HYDERABAD – Joint Meeting: ICANN Board & Technical Experts Group (TEG)
Tuesday, November 08, 2016 – 16:30 to 18:30 IST
ICANN57 | Hyderabad, India

DAVID CONRAD: Welcome, everyone. This is the ICANN 57 technical experts group meeting and a joint meeting between the TEG and ICANN's board.

We have a slight change to the agenda that had been posted. The initial topic, which was on Digital Object Architecture, had to be dropped. The presenter, Suzanne Woolf, had indicated she wasn't feeling well and has escaped with her life. I believe she will be back, but she wasn't feeling up to presenting today, so we've dropped that off the agenda and extended the agenda to allow for our dear friend, Warren, to talk about the IETF issues.

For those who are unaware of what the TEG is, it is focused on forward-looking technical and technology issues, particularly as those issues impact the use of the Internet's system of unique identifiers that in the view of TEG members, ICANN's board and staff should take into consideration when considering ICANN's strategies and operations.

The TEG is an informal group. It is not an advisory committee. It does not have a budget. Its role is to provide input to the

board and the board has no obligation to accept that input other than the fact that it's coming from experts and people who look very nice.

So with that, the agenda for this TEG, we have an update on special names issues and the status of the problem statement and SSAC defined problem space by Jim Galvin; the ETSI NFV, the network function virtualization, by Howard Benn; a work that was actually funded by ICANN, DNSEXTLANG done by John Levine, and Warren will be presenting on IETF-related issues.

So with that, I will hand it over to Jim.

JIM GALVIN: Thank you, David, and I see the slide up here so that's very nice.

I'm making this presentation today in part in my role as vice chair of SSAC, but it also turns out that I am chair of the work party within SSAC that is considering this issue.

So next slide, please.

SSAC has been considering the issue of the domain namespace and the presence of domain names and the collisions that result from their use in the Internet community at-large for most of this year, so this bit of update is here because we have come to a consensus, really, on the words to use to describe the problem

space and where we think we are and the issue as we see it for the ICANN community.

So the first thing up here is to explain what we mean by "domain namespace," and basically it is all of the possible domain names that you can get in a tree-structured hierarchy of individual labels.

This is more than just the DNS, and it's important to understand that. The DNS itself -- most people think of that when they think of domain names -- is really a subset of what we're talking about here.

So the problem space in which we're working is the complete set of names that might exist in this tree-structured hierarchy, and the DNS is just a piece of that.

The next thing that's interesting to observe in this community is that the domain namespace and the DNS protocol that supports the DNS names that ICANN manages and delegates and allocates for use in this industry, those things are used in other places besides just the global public DNS, and that's an important consideration.

The reason why we have collisions and the reason why we have this issue is the fact that the domain names and the DNS has been so successful that it has been adapted and adopted for use

in other places by other people, and this is a good thing. This is a mark of success and an opportunity for innovation and interesting things to happen.

The last thing that is important to understand about the problem space in which we are considering this issue is that domain names as they are defined today by the IETF in use in the global DNS, they cannot be rigidly scoped in practice.

And what that means is, if I have a name and it simply exists, I really don't have enough information to know what to do about that name. So if you think about your browser, you know, browsers, you type some stuff in an open box, a text box of some sort. Most browsers, that's a combination opportunity for you to enter some words for searching or perhaps to enter some labels with a dot between them that look like a domain name and to then look that up in the DNS.

And the question is: It actually is, in the general case, not sufficient. Browsers have to make guesses about how to do that. So that's one example.

In your own local environment, you know, you have a DNS resolver in your local environment, and other applications and services on your phone and, you know, on your computer or on your laptop, I mean, you have the same kind of problem. You may or may not have enough information to know whether the

name that you've been presented, the label that looks like a domain name in the global public DNS, if that's really where it belongs.

And this, in fact, is the problem space in which we're working and that the ICANN community needs to be aware of and consider in its deliberations.

Next slide, please.

So in looking at this, we make the following observations about, you know, these circumstances and these facts as we see them that are before us, and that is that uncoordinated use of the domain namespace is really what creates problems. You get collisions from the fact that there's an uncoordinated use of names.

So the fact that you can have domain names and if you have an application that understands it's using domain names, you can also have local environments that use names that are intended to refer to things in the local environment. They're not intended to reference something in the global DNS.

And because of that, you get a collision of these names.

If you think about it historically in this recent application that -- application round that was opened in -- for gTLDs, for new gTLDs, there are currently a number of names that are currently

set aside. They've been deferred for right now while it is -- we figure out what to do with them, and those of course would have been corp, home, and mail. Applications that were submitted for -- you know, for names.

And that comes from the fact that we have this namespace being used in different places and there's ambiguity as to what to do with them.

So the lack of coordination among more than two groups in the namespace really creates instability. The fact that we have collisions creates ambiguity and that ambiguity is essentially a stability -- it's a security and stability issue for the Internet, and this is something that of course SSAC addresses directly, considers these kinds of significant issues, and creates advice for the community, and we identify these things for the community to consider as it develops its policies and processes to work around these kinds of things.

There clearly are at least two groups that have some influence over names and the presence or existence of names.

ICANN, of course, is the obvious one because it has a role as coordinator of the allocation and assignment of names that go into the root zone. That's ICANN's responsibility is making those decisions.

And it turns out the IETF is the obvious other example of an organization that also has a role. It creates a list of what it calls a special names reserved list of names that it wants to hold back and save and acknowledge for technical purposes. .LOCAL is an example of a name that's on that list. .ONION is an example of a name on that list. And the IETF has its own processes for putting names on that list, and ICANN, of course, which most of us hopefully are quite familiar with, has its own set of processes and policies that define what it allows to exist in the root zone or not.

And then there are other individuals and institutions -- there may be those that we don't know about. There certainly are at least private uses.

In fact, this is the problem that exists with corp, home, and mail. There are a lot of private uses of those particular labels throughout the Internet, and that's why they present a problem, because they collide with names that could be in the root zone, and ICANN, as a community, has to decide -- along with the board and the staff, we have to come to an agreement as to what to do about that and to deal with that instability.

So I think that's the last slide here. This represents where we are. We have created a definition and a statement of what we believe the problem space is, which we think is important, and

we have drawn our findings, what we understand about the facts that we see before us, and the next step at the moment is to develop a set of recommendations.

If you were in the SSAC public forum, it -- that we had in the last session, the last session block here at the ICANN meeting just now, we did actually make the statement that SSAC hopes to have a set of recommendations and a work product available to the community before the end of this quarter, Quarter 4 here, in 2017.

So any questions about that?

DAVID CONRAD: So I open the floor to the board members or audience members that would like to ask any questions on this particular topic.

STEVE CROCKER: Thank you, Jim. This is very helpful. Go ahead.

DAVID CONRAD: Peter, go ahead.

STEVE CROCKER: Yeah.

PETER KOCH: I'm Peter Koch of DENIC.

So Jim, you presented this view of the world that there are like domains and the domain namespace and that there's a separate responsibility for the domain name system within ICANN and another responsibility with the IETF as a matter of fact.

Is that -- who's view of the world is that?

JIM GALVIN: What we acknowledge is that ICANN has a responsibility for the names that go into the root zone, and we're simply observing that there are other groups who are taking advantage of the existence of this technology -- namely, the public DNS and its resolution protocol and the fact that you can have names -- and they have used these names elsewhere. We actually are not having -- don't have an opinion about their authority or responsibility. We simply acknowledge they exist and they are, you know, doing what they do, and we simply have to acknowledge that they exist, and ICANN has a responsibility to -- the ICANN community has a responsibility to acknowledge that and to react to that in some way.

PETER KOCH: Do you allow for another question?

So you said "we" again. Who is that "we"?

JIM GALVIN: Oh, the community, ICANN community, "we." I consider myself part of the community.

PETER KOCH: Okay. So I can say that's your personal view? What I'm trying to get at is that this particular issue is not necessarily uncontroversial and I would offer a dissenting approach to this.

What you call private use can also be called squatting, the same way that I can use your car without your consent and declare that as private use of your car.

There is an MoU between the IETF and ICANN and a clear separation of the responsibilities for the namespace, and the document with which the IETF or parts of the IETF believe they have the ability and the power to assign names by declaring them a protocol issue, I see -- I see a split here and would urge the ICANN board to approach the IETF to fulfill their responsibilities with regard to this memorandum of understanding, and then we can go from there. Thank you.

JIM GALVIN: Right. So thank you for that, Peter. I -- I, you know, will take that certainly as a comment for SSAC --

STEVE CROCKER: Peter?

JIM GALVIN: -- to consider as --

STEVE CROCKER: Peter?

JIM GALVIN: -- it's developing its recommendations.

STEVE CROCKER: Let's continue at least one more round.

PETER KOCH: Sure. At your service.

STEVE CROCKER: Thank you. But it's to -- mostly the other way around.

So I'm not 100% up to speed, but mostly, and I want to untangle a couple of things, and Jim covered it, but I want to -- to cover it again.

There's what goes into the root zone, and one can talk about the domain name system based upon what goes into the root zone, and the IETF has structured the namespace and they also can and, to a certain extent, do talk about the use of names in contexts other than DNS. And the issue that comes up in a practical sense is that names that are intended to be used outside of DNS show up in DNS anyway. "Local host" being an example and there are several others.

So the practical problem is, even though one, from a theoretical point of view, could say the namespace used for the domain name system and the namespace used for other purposes are completely separate, they -- they blend together. They bleed into each other. And rather than ignoring that and sort of not paying attention to the consequences, another approach is to say, "Well, let's pay attention to what the facts are and if there are names that are commonly showing up at the -- for root access, and of course getting a nonexistent domain response, but nonetheless, if they're showing up, let's take that as an objective fact about the way the world actually works," and then we have to decide what we're going to do in terms of whether

we're going to prohibit those names or we're going to have some other mitigation or whatever.

And my understanding is the IETF has not quite focused on a particular policy nor, in fact, is it really the IETF's natural -- I don't want to tell the IETF what to do or not to do, but they typically are -- shy away from policy issues.

