually happening in the DNS ops working list or working group to make it so that, when you contact a name server, an authoritative one, you only send as much as you need to.

So going back to the example of www.example.com, when you reach the root, you would just ask it where dot com is. You wouldn't send the entire string.

This has some implications. There will be some loss of visibility in the folk who run TLD servers. They won't see all of the lookup -- all of the string that's being looked up. It also means that stuff like the DITL data, the day in the life data, loses a fair bit of visibility. That's an unfortunate side effect or, depending on what your views are, a good side effect.

>>DAVID CONRAD: Suzanne.

>>WOLFGANG KLEINWACHTER: Thank you, David. Suzanne Woolf. To the larger question, in addition to kind of the big picture changes that Warren was referring to, I have noticed also, though, that there's been an impact on -- and I'm speaking as an RSSAC member and a root server operator, there's been an impact that's very interesting of this kind of discussion and I think very positive which is that it's increasing the pressure, in my opinion, on root server operators towards more transparency about what they do and what information is available through the network and what root server -- what the root servers actually do. And I think that's actually good, that there's just more drive for people to be able to tell what's happening to their network traffic. Who knows what about them based on what they're doing the network. I think that's actually very positive.

>>PATRIK FALTSTROM: Patrik Faltstrom, and I also happen to be co-chair of the ICG IANA transition coordination group, and one of the things that's happening there is that we are discussing, as everyone knows, what's happening and what the implications are for decreased government involvement because there's contract that might go away. I think in general, I think it's important to watch the discussions regarding accountability and otherwise related to the IANA functions because it might be possible to understand a little bit better what the interest in what direction

various forces want to push things, even though that discussion should be focused on the IANA transition itself. So I think following that could be good to gather data that can help answering this question.

dd

>>WARREN KUMARI: Following on from what Suzanne said, the sort of information going to the root and calls for transparency there, there are some folk who are somewhat scared about even that level of information, you know the TLDs they're looking up, and they're considering doing things like slaving the root zone specifically so they don't leak any queries to the root. Partly the TLDs you're looking at but more importantly the lookups that result in NXDOMAINS, the things you look up that aren't actually in the root.

Somewhat related to all of this though is obviously the people who are very concerned about their information being seen by governments and so are building things like alternate name spaces, you know, dot onion, the Tor network, the dot DNS proposed pseudo-TLD which is specifically designed to not use any of the ICANN-provided DNS and folks like that, you know. This is all -- much of that is being driven from the pervasive monitoring concerns noted in Revelations, et cetera.

>>MARC BLANCHET: Related but a different perspective -- Marc Blanchet -- we can see that some governments but also not only governments but other entities are actually filtering by some means, you know, the DNS. So essentially from the point of view of the end user the DNS is lying. You know, the actual answer from the alternative server is not the same that you finally receive. So I think that's pretty significant, you know, in the DNS ICANN ecosystem because we're delivering something that is not there.

And we obviously have a solution for this. It's called DNSSEC, but then the problem is that we have it done, you know, from signing but we still have the foundation which is, you know, lacking. And so I think that's something that would really be, you know, a good solution to be deployed more than what we have right now with all the caveats.

The second thing I would like to make is regarding not only this one but the other one is that the Internet is being more and more ossified which make it more difficult to deploy new technologies, you know. More (indiscernible) technologies. The only thing you deploy is new web application nowadays.

>>JONNE SOININEN: Yeah, thank you. Of course, I -- basically I do agree with everybody with the things that have been said about the kind of like negative consequences of some governments being involved and getting interested and so on. But I have to just say that there might be, though, most probably remote possibility,

also the possibility that the governments might actually learn something of the Internet and know actually -- learn how it works and through that actually do better policies on the Internet because we have had for a very long time a situation basically where the governments have had very little understanding of how Internet works and many people have worked very hard to teaching them how it works. How much they actually learn from that, that's -- depends on the government. But still, this doesn't completely preclude that they would learn something from this and input it also in a positive way in the policies that they do.

