20140123_StrategyPanelIdentifierWebinar2_ID848738

Alice Jansen: Greetings. My name is Alice Jansen. It's a pleasure for me to welcome you to the ICANN Strategic Panel on Identified Technology Innovations Webinar Session Two. Before we begin I'd like to briefly remind all participants of housekeeping items. This webinar is being recorded. If you have any objections you may disconnect at this time. This session is being streamed by the Adobe Connect room. If you intend to voice comments or questions during the Q&A, please join the Adigo bridge. Your lines are currently muted and will be muted throughout the presentation. You may submit questions or comments via the Adobe Connect chat call during the presentation. At the end of the presentation, you will be given the opportunity to voice your comments and questions during the Q&A. All lines will be unmuted for the Q&A. If you are on the bridge, please remember to mute your computer speakers once the floor is open to avoid echo. If you wish to speak during the Q&A, please raise your hand in the Adobe Connect room to be added to the queue. Should you not be speaking, please mute your line using star-six. Dial star-seven to unmute. The slides, recording, and transcript will be made available following the session. You may find housekeeping indications in the Adobe Connect notes. With that we will turn to Paul Mockapetris, our panel chair.

Paul Mockapetris: Thank you very much, Alice. This is the second version of this panel. I'm Paul Mockapetris. I'll be talking -- if the last session is any guide -- a little over a half an hour following a set of slides. After that we'll have open discussion. A lot of the panel members are already here. Next slide, Alice, please?

The formal charter shown here, one of the things we're supposed to do is develop a roadmap of what we think is going to happen to DNS and other identifiers. Note that it is an identifier technology innovation panel. That was chosen on purpose rather than DNS because there are other systems that exist now and in the future and DNS is only one of them. Develop some recommendations and engage in the ICANN community which is one of the things we're trying to do here today. Next slide, please?

Who follows up after the panel? The panel is going to have a draft report at the end of this month and I think one of the things I want to emphasize is version one is we hope to catalyze or incent somebody to do something other than ICANN and second choice is let ICANN itself chose something. ICANN is perhaps best at engaging the DNS, it's certainly something that's seen a lot of development and the ITS. There's a lot of people willing to give us advice and so forth, things, recommendations we come up with will probably only happen at ICANN because they're close to the root in some sense. Of course ICANN does separately and jointly sponsor things. I just want to emphasize that we hope we can develop some ideas that people outside of ICANN will think are good ideas and we can move forward in some sort of partnership. Next slide, please?

So, the strategy we adopted was to try to look for forcing trends, opportunities, and burning needs. And identify some things that haven't been perhaps pointed out as forcibly as they should've been in the past. That involves avoiding well-plowed fields like DNS activity or existing strategies for collisions, for example. Mostly when we have something new to say we recorded the fact that we looked at a particular issue but we're

not going to try and be as all-encompassing as hundreds or thousands of other people who are working at various aspects of the DNS. Next slide, please?

Our deliverables in the short timeframe which is a shorthand for once upon a time we were supposed to have a year to do this and it's been compressed a couple of other times. Document all the issues being looked at. Work on interesting subsets. Create a roadmap for what we predict over the next ten years, especially ICANN relevant, and make some actionable recommendations. I should point out that whatever we do is actually going to be input in the strategic planning process that ICANN has in place that's getting underway. Next slide, please?

I'd say we, the people here on the panel, some of them are online now, we don't necessarily agree on everything we discuss. I think that's fair to say. Someone may disagree with that. We've met four times with various subsets. I don't think the panel has ever been in the same place at the same time yet. Next slide, please? Thank you.

Just to start us off, what are ICANN current identifiers? They're listed here. Domain names, AS numbers, B4, B6. ICANN as well as the default inside the IANA working with these things. There's also a time zone database. So, starting from here, next slide, please?

The subsets we're going to talk about today are what we've discussed and what we think are drivers for the roadmap going forward. Some information about new identifiers that may be coming along, discussion about hardening the root and a discussion about DNS fundamentals. These are again another subset and it should take a half hour parameter from the start of the discussion we're having here. Next slide.

Talking about fundamental drivers and so forth, I think the panel agrees that there's some bizarre Darwinian struggle that's gone on in the DNS and will go on in the future. And basically the group needs some idea chair under the expansion and contraction although some of the ideas cut both ways. The favorite kind of expansion of the DNS, there's things like organic growth and momentum we have in the legacy base. We are involved to some sense in every computing element that's connected to the internet. So, there's the universality that I think is unmatched by competing systems. We might disagree on what competing system are but things like alabaster for example don't quite have the reach although they may have different properties.