And so as I understand it, there are factions that say, "Well, there's no reason why the IETF should have anything to say about what you're calling squatting. They just go and use the name." It's not quite taking a privately owned vehicle. It's like taking an unused piece of land, if you will, that -- whose -- that hasn't been allocated or owned by anybody yet and -- but -- and they go use that.

So .ONION, for example, is a reasonable example.

How do all those pieces fit together in your mind? How should they fit together?

PETER KOCH: So, yes, thank you for correcting the analogy. We can start from there. At that moment that somebody occupies that land -- and don't take "occupy" too aggressively here, please -- it is no longer available for anybody else. And this is the part where the coordination kicks in. But the responsibility should be -- the

responsibility should be crisp and clear here. This is not about ignoring the facts and ignoring the traffic that happens, same as forged I.P. addresses happen in the world and we see all this, to which nobody reacts in a way that is suggested here. Like, oh, then let's -- people declare their own addresses, come back, and then they get the address they have been sitting on for so long because we don't want to ignore the facts.

The point here is that the document in the IETF and the MOU are in clear conflict. And the conflict at least is that there is no coordination between the two bodies. And all this, there is a namespace and this looks like a domain name and it walks like a domain name and it tastes like a domain name, and probably it is a domain name. So the responsibilities need to be straightened out here. And I don't see this happening.

WARREN KUMARI: So if -- Sorry.

DAVID CONRAD: Warren and then Jonne.

WARREN KUMARI: So, yeah, if I can jump in. Just for completeness, I want folks to know that the IETF has actually been discussing this at length

and recently adopted a problem statement on special use names. We've gone through the process of allocating a name, .ONION. There seem to be consensus that this did not go quite as well as it could have, and so there has been an extended process on sort of adopting a special use names problem statement. And now we're going to try and get that finished and then hopefully move on to a set of solutions.

The document which Jim was talking about, the SSAC document, does discuss things like need for coordination. Presumably the IETF document will also mention something like that. So there is progress.

At the moment, the IETF is not doing any other special use names. The IESG sort of put that process on hold while we looked at it. So there is process happening.

And I can't remember who's next.

DAVID CONRAD: Jonne.

JONNE SOININEN: Yeah, thank you. So the -- kind of like, it might be good to kind of like divide these things. And I think Steve a little bit tried, and I think Jim said it quite nicely -- eloquently there, is that there

are three categories. There is the root, there's special use names in the IETF, and then there's squatting, if I'm blunt. Kind of like something that is -- private use can be that -- something that has leaked from somewhere else.

So the squatting -- or the private use thing, that is, of course, under nobody's control and, like Steve said, IETF doesn't impose any control on that.

On the special use names, I would try -- maybe not everybody knows, but these are not actually something that end up in the root. These are not something that are resolvable by the DNS. For instance, .LOCAL is solvable by something that is called multicast DNS, but it's not in the root itself.

.ONION, which Warren said is the latest one that has been allocated -- and there are very few that IETF have allocated over the years. I think .EXAMPLE, .TEST.

PETER KOCH: There's essentially no other than --

JONNE SOININEN: .LOCALHOST.

Excuse me?

PETER KOCH: There's essentially no other than .LOCAL.

JONNE SOININEN: And .ONION. PETER KOCH: The others were reserved anyways.

JONNE SOININEN: They were reserved before there was any policy.

But none of these that are reserved are actually resolvable by the DNS itself. So there -- like the name says, they're special use.

IETF actually has, like you know Peter, has -- or had a process or a policy how to actually allocate the special use names. And that was what -- under which .LOCAL and .ONION were reserved. And like Warren said, that was found inadequate. And now IETF is working on a better policy.

On the coordination, in the beginning, when IETF started to work on a policy on special use names, IETF actually sent a liaison statement to the ICANN board and to GNSO.

I agree with you that the coordination was probably -- has not been perfect. But on the other hand, there are people from both from the ICANN community and from the ICANN organization actually participating in that work. So I don't -- I think that it's actually -- there is at least some coordination. But looking at what SSAC is going to say most probably in their proposal is that

more -- this most probably needs improvement. And I can agree to that, that it needs more collaboration and coordination. But I don't see -- and I agree with you that there is most probably an issue there.

But I don't understand your point of, say -- that is there's something else that you're trying to allude towards than just that this needs more work together to make sure that we do this responsibly.

PETER KOCH: So without trying to monopolize the microphone or the spotlight, you said this is what comes in the root and, of course, this is also what then can no longer go in the root, right, like .ONION or .LOCAL or anything else on these lists. But in no way is this specific document restricted to doing things at the root level. Somebody could propose a so-called protocol element that would affect a certain second level domain name in any existing TLD. And then it could be declared a protocol element with the consequence that that particular domain name could no longer be resolved in the global domain name space. This hasn't happened yet. But if it can happen for the root, it can happen everywhere.

So, therefore, this carving out part of the namespace and declaring them protocol is not only important for the root, for

the TLDs, but for every level underneath. And there's no boundary in there. And I think this is a policy issue that needs to be looked at from the policy side, and it's not just the technical issue that can be dealt with.

DAVID CONRAD: Yeah, go ahead and respond.

JONNE SOININEN: I think we agree on that. And that's, I think, partly why IETF is working on it, to basically -- to solve this. And you're right, there needs to be dialogue between ICANN and IETF on that. I agree.

DAVID CONRAD: Warren.

WARREN KUMARI: Yeah, I think I was going to say basically the same thing as Jonne. The IETF is working on this, right?

PETER KOCH: The IETF is actually avoiding to work on that, but let's take that offline.

WARREN KUMARI: We adopted a document. The IETF also sent over a liaison statement saying please note there's this thing we should coordinate on. I brought it up a couple of times in the TEG group. The IETF is working on it, right? We've adopted a document. There is progress. It's not as fast as some would like, myself included. But I think we're moving along, so I'm a little confused by some of your statements.

PETER KOCH: Okay, thank you.

JONNE SOININEN: Just to point out, if I may -- sorry, David. But the ICANN community's also working on it. The SSAC is clearly working on it. So I think that there's work is ongoing. And like was said, there is room for improvement but at least we have a start.

DAVID CONRAD: Ron.

RON da SILVA: Thanks. This was a good dialogue so far. But my understanding is you're looking to study beyond just this specific reserve space between the IETF and ICANN but look more broadly at the namespace and where there are other collisions. There are

enterprises that more even -- consumer equipment suppliers that may inject things into namespace that looks and resembles a lot like the DNS.

And, you know, that -- I think what I'm hearing is this more broadly namespace is going to be looked at, which is good. Because, you know, you talked about addressing, just sort of very briefly. And it reminds me of a similar challenge there. There's enormous amount of legacy space -- it's the space that existed before the registries -- that were assigned to various folks in history.

And these addresses have never been routed in many cases. They're not used in the global Internet today. And, you know, some people might use them for internal use and maybe even create names to map to some of those addresses in a static way internally within an enterprise.

It's the same issue, right? You've got collision the moment one of those blocks is then sold. There's transfer markets that are happily happening all over the world. That's sold to another enterprise or service provider to be used. And when the routing is attempted, it fails suddenly because the space is being used in a few places and then it doesn't -- doesn't map to the global Internet.

So there's the coordination effort, but there's also, you know, getting an address or getting namespace that requires cooperation between operators or registries and registrars and folks that are using the namespace. And there's usually terms and agreements that are associated with those relationships.

So I think when you have that combination, there's no -- back to the number analogy, there's no guarantee that if you get a block of addresses that they will be routed until in turn you either arranged through peering agreements with other providers or you purchase services to have that routed. And then it becomes somebody else's problem to go make sure it's globally unique and globally routed in a way it's not colliding with the same kind of idea of collisions because somebody's squatting on it. That's the perfect term to use. So names, addresses have a similar challenge.

DAVID CONRAD: Kaveh and then Jim.

KAVEH RANJBAR: I think, Ron, I have a better example to compare this to I.P. addressing. A few years ago APNIC was assigned 111/8 which included 1111 and 1.2.3.4. And I think Geoff Huston did an

article on that because he receive about 500 megabytes of traffic any given time, and they decided to basically reserve the space.

So I think that's more close to what we see in root than legacy space being hijacked, for example.

JIM GALVIN: Thank you. So Jim Galvin again.

I want to scope a little bit, you know, what SSAC is going to say and the way that we're looking at this. I mean, the problem space is clearly a large one. The IETF is sort of the obvious other example of an organization that cares about domain names. But we're -- the way that we're approaching our recommendations and what we're going to recommend is to consider what is within ICANN's remit to do.

It's easy to suggest that, well, let's coordinate and, you know, that seems like sort of a natural recommendation to make. But you quickly run into interesting questions, like, who would you coordinate with and why? I mean, the IETF is sort of an obvious example. But, again, there are a lot of people who use domain names for their own purposes who are calling them private use or squatting use, whatever -- however you would like to characterize that.

So, clearly, it's not about coordinating all the time and with everyone. You're not going to be able to solve that problem in the general case.

So, you know, ICANN does need to consider what it can control and what it can do about the parts that it does control. So there are issues like what if somebody else pops up and has a list of names that they're using that also creates collisions and thus creates ambiguity? I mean, what ICANN cares about, what the ICANN community cares about, and ICANN as the organization and SSAC as an advisory committee to that is the instability that's created by the fact that there are other people using a technology that is well within reason, you know, and certainly a rational choice for them.

So we need to consider -- the ICANN community needs to consider how it wants to respond to the existence and presence of these other uses and these other lists that will pop up from time to time. They'll change over time. What does all of that mean? New organizations will come up. You know, they're all going to have their own processes for what they do. And ICANN simply needs to have a process for dealing with the fact that these things exist. And that's the issue that -- that's the direction from which SSAC is approaching this as we think about what specifically we want to recommend to the community and to

ICANN, the organization, and, of course, to the board more directly. Thank you.

DAVID CONRAD: You know, since SSAC is taking this issue on, it sounds like it would probably be best for us to wait to see what SSAC's input is on this and, you know, evaluate -- see if there's some input that the Technical Experts Group may want to provide based on the SSAC input.

The only other thing I would note is I believe RFC 2860, which is the MOU between the IETF and ICANN, actually states that IETF has the ability to declare protocol. But things beyond that are policy issues that are not addressed within the context of the MOU which implies that it's outside of the IETF, therefore of ICANN's. That does not necessarily mean that it is ICANN's which does add an additional flavor of complexity to this particular topic.