>>DAVID CONRAD: I will simply say that I -- the idea that the governments would learn is something I usually term the optimism of youth, but I will go on. Jay.

>>JAY DALEY: Thank you. One of the things that I'm noticing is that governments are recognizing that DNS operators are part of the critical infrastructure of a country now, that there are too many things that operate over the Internet that they do not want to see fail at all or under any circumstances and so they are interested in, from an assurance point of view, about everything that goes on in the DNS operator. They want to know service standards, they want to know what backup facilities you have, they want to know those things. For a very good reason that their business and in their view their country depends on it increasingly.

>>WARREN KUMARI: Warren Kumari, just following on what Marc said. I'm sorry to keep banging on about DPRIVE (phonetic). One of the nice side effects of encrypting the DNS is if you can't see what's in the query, you can't filter it. DNSSEC at least let's you see that somebody has blocked your query or has changed the answer and then you can't use it whereas DPRIVE nobody can actually see what you're looking up and so they can't filter it. Whether this is a feature or bug, it's someone else's problem.

>>DAVID CONRAD: And Kuo-Wei, very quickly.

>>KUO-WEI WU: I think on this topic we know the governments, of course, they have their own internal, you know, mechanism they want to see that. Is a much better way is to communicate is to tell the government -- at least share the government the knowledge, you know, what kind of activity would be slow down the whole, you know, resolving issue, you know. And so at least they are then going to make trouble for, you know, the -- the thing we are worried about, you know. Just sharing more information might be a better way to resolve it.

>>DAVID CONRAD: And I'll insert myself into the queue that's closed. One -- right. One observation that I would make, particularly in the context of the DPRIVE work and QNAME-ization (phonetic) is that there are tools available currently made available at registries that allow registrants to identify domains that are receiving queries, even though they are no longer registered, and those are then used to --

people purchase those and put hosting pages and that sort of thing. I personally, again not speaking for ICANN in any way, personally see this as somewhat of a risk, a security risk, and the DPRIVE work, particularly QNAME-ization and the query encryption to the authoritative, would make it unfeasible for that information to be sort of roundly collected maybe.

So moving on to the next topic, thank you very much. Large scale technical evolution. What large scale technical evolution Internet of things internationalization, carrier practices to capture subscribers, popular closed platforms like Facebook and Baidu, is relevant to ICANN's work and how? Anyone want to start us off? Jonne.

### Large scale technical evolution.

>>JONNE SOININEN: Yeah. I'm not sure if this is the request but I'll answer anyway. I think we already discussed the Internet of things is clearly something that might have some sort of impact -- we don't know what yet -- on the naming and so on. And I guess that the internationalization is already something that is impacting ICANN quite a bit. About the others, I don't know.

>>DANIEL DARDAILLER: So one possible impact with the platform that has so much success is if you continue to have that much success you may end up to the extreme with no need for any DNS except facebook.com or, you know, some other platform and then you have a URL that you own below these domains, but you don't need your own domain. So there is a -- there is a side effect to the -- to the domain name with the popularity of the platform and the monopolistic -- the monopolistic approach where it's very hard to displace, you know, Facebook because they have already gained some battle here. So at W3C we are working to making open standards so there can be multiple Facebooks. In fact, you can be your own Facebook provider and you can exchange associate things without losing your data, without privacy risk and et cetera, but that there is -- this is still this possibility that you end up with a -- at the higher level with just a few platforms that, you know, deal with all the domain names and everything.

>>DAVID CONRAD: Thank you. Warren.

>>WARREN KUMARI: Warren Kumari. I guess fairly closely related to that is the increased use of apps where you don't actually have any idea what domain name it's using and so the domain name suddenly becomes unimportant. If you end up with a sufficiently large number of apps and you don't need sort of very unique strings or human readable strings that could have fairly large implications for the commercial side of this.