So, one of the things favoring the expansion is just the momentum we have in the current system, like the new TLDs. There's lots of people who have different opinions about the relevance of all 1400 of them. I personally believe some of them are above and below average. But we're certainly going to try and establish a number of new ideas. Nobody should be completely indifferent to the power of marketing to move stuff forward. We'll just have to see what happens there. There's new capabilities that are being put into DNS every day that may or may not cause expansion.

New data types are another example of things that if you put more data into the DNS it should be more useful to more people. We'll talk about that as a separate issue. I think issues that favor contraction, there's the usual interface on browsers and the absence you load and the people who run those search bars would like to control them for purposes of commerce or just because they want to do something that's closer to artificial intelligence from the standpoint of understanding or having various comments that maybe if every search is triggered by voice input, people are probably going to be pronouncing domain names but that will probably be part of the input.

One of the things that was suggested to us by other people was that we should be thinking about developing 2D barcode components so that domain names could be embedded in 2D barcodes rather than contextual formats. In any rate as we transition to portable

devices and so forth, it makes it less obvious what the role of domain names will be in that time and situation.

The second thing for contraction is protocol and process ossification, Ossification is often applied to the internet in general that says it's hard to change things we have. New systems just have to find new ways of doing things. We'll talk a little bit about maybe some ways to deossify or to try and hit the ball on that but it's certainly the case that a bunch of people have found it difficult to get improvements and new things into the DNS. So, that tends to act against expansion or even force people to go elsewhere to get the facilities they need.

There's commercial identifiers, things like your Facebook ID. One the other day said -- I don't understand why everybody has to type Gmail.com at the end of their email address if it's all the same. There's certainly commerce and a financial interest in owning identifiers. Maybe your Facebook ID will become the new lingo Franco.

Lastly there's some new systems coming from the research world. Those ultimately replaced certain -- or whatever, provide new identifiers, certainly new web address for example, be viewed as the first version of all of this. Next slide, please?

One of the things that's interesting is naming technology itself is kind of white hot in the research world. There's name-based networking, content based networking, information based networking. There's a whole variety of different approaches. I think they all agree that naming as a basic technology to replace addresses or other location based stuff, you just state the name of what you want, not where it is. That's a context that's very hot in the research world, about the only thing the research world can agree on is what to name their own field.

Common themes in the NR access by name and opportunistic cashing. Some of us take a look at things like there was once upon a time a presentation of sending video through DNS and certainly a lot of parallels here. In the research world it's usually access and opportunistic caching. Some of the other things going on there is self-certifying names that don't have any particular need to be registered in the hierarchy in order to be useful. Very often these also have a public key infrastructure and it's allowing you to find the self-certifying name. Other people argue that the public key infrastructure is inherently for example something that makes privacy unlikely because it's just the thing the NSA or somebody else wants to collect to start invading your privacy or cataloging everybody in the network. So, whether or not that should be a component can be debated.

Lastly, there's user friendly names. Some of the systems that are out there in particular you can find a lot of people who argue that the big mistake in the DNS was having names that are user friendly. But I think there are a lot of people who like the fact that they can type the names and the definition of user friendly can vary. A lot of the time URL names are certainly not numeric but they are short. So, what exactly user friendliness is is up for grabs and certainly some of the contraction where people have the any bar as opposed to the URL box has to do with being able to type search terms. Next slide, please?

The reason I think that question is relevant to ICANN is ICANN serves the internet community by administering part of the DNS. The question actually occurs to them what is it they can do in order to make the system more useful to that community, one of the limitations you might try to remove. The present DNS is limited by several varieties of mask. There's more listed here. The most obvious ones are the 576 byte MTU that we inherited and the software that's in my DSL box that hasn't been updated in a decade. It does peculiar things to DNS queries. There's a lot of last mile difficulties experienced by the current DNS which some people say can only be solved by moving all DNS traffic into HTTPS because that's the only guaranteed path you have through this thicket because people will want to do commerce over the internet.

The second kind of thing is the protocol. It's proven hard to define new formats. The original specification approved on the order of 6400 different data types and I think we've gotten on the order of 64 defined. That doesn't quite meet expectations. People who are doing email verification has prototyped everything but putting things in text records and then they went to get a formal record and said -- Why bother? We'll keep it in text records. Putting everything in text records has a certain kind of operationally difficulty as people who are using sub labels if you will in domain names to separate the text information for different types. Whether or not that's a good idea is perhaps something we should institutionalize. I think the protocol has shown difficulty with getting new formats easily adopted across the internet. Again there are people who say -- Gee, other than address records you're going to have a harder time getting anything new through just because of the operationally difficulties ahead of it.