So with that, I will move on to the next agenda item which I have forgotten. If you could put the agenda back up.

Okay. Yeah, I think it's Howard. Yeah, there we are. Yep. So, Howard, if you would like to talk about the ETSI network function virtualization.

HOWARD BENN: Thank you. Can we have the slides up? And the next one, please.

Okay.

So ETSI, as some of you may know, is the standards organization that look after standards for the mobile community. They actually write standards for all of the fixed telecoms in the mobile world but were best known for the mobile which has been the most active area for the last few years.

Now, during the last ten years, we've slowly been transitioning from a world where mobile phones made phone calls to a world where more people access the Internet on mobile than access over any other medium combined. So we are at the stage where the 8 billion users -- the 8 billion registered cards, SIM cards, that are out there with about 6 billion users, we are seeing massive amounts of Internet connectivity.

So within the core networks within the mobile operators, there's been a lot of discussion about can we leverage the work that the Internet industry has done over the years, used data centers for controlling our communications rather than having proprietary hardware and proprietary software which we've had to date.

So within the ETSI NFV group, they've been working on this for the last few years. They've generated two phases of specification. They're working on their third phase at the moment. And there's a number of issues that popped up that I thought it would just be useful to help educate the community here on. Maybe some of the addressing I'll just touch on at the very end.

So the essence is that we share the compute, storage, and network facilities available in data centers today working closely with organizations like the IETF who are also doing work in this area.

Next slide.

So just some of the words that are used, I don't know how many people are familiar with these words -- and it's taken me a little while to catch up with them. So what we have is the entity manager, the EM, which looks after these virtual network functions. These are basically bits of software that are running that perform functions both about, through, and in software and maybe in hardware today. We share resources. The way this is all managed is via something called orchestration. And we have a VNF manager that manages the life cycles. These things can brought up, expanded, contracted, and taken away.

Next slide.

And what we've been doing is basically trying to map the excellent work that the Internet community have done over the years and see how that maps onto the models that are used within the mobile world, and there are several things that have popped up during those discussions.

Next slide.

So first one is reliability. It's kind of interesting on the perception of reliability. So I think we're now at the stage where we're still in the process of trying to work out what people are willing to live with. So if your mobile phone doesn't work due to coverage, then people obviously complain. But they're willing to live with that.

If the mobile phone has coverage and yet a phone call doesn't go through, that definitely is a big no-no at the moment, especially when you look at the number of emergency calls that are made on phones today.

Whereas, Internet-based services, I think although users would like to have really high reliability, they are willing to put up with some thing that maybe isn't 100% reliable.

I don't know exactly where these figures came from. I was given these figures. But most mobile networks, we talk about minutes of downtime during a year, not hours. So we can see that

perhaps some of the data provided here shows that the current Internet community isn't quite as reliable. Again, I can't validate this data at all. So next slide.

We also need to make sure that there is interoperability between these systems, and that's really is where the world of standards comes in to ensure that we have the protocols to allow different vendors to provide different parts of this infrastructure and yet they all works together in a reliable, seamless manner.

Next.

And one of the things that we've been doing at both ETSI and working with the GSM Association is to start looking at how do we start benchmarking some of these systems. How do we work out things like reliability. But also issues like latency. So if you provide a voice-based service that is encrypted, then latency is vitally important. We have to get very, very low latencies to get good voice quality, so we need to start benchmarking these services. Next slide.

And the really big one is security. So there is a great concern out there that if you move away from a mobile operator who has everything locked up in their own data centers with no access to the outside world to somewhere where you're in a data center that possibly has other publicly addressable systems running in

it, that there could be an opening up for cyber attacks, denial of service attacks, the whole range of different things that goes on today in the Internet. And so there is real concern from the operators about that happening. And so the ETSI NFV group have a security group that are looking at that and trying to propose some solutions. But, of course, we have to work together with the Internet industry on this. Next Slide.

And again, one of the other interesting areas is that in the world of voice communications today we have things like legal intercept, which is a requirement in most of the countries that we operate in. This is starting to filter through to the Internet community as well, so perhaps we can share some of our experiences here. Within ETSI we have a group called TC Cyber looking at all of the cybersecurity issues, including legal intercept, and how they are provided while maintaining privacy for the end user and security for the system. Next slide.

Migration is obviously another interesting area. So we're in a system where operators today really like to have no downtime whatsoever when they're upgrading their networks. We're starting back to work a lot closer with the open source community who are not used to such rigorous requirements. So we're working very closely with the OpenStack development teams at the moment to try and resolve some of these issues. Next slide.

And again, integration linked in with security. So I think that what we've seen is in the Internet environment there is -- the virtualization has been going on for quite some time, so today there are many services out there providing Internet services where you can send folks off one application from another application, where you can partition memory and storage, and you can make sure that the two applications can never see each other. We just need to make sure that that is really, really reinforced and we can guarantee security. You can only imagine what would happen if somebody could get into a mobile operator's network. So today I know some people experienced some issues with mobile roaming in India when they arrived last week. But one of the issues around roaming is that you can dial any phone call, any phone number anywhere in the world and get through to anybody more or less anywhere where they are. So from a cybersecurity viewpoint, that's obviously a great concern, if somebody could access that network, a vast amount of damage could be done in a very short amount of time. Next slide.

So within the standards side we are continuing to develop the standards in the ETSI NFV group. We're working very closely with the GSMA. So the GSMA is the organization where all of the mobile operators are members of, it's where all the roaming agreements take part. It's where a lot of the issues on security

and user management to handle and then ETSI write the standards on things like the NFV, on cybersecurity, whole range of different areas. Next slide.

So I think I'll make this last slide. There's a few more slides in the pack, for those who want to have a look. But just a quick mention on the NFV security working group. Again, anybody who wants to join this group, you cannot see because it's an ETSI ISG. So anybody can join. You have to fill in a little form. But you can -- basically anyone can join these -- this group.

So one of the things we've been trying to do is pull together the security experts from both the Internet world, from the mobile communications world, and then bring them to generate a whole set of standards. So working right the way down to the guys on -- in organizations like OpenStack and the open source community right the way through to the way that mobile handles security. So we're transitioning away from the use of sim cards. There's a whole set of work been going on so you can have downloadable credentials into a secure environment for authentication. And so we want to make sure that we can take the best of what we do today in the world of security for mobile. That is authentication. We know everybody who accesses a mobile network, we know that subscription details. We don't necessarily know who they are. I think this is very different from the Internet world where you have a very open connectivity.

And it would be interesting to see whether we could work together to pull these -- so we can get a securer Internet moving forward. Thank you.

DAVID CONRAD: Thank you, Howard. Okay. I open the floor to board members or members of the audience, if anyone has any questions for Howard. Yes, Kuo-Wei.

KUO-WEI WU: I like to make a comment about the real security we are facing in the futures. And I think I share with some idea with some of you. As IoT and home devices are getting popular, we have to remember those are home device and IoT device, the price getting cheaper and cheaper. And to be honest, I look at the manufacturing chair and the industry, how they're making home device and IoT. They don't spend pennies in their software at all. They just go to the Internet to get free software. So easily you can expect those home device and IoT definitely is the problem of the security come from. And particularly I have to say, in some of the country if you buy the hardware PC or Mac and they give you any kind of software for free, including the virus. And so I think that the people here, you definitely know the parts can be bought with a very limit money to file the DDoS.

So I think if we really want to thinking about how to resolve the security, we have to figure out how to make these manufacturing industry on the right way to maintain the stability and security of the whole overall Internet. So that's my personal comment for this.

DAVID CONRAD: Yes, Howard.

HOWARD BENN: The -- it's an interesting point because ETSI in particular has been working on security standards for IoT devices for a number of years. It's very difficult to make sure that all manufacturers comply with the guidance that is provided. And I think this is one of the -- the really, really big issues that we all face moving forward. And ETSI have a group called NGP that I talked about in previous TEG meetings. So one of the aspects that they're looking at is, what would the Internet look like if we started from scratch today. And one of the things that came out of that work is that, you would have to associate with the Internet. So you couldn't just randomly have devices without some sort of association. And therefore, devices that caused issues could be disassociated in a very secure manner. And it is something I think we have to face up to, is that this is kind of this line that we have to draw that's getting harder and harder between privacy,

security, and trying to stop some of these attacks. And just the latest one was actually on a Dynamic DNS, wasn't it? That was what caused the issue in the first place. So maybe that's another issue we need to talk about another day.

KUO-WEI WU: Can I respond, please?

DAVID CONRAD: Yes.

KUO-WEI WU: Actually the other day when the DYN was being attacked, you know your friend, John Klensin, wrote an email to me. Many years ago when IETF have a meeting in Taipei actually John Klensin work very hard to link the manufacturing with the IT people. But very pity it doesn't happen because I have to say many of those home devices, you know, is manufacturing in Taiwan and assemble -- well, manufacturing in China but (indiscernible) Taiwanese. So actually John proposal one idea, say is it possible to build a link or communication channel between IETF and the manufacturing peoples. So that is my comment.

UNIDENTIFIED SPEAKER: To the second part, because the first part is more IoT issues. Probably we could deal -- the work ETSI is doing, we can do this in Copenhagen. I want only saying about this brief format affecting people participate in the IETF meeting. My company, my company's World Wide Technologies, we participated so many people and I was so happy, budget for participating in IETF activities that I would say that this ideas of the manufacturers or vendors are brought there. But I have individually brought them, not as companies because the contribution in IETF feels mostly formally individual driven. Thank you.

DAVID CONRAD: Jonne, did you want to comment?

JONNE SOININEN: Not necessarily on this, more on the NFV but let's see what Jay has to say, first.

JAY DALEY: Thank you. Jay Daley. Actually thank you, Howard. It was a great talk. It's only a small question. Can you remind me on the intellectual property status of work in ETSI?

HOWARD BENN: I certainly can. So the ETSI IPR policy is based on FRAND, so fair and reasonable, licensing and non-discriminatory. And there has been an awful lot of discussion about how ETSI interacts with the open source community because many open source projects have a free IPR policy. So those discussions continue. I think it's what -- what's clear is that the open source community and the ETSI and 3GPP communities are working closer together and more and more projects are coming to fruition.