>>DANIEL DARDAILLER: Yeah, if I may add something, it's exactly that. I mean, if you look at the platform interaction, there is very few link that you can actually use. That have -- you have to use the link system of the platform which is uniquely constrained. You cannot link outside the platform sometimes, things like that. So it's really related. People at the platform level, they want you to stay within their platform.

>>DAN YORK: It's Dan York, and just a note for the transcript that the previous comment that was made was attributed to me, I saw it in there, but it was actually the other Dan down here from W3C. So whoever is writing those notes may wish to do that.

My comment was -- Warren stole part of it so we can, you know, trade on part of it. There's a couple of interesting things as we look at moving to mobile apps and the processes that are happening there. One is what Warren said was this increased -- or the decreasing dependence on domain names in space of just the app and you're using it and you don't really care about any of that and you don't know about that. The other interesting thing is about the actions of those apps themselves. At the DNSSEC workshop earlier today we had an interesting discussion that was brought up by Jeff Houston, right? It was in his discussion about some of the changes of things indicating that the Android app for Netflix tunnels its own -- it has its own built-in DNS resolver which was sending all of the traffic, all of the queries, to Google's DNS -- public DNS servers. They were doing that so they could route around the fact that they might have blocking and whatever network it's being connected to. And so, you know, to get the best DNS answers they were doing that. But that's an interesting dynamic, that the app is making those kind of direct connections and direct things without using the typical resolver and the pieces that might be on the local network.

So from the ICANN perspective, and I guess from the Board perspective, it's just interesting to understand that these -- this mobile environment is changing a bit of the underlying infrastructure that we've been relying upon for the typical ways that we work with that.

I think there's several other, you know, meta kind of levels of things we look at in terms of increasing reliance on CDNs, contents distribution networks, and other places that are starting to make these sort of overlay networks in terms of, you know, the layers of the Internet architecture. What impact that has around some of the DNS, I'm not entirely sure with some of those pieces. But I think it's interesting to take a look at part of that and similarly, VPNs, you know, and the changing nature of where DNS queries are being issued. I had somebody here who was asking me the question of why it was that she was unable to get to [www.facebook.com](http://www.facebook.com) from her home network in China but when she set up a VPN back to wherever, she was able to go and see it. Again, we're having different things of those DNS queries

originating -- that may have originated in one network are now originating somewhere else in another part of the network because of this translation that's happening there.

I would also point out for the Board's interest and sake, there's a recent report that came out from UNESCO which was called fostering freedom online, the role of Internet intermediaries, and it's a report that came out just in January that talked about this issue of Internet intermediaries, the platforms like Facebook, like other different systems like that. It may be an interesting document and report that -- for you folks to take a look at and understand -- look at this on a macro level.

### Problems with centralization of data and routing

>>LARS-JOHAN LIMAN: Lars-Johan Liman. I see several issues with these actually large centralization of systems. The Internet was designed to be a system of smart end nodes but stupid routers in the middle, and we're now seeing the pendulum swing back to the telephony system where you had a somewhat smart switch in the middle and very stupid telephones on the end. It will swing back again. It will take 10 to 20 years. It swings back and forth, I've come to realize.

But the centralization does carry a few problems with it. One is actually a stability problem. We talked about that -- Jay talked about that before in the different setting -- circumstances. But when you start to rely on really huge central systems like Facebook, Google, you also risk the stability because it's -- if that system breaks, you have lost a lot of the communication that you want to do. And I can just toss in a little remark that it's interesting that so many people and so many societies chooses to rely such high degree on a system that's build on best effort. But that's a different thing.

I also see integrity problems with these large systems. You have the intermediaries, you don't know what to do -- what they do with your data and that boils down to actually also your -- your trust in the code that you run on your machines. Are you -- do you feel in control of your computer? I don't know how much crap Apple puts into my computer nowadays, and I'm seriously actually considering going to a different operating system. But I see more and more trends that organizations send my data and my requests to places that I'm not aware of. That's an integrity problem for me.