JAY DALEY: Okay. One more. On the OpenStack work you talked about, I was a little bit surprised by that because there's so many government services in my country, for example, that are -- now have OpenStack down the bottom level of them including electoral things, for example. So there's already that significant trust there. Now, I'm well aware that telecommunications people sometimes have a higher level of requirement there, but is that ETSI contributing code or is that ETSI trying to -- no. Okay.

HOWARD BENN: For open source -- for OpenStack, no. It is -- it is simply the -- the ETSI NFV group are informing the OpenStack community of the issues that they're seeing, how the -- they're obviously individual

people who are the same people contributing, but OpenStack still has the same IPR policy and I'm sure that will continue.

The only program that really is getting discussed at the moment is the Open MANO where some of the work is actually -- ETSI is looking at directly contributing into that. And I think that's the one that is causing most discussion at the ETSI board level at the moment.

DAVID CONRAD: Jonne.

JONNE SOININEN: Yeah, if I can add to that a little bit. So I think what Howard tried to say is that what ETSI does is they do specifications. Some of those specifications are aimed at giving guidance to, for instance, OpenStack or OPNFV. OPNFV is Open Platform for NFV which is an organization that is basically creating a framework for NFV and contributing to optional projects like OpenStack.

The contribution -- how my company that I work for when I'm not here and how Howard's and Francisco's companies also work is basically we contribute directly to OpenStack or Open NFV and usually use a lot of that guidance that has been agreed among the industry in ETSI. ETSI itself doesn't -- ETSI is a

standards organization and it is contribution-driven by its members. So ETSI itself doesn't do contribution.

What Howard referenced to is that there's a group called Open Source MANO, so Open Source Management and Orchestration, which is actually an open source project within ETSI. So something that ETSI runs. They kind of a full industry. This is again not an open source project that ETSI itself has started but members of ETSI have started within the ETSI context. I hope that helps.

What I would like to actually emphasize on Howard's presentation is that a little bit of the NFV story, or the network function virtualization, that there's a quite big transition in telecommunications now where some of the technologies that have been generally -- what have been used already for some time in the so-called IT world, so like OpenStack, cloud, and virtualization are taken in to use also in the telecom world. And moving away from specialized hardware and specialized network elements to more data center driven architecture with generic hardware and then software that is -- has a lot of open source components but also might be also proprietary, basically creating a virtualization platform on top of which these -- what used to be discrete network elements are run as either -- as virtual machines or as software basically.

UNIDETIFIED SPEAKER: K.S. RAJU: On second issue I want to ask you one other thing. Lots of telecos and cable operators, they do apply the refurbished routers which is unknown and then which I have seen in India lots of companies who are broadband cable operators, TV cable operators, they use a set of boxes to offer the Internet and all. They use the refurbished equipment.

And so one more thing is really affecting the cybersecurity issues in this region.

And one more thing is the biggest electronics recycling (indiscernible) is India and Asia-Pac region. Okay? Thank you.

DAVID CONRAD: Any other questions related to NFV? Okay. Then we will move on. Sorry, there is one online. Yes.

REMOTE INTERVENTION: Thank you, David. This is a question from Wolfgang Kleinwachter from University of Aarhus. He says, car manufacturers have to comply with internationally-accepted security standards. Why can this not be done for hardware and software manufacturers?

DAVID CONRAD: That is an interesting topic. I'm -- I imagine that organizations like ETSI would be able to come up with criteria and standards upon which such regulations could be made. But I don't think that -- I mean, Howard, do you want to touch that?

HOWARD BENN: Dangerous subject. I think it's -- it's a really interesting question.

So does a device have to be proven to be compliant to a set of standards before you can connect it to the Internet? And that's what we're talking about here. At the moment, no.

DAVID CONRAD: Yeah. Very true. Although, looking at the way denial of service is -- capacity is increasing, that may end up not being our choice in the future.

Yes, Steve.

STEVE CROCKER: I'd like to understand Wolfgang's assertion in a little more detail.

What is the -- what are the international standards that exist and that compliance is required? I'm not sure exactly what he's referring to.

DAVID CONRAD: I believe he was referring to within automotive standards. If you put a car on the road, it has to abide by certain requirements.

STEVE CROCKER: Ah. I missed that.

Well, the -- the automobile is a lot more advanced than the Internet. Sorry. The automotive -- the automobile is an Internet device, isn't it.

DAVID CONRAD: It's getting that way, yes. John?

JOHN LEVINE: Automobiles are fundamentally different because in most countries you need a license from the government to put -- to place your car on a public road, and it would be nice if we didn't have to end up there for the Internet.

DAVID CONRAD: It definitely would be nice.

Yes. Howard?

HOWARD BENN: Okay. So to place any electronic product onto the market in Europe, you need a CE mark, and a CE mark basically says you're compliant with all the standards you have to be compliant with. So every mobile phone has to prove compliance.

At the moment, none of those compliance documents are related to the way that you access the Internet.

DAVID CONRAD: Yeah. So I think we'll move on now.

[Laughter]

DAVID CONRAD: The next topic is DNSEXTLANG by John Levine.

JOHN LEVINE: Thank you, David, and I'm glad to see that according to the agenda, I am probably here. I presume that means that -- that you have determined about my mass, you cannot determine my location.

So this is -- this is a very different kind of operational issue.

If I could have the next slide, please.

The DNS data consists of records, and the records are of various types, and there are about -- between 70 and 80 types defined, of which perhaps four are in common use.

And there's been a long-standing question of why don't we have new record types, because when we invent new services or distribute new types of data over the -- or -- over the Internet, it frequently makes sense to coordinate them with different -- different record types.

I mean, for example, Paul Wouters has been quite active in DANE, which has defined new record types to publish SSL certificates and stuff like that.

So the reason it's hard is if you can look at my slide, we have this four-step process for getting your records, you know, from your brain onto the Internet, and the first part is that you somehow have to get the DNS records into a master file which defines the data to the Internet. And for most -- historically, people manually edited the file with a text editor, but these days you go to your registrar or to your DNS provider and they have some sort of Web-based thing that lets -- that lets you type in some amount of DNS data. And the Web-based stuff tends to be pretty bad, which is why we call it crudware.

The crudware then somehow creates its own files that are passed on to DNS servers which -- master servers which are -- that's software like BIND and NSD and PowerDNS.

That then puts the records on the -- on the -- on the public Internet. Then for an application to use it, coming up from the bottom, the applications have some sort of DNS libraries that allow you to request records, which then go up to DNS caches, which then retrieve the data from the masters. And this is the way the DNS has worked for a very long time.

Next slide, please.

Now, when you define a new record type, here's what happens now.

First, the IETF publishes an RFC that defines the -- the record type, and the implementation -- the implementation and the publication tend to overlap somewhat, okay?

So the first thing you have to do is you have to update the libraries to know about the new record type, which means that whoever maintains the library has to add the new record type, debug it, come up with a new distribution, send the distribution out, and then everybody who uses the libraries then has to update their software, which they may or may not do.

Caches fortunately don't need updating, so we're not going to talk about them anymore.

The master software also needs to be updated so it understands the new record type. That turns out not to be so much of a problem because the people who update the DNS servers are vigilant and tend to update them fairly quickly. But again, once they come up with a new version, they then distribute the new version of BIND or NSD or whatever, which people may or may not install. And the crudware is never updated.

So that typically, if you have some Web-based DNS -- DNS console, you know, you can -- you can use the same four record types now that you could use a decade ago.

So next slide, please.

So here's our goal, which is that when a new record type is defined, we want the -- we want these three pieces of software to be updated automatically so they can handle record types all at once.

Next, please.

Okay. So what -- what that means is that the -- the master -- the master servers and the library software, it has to understand the syntax of the new record, which will be -- well, there will be the

name of the -- the name of the new RR type and then a bunch of fields.

It has to understand the binary form and it has to be able to translate the text form into the binary form and back.

The master software and the library software has to be able to do that.

And if you actually want people to be able to find these things, since it's Web-based, you need some way that it can prompt people for the necessary fields and syntax and so forth.

So next slide, please.

So here's the idea. We come up with a language in which we can describe our record types. Initially, I said we would put them in text files. Paul Vixie had the brilliant idea of actually publishing the descriptions in the DNS itself, so when you come up with a new RR type, it published in the DNS and then the system can automatically find them, which I'll describe a little more in a minute.

And once that's done, you need to upgrade your software once to handle the extension language. Once you've done that, new record types come in automatically.

Next, please.

So here's a description of a couple of record types. The first one is a mail exchanger, which is a pretty familiar record type, and you can see it's an MX record. It's -- we're describing it as a mail exchanger, and it has some records. And then a text file also has some -- has some fields, so the text file also has -- is a record with some fields.

Next, please.

So in each description, the first line is -- for example, here's an SRV record, which is a relatively complicated one.

So it says the name is SRV. The type number is 33. The "I" means that -- DNS records have classes. This is only good for the Internet class.

Then server selection is a comment, which is intended to be used in -- to prompt the user.

And then the first field is priority, the second is weight, the third is port. Each of those is a 2-byte integer. And then there's a domain name, for the target -- which is the target.

And I've come up -- and I've come up with descriptions of pretty much all of the existing record types in this format.

Next, please.

It turns out that to handle nearly every record type, I came up with 14 types. There's three sizes of integers, there's text strings, there's domain names, there's v4 and v6 addresses, and then there's all sorts of other stuff. There's timestamps, there's 32- and 64-bit hashes. There's arbitrary hex fields. There's Base64 and there's a few others.

And then I also have an escape type called "Z," which says -- this is for a particular type that can't be described in -- reasonably in any other way. But there aren't very many "Z" types and they don't apply to records that are widely used, so it turns out to be much of a -- not to be much of a problem in practice.

Next, please.

In the descriptions of DNS types, there are options, and it turns out that the kinds of options I put in, there's three of them.

Like if you'll -- here's this -- is a description of the NSEC3 record that's used for -- in DNSSEC, and the first field, the hash algorithm, you can either define it as a number or you can put in a mnemonic.

So here it says that the -- in fact, the only algorithm that was defined initially was SHA-1, so this says, "Well, if the user types in SHA-1, that means 1." In the second field with the flags, there's only one flag opt-out. In fact, you can have multiple fields,

multiple values, separated by commas. Some fields have various types. For the salt, the fourth field, this happens to be a hex field which is stored in the record with a count, so the "C" says store it with a count.

And then there's a 32-bit hash.

And then the last field is the types, and there are some records that have -- these are types that -- these are types of records. In this case, it's the types of records that are stored at this particular name.

And for NSEC and NSEC3, there's actually a list of all the types.

So the "L" means that this is a list of types, not just a single type. And I won't go -- I mean, you can look at my draft to see -- to see the details of this, but what I'm trying to show you here is the field options are not really very complicated and you can generally -- you look at the RFC that defines it and see what the record types are and write a description like this in a couple of minutes.

Next, please.

Now, to update --

Actually, if you can just flip back to the previous slide. Thank you.

This description gives you enough information that the libraries and the master servers can parse and de- -- and de-parse the records. I mean, because it's -- it's the record type and then it says, in this case, a 1-bit -- a 1-byte integer, another 1-byte integer, a 2-byte integer, a counted hex field, a Base32 field, and a list of types, and -- and with this description, it can then -- that turns out to be enough so that application software that I'll describe in a minute can -- can parse the master files and -- and de-parse the binaries.

Go ahead, please, again.

Now, for the users, that's where the -- that's where the comments come in.

So if -- if we have a user who's going to be defining an MX record, the idea is that you'll click on, you know, "New Record" and it will say "What type: MX," and then it will show you a little form like the one that I've mocked up at the bottom of the screen here.

And it's taken priority and host name out of the description, and then the user has typed in 100 and the name of the server.

And since it's typed, it knows that the value in the priority field has to be an integer that will fit in 16 bits and the value of the host name has to be a domain name.

So that, I mean, the user has to know something about -- about what he wants to -- what he or she wants to do, you know, but this way you can -- you can -- it can prompt you pretty -- pretty strongly so what you get is at least syntactically correct.

Next, please.

So the final bit is getting the data from the DNS -- getting the data from the DNS, and Paul's idea was to -- was to publish the record descriptions in a fixed place in the DNS for which he proposes RRTYPE.ARPA for the number -- if you're looking up by number, and RRNAME.ARPA if you're looking up by name.

So here we have a -- a hypothetical foo record which is Type 999, so the description is at 999.RRTYPE.ARPA, and at FOO.RRNAME.ARPA and then the actual description is simply stored as an ordinary text record in the DNS. And it says RRTYPE equals 1, just so you know that this is really an RR type. EN says that the comments are in English. If you want to internationalize this, you can have different versions of it with different versions of the prompts in different -- in your local languages.

And then it's just -- and then there's just -- the strings are what I -- are what I showed you before, the description of the name of the record and the individual record types.

So -- and it was really easy to write software that parsed this out of the a text file and turns it into a zone file to publish. So once you've done this, then when we define a new record type, waving my hands like crazy, we arranged that the -- once the RFC is published, the description is -- is placed in the DNS and that any software that uses this stuff can -- can look it up.

Next, please.

Okay. This is not a total panacea for every possible new RR type you would want to define, and there's two reasons. One is that there's a few RR types that just have weird syntax that's hard to define. And in particular, there's a few where the order of the fields in the master field and the order of the fields in the binary record don't match. You know, so if you really need to do those, you know, you can write code to interpret one of my special "Z" record types, but in general, all those types were already handled by -- by servers and they are not types that users are likely to -- to want to -- to want to put into a realistic zone file. They're either -- they're -- they typically are obsolete like SEC, the predecessor to NSEC.

And the other is that some new RR types actually require the server to do something special. I mean, when we defined the latest version of DNSSEC, when you -- you know, when -- when

your cache does a DNSSEC lookup and it finds the NSEC record and the signature record and stuff, it has to do stuff with them.

So I can describe the syntax of the records but I can't tell the cache what to do with it. But again, this doesn't happen very often. I mean, we only have to invent DNSSEC once, and -- and changes that -- that require semantic -- new record types that require semantic changes like this happen maybe once a decade, so I'm not going to worry about them. This is definitely a 90/10 solution, so...

Next slide, please.

So having invented this, I started implementing it, and David Conrad very kindly arranged to get some support for implementation. So the draft of the specification is done. I've modified the perl DNS library so that it can -- it can read record types out of the files, it can read record types out of the DNS. It can automatically, on the fly, when it sees a record type with an unknown type name or a binary record with an unknown type number, it can actually go out and look on the -- look in -- look on the -- look in the DNS, find the type, fetch it in, compile it to new perl code, install it on the fly, and -- and then handle the record type. It's actually pretty -- it's pretty slick. And I'm currently talking to the people who maintain that DNS about how to best to integrate it into the standard distribution to the

library. I'm doing a proof of concept in Python to -- to show how the Web -- the Web thing would work. And these will all be given away for free as -- as open source.

So the hope here is that once we've done this, that adding new record types will be easier and people will be more -- more willing to do it.

I mean, we've had, you know, very few new record types defined, simply because there's been a perception that you can define a new record type but no one's going to use it because -- you know, because -- because the provisioning software can't handle it and -- and then there have been some unpleasant workarounds. In particular, a lot of new services have actually been done by reusing text records which in some cases has worked okay but in most cases has turned out to have bad side effects.

So this mostly works. I'm happy to give the software away to anybody who's interested in it and -- and I hope people use it. Steve.

STEVE CROCKER: Thank you. Very, very cool. Having worked through the deployment of DNSSEC and the problems of having new resource record types that it's uncertain when people would

accept them and so forth and the open set of problems that have not yet been solved as to whether or not there are additional things that are needed and whether or not to use text records or additional resource records of some sort, I'm -- I fully understand and appreciate the problem.

I jotted down a set of questions, one of which is the prototype, and you say you're working on that. That's good.

Looking ahead a bit, two kinds of questions about sort of success and failure, in a way.

I could imagine that a new record type is defined and that a whole bunch of uses of it are -- populate the DNS, and so then all of a sudden resolvers all over the world are faced with seeing this new record type and have to go into this fetch-and-then-reconfigure cycle. So that leads to two possible bottlenecks. One is on everybody fetching from the same place at the same time and the -- kind of a potential meltdown, if that's not anticipated, so simply putting it under .ARPA may be shortsighted if there is a big load that's going to be placed on that.

And then the other is, how long does it take a working resolver that has considerable load on it to absorb, reconfigure, and be able to respond to that? Does it -- does it have time? I mean,

these things are under relative stress for the high-frequency ones. So that's one set of questions.

And the other is, I can easily see why this is motivated by solving yesterday's problems. Do we know any things coming up where this actually has -- where there are new record types that are likely to be used, and with what frequency is that going to be exercised?

JOHN LEVINE: Well, in answer to the first question about performance, I haven't a clue. You know, it really depends on kind of the caching -- the caching strategy out of -- out of the leaves. You know, and if I have -- you know, if I -- if I have a share- -- you know, if I have a busy server and I have a shared library, it's going to make a big difference whether the shared library can fetch a record type and compile it for use on every process on the system or whether every type of process starts up and it's got to do it over again. But that's -- I think that's sort of a quality of implementation detail.

As far as actual record types to do this, I mean, there's a few new record types coming along now, like there's SMIMEA which is new, and there's -- which this could easily describe, you know, and people have asked -- said, "Well, you know, this will -- this works fine if the new record types use -- use fields like the fields

we used before," and it -- and having gone through and made an inventory of every record type anybody has ever attempted to define, in recent years by and large people have reused field types. There aren't -- I mean, there was a new one for EUA48 and EUA64 to put Mac addresses. But that was several years ago, and there hasn't been a fundamentally new field type since then. So I think for the kinds of record types I see, it seems to work pretty well.

It's also sort of a chicken-and-egg thing. If people know it's easier to get a record implemented if it uses an easy-to-describe field type, people might be more inclined to do that.

DAVID CONRAD: Jay.

PAUL WOUTERS: Just to comment on this, I think there was a little misunderstanding between the question and the answer. The resolvers actually do not have to do any new work because they're just doing wire-format DNS and so they get new R types and a number. And they just look up the number and give the binary data as an answer. So there is no additional work to be done under regular DNS queries.

The only additional work is when a human would want to add this new record type to the DNS zone that they own to put this through the provisioning software. So this doesn't produce any kind of discernible load on a server.

JOHN LEVINE: That's not quite true because all the way -- the application typically need to depart -- need to parse up the record so it can actually fetch the useful bits out of it. I mean, like if I -- in the unlikely event I wanted to write an application that used SMIMEA, I would need -- my application needs to know, like, here's the hash and here's the type and here's the data. So the application needs to know what the fields are. But, again, that's the sort of thing we should be able to compile once.

PAUL WOUTERS: Okay. So if you are doing that, that is actually really scary because then you are hooking up a word to the DNS that's actually part of the DNS. So that would actually be really scary.

JOHN LEVINE: Well, yeah but...

[Laughter]

DAVID CONRAD: Jay.

JAY DALEY: Thank you. Jay Daley.

It's nice, John. I know someone who once tried something very similar with sort of DNS schema and did a beautiful way of describing DNS in DNS schema and all the breakdown of all the different types of fields and things. And DNS schema gives you a bit more sort of depth to that, but, anyway, that didn't get anywhere. A couple of things about this.

Firstly, how are you going to internationalize the way that things are presented to the end user as expressed in the binary data here?

JOHN LEVINE: I'm trying to think. There are -- the individual records tend not -- the only thing that might need internationalization is a string field that's going to be presented to the user.

JAY DALEY: That's specific what I mean, yes.

JOHN LEVINE: Yeah, I mean, other than that, like -- you know -- I.P. addresses don't need to be internationalized. That's a good question. It's also -- it is a topic that nobody has really thought about. I mean -- I mean -- the strings in text records are eight-bit clean. You can store Unicode there. You can store UTF-8 there, if you want to. But as far as I can tell, nobody does.

So I think the answer would be if we figure out sort of at the IETF level that we want -- that we want to store non-ASCII text data in the DNS, then whatever we decide to do I need to figure out how to describe.