And my third comment is that this also involves the freedom of speech, to some degree. To begin with, you have these large central systems who may apply their own censorship according to their own rules when you try to publish the information you want to publish. And also, if they -- if you have this pervasive monitoring, you may restrict yourself when you write messages to various Internet publishing mechanisms. For instance would you use the word "bump" in an email?

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

Hmmm, maybe not. Depending on where you live and where you're sending email it may be different. So these are issues that worry me in the large scale. Thank you.

>> My pleasure. So two comments. One related to all of these, as I said before, the Internet being classified. My point being that if people think that we can deploy new stuff easily, it's not the case.

My second point was about (indiscernible). I think I've been involved in this for some time, and I recently within the IETF so I -- IDNs have been gone through two versions essentially of the protocol ID. I was co-chair of the first working group, and so IDN 8 2003, IDN 8 2008 and recently IETF has -- you know, people found some issues and we may end on the next IETF talk about this in a BOF or something. So the overall arching point here is that it is and will continue to be difficult to map languages into, you know, the constraint environment of DNS and identifiers in general and therefore it will be always a compromise at some point which includes security. And, for example, within the ICANN currently there -- the LGR work which is related to defining what is acceptable in the code in the root zone and this is again all about compromise because you cannot map all the languages features into the DNS.

>>PATRIK FALTSTROM: Thank you, Patrik. To continue on what Warren said, I think what we have to think about carefully is the separation between domain names themselves not being important any more. But the DNS protocol continued to be extremely important. So I'm comparing with RIRs where IANA now handed out one IPv6 block or two to each one of the RIRs. And one day we might have five domain names registered. Done. But the DNS protocol will still be important.

>>DAVID CONRAD: Thank you. Howard?

>>HOWARD BENN: Okay. Very quickly. The GSMA and, hence, the mobile operators are getting very concerned about the way that more and more encryption is happening. So the increased use of TLS. Because they can no longer control the quality of service of their radio networks, because they have no idea what the content is. So there's been a big move to doing local distributions. So in the U.K., the popular one is some of the streaming services on the BBC. You can detect the iPlayer content. You can store cache out locally. But, of course, the operators can no longer see that if you use HTTPS all the time. So there is some concern there.

>>JONNE SOININEN: Yeah, just very quickly. I think that a little bit here we also have to see what are, basically, implications to ICANN as well and not just the -- and not to the Internet ecosystem as a whole. So many of these implications, for instance, privacy and so on, have little to do with ICANN.

But the idea that the inter name space would fragment into different pieces of that, some -- basically, some places would resolve to different names because of blocking or because of that -- some, basically, intermediary is providing -- could be providing their own name service for their own application and so on and have their own names that are not resolvable from the global DNS. This might have implications for ICANN as well. And we have to be a little bit -- look, follow that trend and see how that evolves.

>>DAVID CONRAD: Okay. Thank you.

And we have eight minutes left. So the next slide, please.

### IPv6 deployment update.

How are we doing on IPv6 deployment, and what implications does it have for those other trends in the access issues -- IoT, closed versus open standards and platforms. Anyone want to start us off on this topic?

>>DAN YORK: Yes. Dan York. Obviously, this is an area that I work with within the Internet Society around the deployment of that. The reality is that, if you look at measurements, statistics we're seeing, we're seeing large amounts of IPv6 traffic. If you go to the [worldIPv6launch.org/measurements](http://worldIPv6launch.org/measurements) Web site, you'll see that there's large percentages coming out of some networks. For instance, the one that I'm aware of with a lot of traffic coming off of Verizon's mobile LTE network, you know, is all -- is 60% I think, now around IPv6 coming in there. Google's stats show that. There's a lot of things.

I think that the question really here is where's the intersection with ICANN? And I think what's useful and interesting, I think, from the ICANN perspective is I think what ICANN's done very well is -- I mean, you have the interaction with your registrars and registries.

And the pieces in, for instance, the 2013 RAA requiring for operational perspective, IPv6, I think that's good. And I think ICANN should continue that and continue working with the registrars and registries to ensure that they are operating in that mode.