JAY DALEY: Yeah. As well as that, what I meant, though, is that currently those -- the names for things don't actually appear inside the DNS in any way. And, therefore, if somebody presents them to somebody, they're getting them from elsewhere or making a choice what language they use about it. If it's coming inside the DNS, then you need a different language version for each of those. You suddenly have to add another dimension to the level of data you're providing.

JOHN LEVINE: Oh, actually -- the slides are gone. But actually in the DNS version, there's a language tag in the record.

JAY DALEY: No, no, no. But there's extraordinary length of things that encode with that.

The other point I was going to make is that there's an analog here to EPP and what takes place in EPP that should be considered. EPP has a -- at the core of it a very fixed and defined data model.

And, again, I know somebody who once had a neat idea for suggesting that EPP should specify a -- have a mechanism by which new data -- instead of actually including the data, it describes the data that should be given to it. Because when people in different registries add new things in there such as a company number or something like that, that has to come through an extension.

JOHN LEVINE: Yep.

JAY DALEY: If added on. But if EPP were at a different level where it was more descriptive about that so contained a list of these fields in a standardized way --

JOHN LEVINE: Yeah.

JAY DALEY: -- that would be better. Just suggesting that those two things might actually do some benefit of bringing together because there is a time again when new records may need to be encoded inside EPP as well in order to be transferred between parties and so there's a little bit of a tie-up there.

JOHN LEVINE: Certainly. The concept sounds similar although I'm not sure how much commonality there would be in any of the implementation.

DAVID CONRAD: Wes.

WES HARDAKER: Thank you. Wes Hardaker, USC/ISI. A couple of comments. Cool idea. I love it. I have a couple pleas. First off, don't put the internationalization format in the record itself because there's a whole lot of them. So why don't you put it into a label so that when I want to query for English, I only get once response instead of a very, very, very large packet.

JOHN LEVINE: I thought about that. The problem is twofold. One is: How do you do a default? And you could sort of do it with stars, but it gets kind of ugly and nasty.

WES HARDAKER: I (indiscernible) stars, but --

JOHN LEVINE: Yeah. Well, beyond that, also, if you really want to do it right, the example is a two-letter language code. But, in fact, English is in one language, it's en dash many countries. And that's the sort of stuff that would be trivial in a real database and is hopeless in the DNS.

WES HARDAKER: Well, It would be hopeless in a very large packet coming back to you, too, to accommodate for all those.

JOHN LEVINE: Yeah, I know.

WES HARDAKER: Think about it.

JOHN LEVINE: No. In fact, an earlier version of this put the language tag into the name. And I moved it into the data just to make the lookups easier. If it turns out they are unwieldy, I can put it back in the name.

WES HARDAKER: Also, don't forget that some -- the more recent trend for display formats within DNS in particular is to move away from bits and actually start putting in individual words. If you, like, go look at DANE, for example, we actually updated that so that rather than having type codes of 0, 1, 2, and 3, we actually put in keywords of what those actually map to.

JOHN LEVINE: It does that.

WES HARDAKER: Okay. Good.

And then, finally, the most interesting thing is that you might have an interesting time with security ramifications when somebody spoofs a record to, say, a registrar and actually reverses fields or gets the user to put in completely the wrong stuff and ends up leaving their inserted data into their thing --

into their zone being questionable, if not a real security problem. So food for thought.

JOHN LEVINE: Yeah, you're definitely at the risk -- you're definitely at the mercy of the people who maintain the descriptions, you know? You have the same problem when you are updating libraries. It's just sort of at a lower -- slower.

WES HARDAKER: You're missing my point. If I can spoof the .ARPA data, then I can cause whatever application you're working with to actually cause you to put in potentially very different things.

JOHN LEVINE: Yeah.

WES HARDAKER: I could change the word to "password," for example.

DAVID CONRAD: That's why we have DNSSEC.

Steve?

STEVE CROCKER: Very good point, Wes.

I was just thinking about if you publish a description and then you need to edit it, either because you made a mistake or because there's a later update, it seems to me you're going to have to change the keyword in order to cause a trigger to update across the network. Otherwise, the old description will just be used forever.

JOHN LEVINE: I hadn't thought much about that, but I was -- I mean, it's been very rare that we've needed to update an RFC that described an RR type because it described it wrong. So I was hoping that if we could get people to apply a similar level of care, we're not -- we cannot screw up that way.

STEVE CROCKER: But this is the Internet.

JOHN LEVINE: Well, yeah.

[Laughter]

You can certainly imagine ways to do version tags or time-outs or something. But -- but I would hope -- I would like to avoid

solving that problem until I'm persuaded it actually needs to be solved because it makes things more complicated.

DAVID CONRAD: Jay.

JAY DALEY: Yeah, tell me to go away and read your I.D. if you want. But I'm really not clear why this has to be in DNS and what then the relationship is with discovering new RRs and TTLs and things like that. You know, in terms of this being an operational system, a live system for lookup in this way, I don't understand why it isn't sort of a static file that's published and people don't get it for every lookup.

I mean, are you expecting a piece of software to look at this every couple of hours to see if any new --

JOHN LEVINE: The implementation I have now goes and looks up -- whenever it sees a record type it doesn't have a description for, it goes and sees if it can find one, you know, and then it caches it locally.

JAY DALEY: Okay. Right. So it's a when you see something you don't know, then go looking for that.

JOHN LEVINE: Yeah.

JAY DALEY: Okay. Thanks.

DAVID CONRAD: Paul, did you want to comment? No? Okay. Thank you very much, John.

And now we move on to Warren Kumari to talk about work being done in the IETF. Take it away, Warren.

WARREN KUMARI: So, yeah. Hi. I'm Warren Kumari, and this is Paul Wouters. We are the two IAB-appointed reps to the TEG. And this is going to be a quick update on some work happening in the IETF. I'm going to try the clicky thing and see if it works. Whew, it does.

So this slide deck actually covered a few different things. I'm going to skip over the first presentation, and then depending on time we'll go back to it or we'll go to a different slide deck that also contains that stuff.

So signaling trust anchor knowledge in the DNS. So what's the problem this document is trying to solve? Well, soon we're going

to rolling the DNSSEC KSK. This is a good thing. If you want more info, there's a URL up there and some dates.

Unfortunately, the process for actually introducing the new key is an RFC called 5011 and some name servers don't support 5011. Either they were written before 5011 came out or they just decided not to implement it.

Most implementations do support 5011, but many of them have the 5011 support disabled. This is because when we first started introducing DNSSEC and doing all the presentations and DNSSEC workshops and things like that, we had a bunch of examples which included a configuration which said this is the root key, always believe this is the root key, don't bother trying to swap it. People who have just cut and pasted that config are going to use the old root key. It's not going to roll over.

So here's sort of a Venn diagram because Venn diagrams are cool, which says all DNSSEC resolvers, some of them support 5011 and some have them actually have it enabled. Unfortunately we have no way of measuring the size of any of these circles. That's not quite true. We know how many DNSSEC resolvers there are. We don't know how many of them do 5011, how many of them have it turned on.

So this is an extract from the KSK rollover plan. It largely says what I said, measuring this is really hard. We do now, however,

potentially have a document which will help us with that. It's draft "DNSOP trust management," I think, is the key tag management. I can't remember the exact name.

Basically what it says is this. Resolvers every now and then when they would normally do their RFC 5011 processing will send a query which encodes a list of the trust anchors that it knows about. So in this example, we have a KSK at the trust anchor called 1984. We are rolling to one which is called 4242. So initially the resolver sends queries just looking for ta-1984.

When the key roll starts happening, it starts sending queries that have 1984-4242. And once the key roll has actually completed, it will just be sending queries that contain ta-4242.

What this does is it allows somebody who's watching traffic at the root zone to have a look and see what percentage of users have the old key, what percentage of users have the old key and the new key, and then what percentage of users have just the new key.

The same information is also encoded in a different way and stuck in an EDNS option. It's basically the same thing, just a different way of encoding it. And the nice thing about this is you can tell before the key roll completes who is actually going to break and potentially who you should talk to about fixing it.

So, yay, has this actually fixed the problem? Unfortunately, no, not really. Deployments which came up before the RFC 5011 support are kind of by definition going to have come up before this document was published as well. This means that we still can't measure that percentage of users. Also, this document is currently going through the IETF. It will be published sometime soon, we hope.

It's actually finished working group last call so it should be relatively soon, but it's still going to be a while before people actually implement and resolve a code. And then once it's implemented, it's going to be a while before it actually gets deployed.

So, hopefully by the next time there's a KSK roll, however many years from now, we'll actually be able to get some more useful stats.

Questions? And sorry I went through this fairly quickly. We're trying to squeeze in the other slide deck.

DAVID CONRAD: Steve?

STEVE CROCKER: Two things, one of which is directly related and one of which that branches offer. One that is directly related is with respect to signaling what keys you have strikes me as comparable to signaling what algorithms. So I don't -- okay. So the fact that you're saying yes means that some coordination are thoughtfulness about what mechanism to use and how to do it.

Possibly time for another discussion. But it occurred to me that your comment about we don't know where all these resolvers are we don't own is similar to discussion we had a little while ago about devices that are on the network and we don't know what their status is with respect to security.

One could imagine trying to register devices, all devices, on the network. All that's a big, hairy thing. Possibly one might have a discussion about registering in some fashion or having a location of all the DNS resolvers on the net so that they could be contacted if there's a problem or they could meet certain standards or something. Just tossing a small pebble into the pond and standing back in case there's a splash.

WARREN KUMARI: So, yeah, we had at one point discussed possibly including things like resolver version or new algorithms or things like that. But we decided it would be better to get this published and then

have possibly a second document which describes a way to encode which algorithms you know.

Jay?

JAY DALEY: Yeah. Sorry if I'm talking too much. Tell me. It does seem to me there's a lot of things we don't know about resolvers. And somebody I think, looking at David, would be nice to do some work so that we understood what versions of what resolvers did some things. This is one of them. Parent centric or child centric is another very important one.

And perhaps if we've got more data about that, we can reboot attempts to fingerprint them. And we can at least also then know that surveys that we do, if we get statistically correct surveys done, we can then match that up to what we know about those things and that can then provide us some extrapolatable numbers.

DAVID CONRAD: That is actually an active area that my team is looking into. Paul Hoffman is doing some research on resolver implementations, and Roy is looking at DNS analytics to help inform an understanding of a demographics of resolvers.

JAY DALEY: Okay, great.

DAVID CONRAD: And we'll be publishing that.

JAY DALEY: So Fantastic. We are doing stuff on parent-centric and child-centric identification so I can probably help with that.

DAVID CONRAD: Okay, cool.

Ron?

RON da SILVA: Yeah, the last comment about trying to figure out some analytics of resolvers and making some decisions is good.

I'm wondering also what proactive communication steps are being made to try to reach subsets or collections of folks running different resolvers. I mean, how do you address that issue? I know it's a big gap, and it's unknown what's going to happen. What steps are being taken to at least do some proactive communications?

DAVID CONRAD: So in the KSK roll, we actually have a fairly elaborate communications plan. It's published on the ICANN website, /kskroll, I think, /#communications. But you can just scroll down the page. It's toward the bottom of the page.

One thought that we had that we're actually looking into right now is now that we have access to the mail root server query data is actually looking at the I.P -- the source I.P. addresses of the queries that we're receiving. And modulo, we will winnow out the trash that's the root server. And then doing reverse lookup of those DNS addresses or looking it up in the WHOIS to try to identify the ISPs that are running those or the networks that are running those resolvers and contact them and say, Oh, hey, by the way, something interesting is going to be happening in about a year. You might want to be aware of it.

And we're also looking at whether we can determine if the resolver in question is actually doing DNSSEC, which would obviously make them more interesting than just your everyday resolvers. But that's an ongoing area of research. Daniel.

DANIEL DARDAILLER: Just one question. Do you have any restriction on who can request the KSK of the resolver?

(Off microphone.)

WARREN KUMARI: So, I mean, this is actually the resolver advertising it to the root by doing a query for a name that looks like that. So that query is a fully-qualified string. And it will go up and hit the trust anchor point. So it will hit the root -- it will show up at the root servers. That's where it will show up only.

DANIEL DARDAILLER: Because in case of the key leakage it looks like advertising I've got the wrong key.

DAVID CONRAD: And Jaap.

JAAP AKKERHUIS: Jaap Akkerhuis, NLnet Labs. I heard last week that Geoff and Joel actually claim to have the map of 95% of the resolvers which I'm active in this field, so you might want to check out what they've been doing.

DAVID CONRAD: Sorry, who has that?

JAAP AKKERHUIS: Geoff Huston and Joel.

DAVID CONRAD: Okay. I talk to them occasionally. Any other questions on this topic? If not, I believe we have more.

WARREN KUMARI: Yes. I guess if we can change to the other slide deck. So this was a -- no the other, other slide deck.

So this was originally a presentation which was supposed to be about half an hour. But I've got like 15 minutes to try and do it in, so I'll see if I can squish it in. So I'm going to go through it really quickly. Please tell me if I'm going way too fast.

So DNSSEC provides authentication of both positive and negative answers. So a positive answer is something like you look up www.example.com and you get back 19216811 and a signature that proves that that's correct. Less well-known is that it also provides authentication of negative answers. So if you look up something like login.example.com, if that name doesn't exist, you get back an answer from the DNSSEC telling you it doesn't exist and you also get a signature proving that.

Generating signatures is a fairly expensive operation in terms of CPU, so DNSSEC tries to avoid doing that wherever possible.

And one of the clever tricks that it does is something called NSEC which is short for next secure. And what that does is it takes all of the names that do exist in the zone, it sorts them alphabetically, and then it signs all of the spaces between them. That means that it doesn't need to know what query somebody might look up and it doesn't need to assign answers on the fly.

So that's kind of a confusing thing. I have an example to demonstrate it. So here is a look-up for .BELKIN. I chose this because it's a very common string which gets seen at the root and it's a TLD that doesn't exist. So here's the look-up for .BELKIN. I get back a response which says, NXDOMAIN, basically the domain does not exist, and I also get further down an NSEC record. And the NSEC record says, there is nothing that existing between .BEER and .BENTLEY. And then further down there's a bunch of cryptogoop which proves that that's true. So now my resolver can have a look. It sees that Belkin is between .BEER and .BENTLEY and it knows that that doesn't exist because it's got a signature that proves it.

So that's all really interesting but why is that useful?

So this document in the IETF draft, IETF DNSOP aggressive-NSEC says that recursive resolvers can use the information in the NSEC record to synthesize answers. Currently, if the resolver had got a look-up for .BELIEVE, for example, even though that

falls between .BEER and .BENTLEY, it would still do another look-up specifically for .BELIEVE. It would do a look-up, would have to send it to the root, the root would have to send back the answer, et cetera. This document says, don't bother doing that. If you already have an NSEC record that proves the name does not exist, just use that and reply with it immediately.

So this does a bunch of nice things. Let me check time. It improves user privacy because it means that names that users look up which don't exist don't leak out to the Internet. It decreases latency. The resolver can reply immediately. It also performs performance because the resolver doesn't need to send off queries. It also has a nice other feature which it improves DDoS resilience. Currently there is a bunch of DDoS attacks, where users look up -- or sorry, attackers use up -- look up a bunch of names that don't exist, they ask their recursive server, the recursive server goes off and asks the authoritative server, and if you do this enough times, the authoritative server gets overloaded. By having the recursive resolver just answer straight from cache, there's no additional query, the authoritative server never sees these queries. And apologies again for going through this really quickly.

So is this actually useful? Here's an example from May 12, which was a Friday afternoon because things always go boom on a Friday afternoon, where Collin and Kaveh from RIPE sent me a

question saying, Google is suddenly accepting K-root a bunch more junk queries. These junk queries are sort of a random string and then something that kind of looks like an IP address but it's not formatted correctly. Please stop this. It's making us annoyed. So when they actually contacted me, it was kind of around noon UTC. Not sure how well people can see the graph, but it was just when the number of queries had started going up. I work at Google, and so we started looking around in Google Public DNS to see what's causing this. Was there potentially a bug, had somebody made any changes to the code, or what's going on? Potentially are we being used as a DoS reflector, and are people sending us these queries and then we're forwarding them on and that's what's causing it? And more worryingly, why does this look like organic growth. A DoS attack usually sort of starts up at four eight, goes across and then turns off. This was more worrying because it looked like it was growing and might just continue to grow.

After some more looking around, we discovered it's not just Google Public DNS that's sending these. A lot of resolvers are. So shoo, at least it's not us, but what is causing it, and can we make it stop?

A little bit more looking around, and we discovered that there was a new worm which was spreading on the Internet and it was infecting access points and sort of home routers made by a

company called Ubiquity. And as part of their attack, they would infect a machine or infect an access point and then would do a look-up for a specific type of string to see if it could reach the Internet. And the string happened to be one that looked like that. Random string and then a random set of octets. So now at least we know it wasn't just our fault. But let's see if we can do anything about it.

So how am I doing for time? This is a graph of queries from Google Public DNS to the B-root server, which is operated by USC/ISI, Wes down there. And I'm not sure if people can see the letter -- the numbers, but over on the far left before the attack started Google was sending around 500 queries per second to B-root. When the attack started, that's the big spike going up, it spiked up to around 2,500 and Google Public DNS had this software already built in. We just hadn't enabled it. So we turned it on at 100% at the top foremost affected locations, and you can see the big drop down there. As I said, it was a Friday. We avoid making production changes on a Friday. So we waited until Monday. And then we turned it on on half of the machines in all of the locations that we have. That's the next dropdown. We let it bake for a week. And then finally, towards the right, we turned it on at 100% at all locations. And then you can see towards the far right, the number of queries which were sent to

B-root now is more like 30 or 40 queries per second. So, you know, about a 10X drop or so.

What does the document actually say? Largely what I said at the beginning. If you have an NSEC record which proves that a domain does not exist, don't even bother looking it up. Just reply with that. Also, if you have a wildcard record which covers this, don't bother looking it up. Just use that information and return the answer immediately. And that's about it.

So just sort of as a summary, currently the root gets around 60% of queries which result in an NXDOMAIN, the domain does not exist answer. These are sort of bogus, junk queries. If everybody were to do this, the bogus queries hitting the root would be more sort of like 1% or so. And sorry I had to rush through that. Hopefully it was vaguely coherent. Questions?

DAVID CONRAD: Anyone have any questions for Warren? I'll throw you a question that -- the NSEC 3.

WARREN KUMARI: Yep. So this does NSEC 3 as well. It does not work with NSEC 3 with opt out because you can't actually do that. But for NSEC 3 -- NSEC 3 works almost identically to NSEC. It's just instead of sorting names that exist in it, so you sort all the hashes that exist

in the zone and you just look -- check the name you looked up, if the hash matches. Basically the same thing, just hash things first.

DAVID CONRAD: Ram, did you --

RAM MOHAN: Warren, I'm channeling some of my board colleagues who say that the level of tech in this is so deep that they're a little lost. So perhaps, you know, a couple levels up summary of what -- what this means and what the problem is might be helpful.

WARREN KUMARI: Sure, yeah. Apologies. I did rush through that way faster than it actually deserved. The really high-level summary is, if deployed this cuts down the number of junk queries to the root and to other domains as well. It increases user privacy. It increases performance. It decreases the number of look-ups that end up hitting authoritative servers. I think that that's largely the summary. And I'm happy to go through it much slower if people want more details.

RAM MOHAN: Thank you, Warren.

WES HARDAKER: Hello, I am B-root from USC/ISI, and I just want to say, thank you. You caused my pager to stop going off those days.

JAY DALEY: Ram, I think a take-away from the board from this and the previous one is that resolver development has not had a lot of sunlight on it for a number of years, and if it had more sunlight on it in a more structured way from the industry, then there are a set of problems that could either be resolved or safeguards put in place to make solving them -- other new problems later easier to do.

RAM MOHAN: Thank you. And I'd encourage some of my board colleagues who are channeling this to feel free to speak up for yourself rather than channeling through me.

DAVID CONRAD: Any other questions? Yes, John.

JOHN LEVINE: Do you know where it's been implemented so far?