The other piece I would say is that ICANN can help with the deployment in certainly helping assist that registrars and registries allow registrants to enter the information in the provisioning interfaces. There are some out there that still do not allow the entry of a AAAA record in some of the user interfaces. So, to the degree that ICANN can help promote the registrars and registries to fully support IPv6, that will help with the overall deployment.



>>DAVID CONRAD: Okay. Warren.

>>WARREN KUMARI: So, following on from what Dan said, yes, there are some networks that are really good. 60%. But we're not completely there yet. On a worldwide scale, it's more like 5%. Whether or not this is a problem is unclear.

What ICANN can do to help promote this potentially as well as requiring that registrars and registries accept these records, maybe they can push for things like the DNS hosting providers or the registrars who provide web hosting also provide IPv6 access. That's one place where ICANN might have some control.

>>DAVID CONRAD: Mark, do you have any thoughts on this?

>>MARK KOSTERS: Mark Kusters, ARIN. So one of the things that the original registries do is they do a tremendous amount of outreach, as well as ISOC, to various communities -- hardware vendors, software vendors, et cetera. We go to various trade shows to say, hey look, we're running out of v4. Go join the v6 bandwagon. What's interesting is that the number of people who say what is this v6 stuff, what do you mean v4 is running out is, happily, going away. So people are actually coming to the realization that they do need to build these things.

We're just not there yet across the board.

And this is going to take a long time.

>>DAVID CONRAD: Kuo-Wei.

>>KUO-WEI WU: I think one of the things that ICANN we can do is I remember in the very early days, maybe six, seven years ago, the RIR actually go to the ICANN GAC to present the IPv6, you know, forecasting. But at that time, in general, I think it's a little bit too early because the data is not solid enough. And so you can see the government actually are very hesitant, you know, to turn into the IPv6.

But right now, even just 5%, even just 5%, I think we really see dramatic growth from last year to this year.

You know, so I would suggest that RIR -- you know, because they have all accept IPv6 measurement data. And it is a very good time right now to, you know, for RIR to reinitiate to present IPv6 status to the GAC. Because once the GAC -- the government accepts the message and the data, you know -- for example, as I know, the many of the government Web sites in Taiwan, even didn't turn on v6 yet. Because they kept waiting the real information.

Joint Meeting of ICANN Board and TEG  
11 February 2015  
ICANN 52 - Singapore

So I think that might be a very good time for RIR to revisit the ICANN GAC to tell them and bring the real data to them.

As we have three minutes left, I am going to take the chair's prerogative and defer the remainder of the slides in order to ask the board members who are present if they have any specific questions, thoughts, comments related to what we've heard so far. And I will start with Dr. Crocker.

>>STEVE CROCKER: I yield my time. I'm interested in what the others say.

>>DAVID CONRAD: Other board members, would you like to have any comments, thoughts, questions in any of the areas? Yes, Asha.

>>ASHA HEMRAJANI: Thank you. This is Asha Hemrajani, ICANN board. I'm still new to this. I do have a couple questions, which I think will take more than 10 minutes. So I just wanted to say I'll be in touch with you and that gentleman over there -- I can't see you clearly -- to ask you a little bit about what you just mentioned. I hope you don't mind. Thank you.

>>DAVID CONRAD: Kuo-Wei. Rinalia?

>>RINALIA ABDUL RAHIM: Thank you. Rinalia Abdul Rahim. I just wanted to say I really appreciated this session, especially the question on large-scale technical or technological evolution relevant to ICANN's work and the other question earlier. Because it helps us in the risk committee to look at what's coming and to help us prepare ICANN for the future. So thank you very much.

>>DAVID CONRAD: Okay. We actually have one minute and not enough time to go into the last question. I will actually try to defer that to the mailing list and look forward to any input people might provide. Thank you, again, for participating in this session. And I suppose the next time we will meet will hopefully be before Buenos Aires; but, if not, in Buenos Aires. Thank you very much.

[ Applause ]