WARREN KUMARI: I know some of the places. So Google Public DNS does this. It has also been implemented in Unbound which is sort of one of the two standard big recursive platforms. One of them being Unbound, the other one being ISC's BIND.

DAVID CONRAD: Okay. Well, thank you very much. If there are no other questions, then we move into the any other business. Is there anything anyone on the TEG or in the audience would like to raise? Looks like Yoneya has one.

YOSHIRO YONEYA: This is Yoshiro Yoneya. So during the DNSSEC workshop there was a question about how to deploy BCP38 that filter out the spoofed source portal risk queries. And that kind of spoofed queries or spoofed packets used for those attacks. So that the deploying BCP38 is very important to decreased such kind of those attacks. So I think here is a good place to talk about it because the operational practices is how to explain the IETF but the operation -- so the operators groups is also important and ICANN is also important place to think about it.

DAVID CONRAD: So I know SSAC has published a couple of documents, I believe, on the value of doing something like BCP38. It's -- you know, has

been mentioned that it's probably worthwhile for SSAC to reiterate the value of BCP38. But it's sort of not directly a topic that the TEG itself would tend to focus on. Ram, did you have something?

RAM MOHAN: Thank you. I wanted to provide some -- so I'm going to put my board hat on and take the technical hat off and provide some feedback to the TEG, which is kind of a weird thing since I'm technical, right? But it feels like there might be a couple of things -- in our next iteration there might be a couple of things that we should consider doing to make the -- to make this an even more of a dialogue and a discussion. One suggestion is, that as we get the agenda and we get the topics that we -- we have some sort of a high-level summary, executive summary that -- that explains what is the issue, why is this important, and why should you care. Because I think -- I think that is a core -- a core thing that is missing. Because for us on the technical side, you know, you read the topic and you understand why you should care. But if you're not, then the -- I feel like sometimes some of these topics, the way we're doing it, it is perfect for somebody non-technical to say yeah, that's a techie thing. Let those guys go and work on it. So that's one piece of feedback.

The second is that in the agenda gathering phase of this, I think it might be useful to invite input on the board side from the -- especially from the non-technical folks, to get a sense of what types of topics they might want to be -- they might be interested in. I think that may be something that is useful.

The last thing is, what it points out to me is there's a crying -- there's a strong need for some kind of consistent tutorial type sessions that perhaps we video and make available, not only for that individual session but actually downstream so it can become, you know, some sort of a repository of information and kind of an onboarding. Not only for the board but really for others in the community where these are important topics. And often, you know, I hear folks from the community come and say oh, ICANN, you just do policy. But we're doing tech here, and I worry that even the level of tech we're doing is somewhat inaccessible to many of the people who attend this meeting.

DAVID CONRAD: Warren, go ahead.

WARREN KUMARI: So thank you, that's really useful feedback, and it's kind of similar to something I was going to say. The purpose of the TLG, which is a sort of a subset of the TEG, is to connect the board

with technical resources, or it's worded something like that. So I think we'd very much welcome questions from the board on what the board would be interested or like to know more about. And sort of your thing on tutorials, a short while ago BCP38 was mentioned. Is that something that the board would like a sort of quick briefing on what it is and what it's about, would that be useful, or are there technical things that the board would like more information on, you know, updates, mini tutorials. They could even -- they wouldn't initially have to be here in a great big meeting room but sort of places, things you want to know more about that we could go off and research and try and provide in an easily digestible form.

DAVID CONRAD: So I think it's clear that there is interest in that level of tutorial. One of the things that we have done is a -- the how it works series that were provided initially to -- to the community aiming at newcomers. We have internally been discussing, you know, perhaps extending that in various different ways. I think if there is interest at the board level of having, you know, directed tutorials on specific topics, I know my team would be overjoyed to do that. And I'm sure there are numerous resources in the TEG and TLG that would be able to provide that as well.

With regards to the agenda, I have -- I have been struggling to find the best way of obtaining agenda items for the TEG. I've tried a bunch -- several different approaches, asking board members directly, asking the TEG directly, asking -- just coming up with stuff, you know, out of, you know, bodily orifices. So far none of those really worked ideally. So always looking for additional input and specifically what sort of things the board would be most interested in because, you know, this is a group that's specifically designed to provide input to you. You know, techies will talk to each other, you know, in various places, usually inappropriately. So definitely looking forward to additional input. And I see a number of hands. So I guess I'll start with Ram.

RAM MOHAN: David, a quick example. A few weeks ago, there was widely reported in all the press an attack on network infrastructure and, you know, from -- from the board, there were several questions on -- not on what's already reported but questions on what's it mean. You know, how do we pay attention to it.

So it's -- it's kind of interpretation and analysis type of stuff that seems -- there seems to be a need for it.

DAVID CONRAD: Warren?

WARREN KUMARI: So yeah, obviously board members are really busy and taking up sort of two hours of their time with something that's not providing them value is not a good use of that, so please, would really appreciate any feedback if this is too technical, completely off topic, you know, what would be useful instead, et cetera.

DAVID CONRAD: Maarten?

MAARTEN BOTTERMAN: Yes. Thank you. I came here in my innocence. This is the first time I joined this session. But I came here because this relates very closely to our mission, and that's why I would like to use more.

And in the beginning, I thought, "Hey, I sort of can connect to what you're doing," but if this is really also intended to inform people like me, yes, please, let's have first a tutorial, because I agree that some of the issues can be simpler then with a little introduction.

And second, please try to aim the presentations at the level that you think, well, maybe people who have good interest and a bit of insight can benefit from it. I would really appreciate that.

Thank you very much for what you're trying to do here.

DAVID CONRAD: Steve?

STEVE CROCKER: Thank you. I agree with all the comments about the adjustments, but I want to make a point that that's within the context that I think this engagement, even -- even with the comments made about this, actually has been quite valuable.

This provides a fairly different exposure for the board to the technical issues that are coming along and plays a very big role in raising awareness and sensitivity to these issues, even if we can't quite all keep up with the details.

So I'm very happy and want to make sure that it's not just criticism or negative comments that are passed along, but I think this is a quite valuable process that we have here. Obviously, it can be tuned and will evolve over time, but I'm pretty happy with it as a baseline.

DAVID CONRAD: And then Patrik?

PATRIK FALTSTROM: Thank you very much. Patrik Faltstrom, TEG member from SSAC.

So Ram, I just want to have a clarification. You were more asking for a problem statement before the presentation more than having the -- to -- to a deep technical level?

RAM MOHAN: Yeah, that's exactly right, Patrik. It's not that it's too technical and that we should be less technical. It's much more of begin with what are we -- you know, why are we even raising this, why do you think this is important, and then dig into the technical level, so it provides some context.

DAVID CONRAD: And Cherine?

CHERINE CHALABY: So I enjoyed very much this session, and particularly, the first and the -- and the last topic. I find them very useful from a contextual point of view.

I think we -- what is unclear to me when we say that the TEG is meeting with the board, I think you are probably meeting with a subset of the board, those that have an interest in the topic or those that can understand the topic.

So if we want to make it a bigger involvement of the board so everybody can get involved, I think we will have to do one of two things. Either send some material in advance prepping people on the key issues and maybe raise the level of discussion a touch higher.

So this is a job for us, Steve, to really make it clear to the TEG what do we want to achieve from that in terms of that level of interaction. To me, it's not -- it's not very clear at this point yet. Thank you.

STEVE CROCKER: Yeah. So with respect to how much of the board is involved, the basic approach -- I'll take responsibility -- that we took, and particularly in interaction with David, was the technical experts group comes and expects to interact with the board, so there has to be some board interaction.

On the other hand, the board has, as you and I both know full well, a fairly heavy schedule and we have not made it a formal

requirement that every board member show up. We've not scheduled it as the only thing the board does.

So the practical situation, which is exactly what we have here, is that there is a substantial showing of the board, and in any case, that fits broadly with the -- what we try to do in the board anyway, which is not have everybody do everything, but we break up into committees and working groups and so forth.

So this is a kind of de facto ad hoc version of that, self-selected.

I did the count, including Goran who was here but went away. I think we had 10 -- if I recall correctly the count, 10 of the board. The full board is 20, including liaisons, so that's half. And it's the better half, obviously, because we're the ones who selected to show up here.

[Laughter]

STEVE CROCKER: I say that facetiously, but the self-selection actually has a positive effect in there.

So I'm not uncomfortable with the breadth. I mean, there's several of us who can keep up with this, but more important are the next set of people like yourself who don't come from a

technical background but have a pretty good feel for it and -- and several other people around the table.

We can always adjust the process over time, and so -- but I'm -- as I said before, I'm pretty happy with the level of engagement that we have and the effect that it's having, and of course it could be tuned up.

DAVID CONRAD: Warren?

CHERINE CHALABY: Can I quickly respond?

DAVID CONRAD: Sorry.

CHERINE CHALABY: Thanks, Steve, for clarifying. I think it's important to manage expectations, and I think you've made that very clearly, and it would be interesting to see from the TEG that -- what they're feeling about this interaction with the board at this level. That would be good feedback as well. Thank you.

DAVID CONRAD: Warren?

WARREN KUMARI: And one very last quick thing. I know everyone's busy, but if you have time, please provide feedback to David or Barbara or someone on how we could make this better and more useful to you -- you know, which were useful, which weren't, et cetera -- and we'll try and make them more useful in the future.

CHERINE CHALABY: Immediate feedback. This was very useful and very, very helpful. Thank you.

DAVID CONRAD: So now we're actually eight minutes into the reception for thanking the community regarding the IANA transition, and I would also like to remind everyone that there is a cocktail reception at the Casbah at the Westin. We have two buses, one that's leaving in I guess about five minutes, seven minutes, something like that, and then a second shuttle that's leaving at 7:15. The reception at the Casbah at the Westin starts at 7:30 and goes to 9:30, and it is -- does have alcohol, so -- yes. Wait. Sorry.

(Off microphone.)

7:00 and -- okay. That's interesting because that's not what my calendar says, but okay. So 7:00 and 7:30. So two shuttles, 7:00 and 7:30. Please, hope to see you there, and if not, I'll drink all the stuff that you don't.

[END OF TRANSCRIPTION]