Lastly, it's not all protocol related. Some of it is process related. There are some of us who view it as peculiar that the DNS ops working group use things called mechanisms, specifically so there won't be protocol, to avoid part of the ITF process. There are many different working groups in the ITF that are doing ITF stuff but the overall architecture seems to not be part of that process. Maybe that's right or maybe that's wrong but it does seem to be something that might need a looking at.

So, I guess that boils down to try and do anything to help with this, breakthrough the bottlenecks, DNS, perhaps figure out the evolution to the next research direction or new features like privacy and confidentiality. ITF is about to have a boss as I understand that -- in London, dealing with some of these issues. It's just something that wasn't conceived of early on and I think might be partially inspired by Snowden. But it's there. On to more practical issues. Next slide, please.

One of the things that came to us unbidden in this panel was risks which was not the language of parenthesis that was used. This is a tool used in the AI world. But this is the separation protocol that separates locaters from identifiers. The technical thing that it does is it allows you for example to have an address that gets encapsulated in another one perhaps, sent to the end point, and then encapsulated so you can do things like have an identifier that can move around the internet and keep its locator which looks constant.

So, it's a V4 or V6 space kind of thing. And ICANN was offered the opportunity to be the root of the mapping database. But as it turns out, all things addresses are politically owned by the IRR. There's also practical questions about what software we should use. We had both proprietary and sort of open software offered to us. If ICANN's going to use new identifiers, one of the steps we would recommend about qualifying the political approach so that the address mapping of V6 and V4 are owned by the IRR, one of the things we can do is map to our ID tags. Who does that? So, there's -- I think one of the other questions is with a lot of the new technical capabilities, they might be formidable in the textbook space but they prove even harder problems for people in the political and government space.

I think another issue that came up was should ICANN publish more in the DNS? When you take a look at the new GTLDs there's a list of reserve labels that could be updated in an emergency. Should those things be available online? There's a group of people who are creating a repository for that. I mean, is ICANN going to have a requirement that you avoid certain labels? Particularly labels that are not in the database and are in multiple different foreign alphabets or worldwide alphabets. You're referring to a paper publication. It's unlikely that a system administrators will be able to understand how to generate the appropriate binary for five or six different character sets including Chinese, Arabic, and the others as well. It's just an error prone thing. Should those things but published, it would be appropriate to publish them in the DNS so you could check immediately to see whether or not a label would be prescribed or had become prescribed.

Other people have noted that certain kinds of information like domain birthdays for example, one in particular domain was created or allocated, they're particularly relevant to their reputation and they're not private information. Could those things be published online in some format? If so, presumably you'd want them to be accessed immediately and maybe the DNS would be the right thing. So, should some of this tabular information be published in the DNS? Or some other method? Okay. Next slide, please?

Obviously the root is one of the concerns of ICANN. We talked a great deal about the possibilities here. I think there were two possibilities we looked at, one was hardening the process of generating the root file. The second was hardening the process of distributing the root file. Another issue that we talked about, you have particular ideas was the analysis that takes place on the root. Certainly there's lots of possibilities there. When generating the root file, one of the issues was whether or not there are more robust or secure hardware. Certainly saying they've done a lot to make people thing about whether or not we have to worry about new attack vectors and there are vendors who would like to sell you more secure hardware components. Should they be a priority? We didn't end up thinking so, perhaps because we didn't know enough. But it seemed to people that hardening the root generation process might be a good idea but there's actually other opportunities there to globalize zone control. We'll talk about that more in detail.

From the standpoint of distributing the root file, there's a lot of people, and this goes somewhat in line with the research direction, that says that what you want to do is rather than worry about the 13 different addresses of the root servers and whether or not you have a path to a root server, you should distribute as many copies of the root zone as you can and not just root servers. There's sort of a specific example proposed by Bill Bixby which basically talks about distributing copies that are assigned and have two particular addresses and anyone who wants to can join those roots. There's calls on the line if people have specific questions or see what's going on at the end of my presentation here, we can go into it in more details.

There's other things. One is if you distribute a root zone copy every time, you don't need to have the addresses and servers. It's a free standing kind of thing. It needs server information perhaps for the glue but it needs to be digitally signed. There's some fine points about this. Some people said why can't you deliver sort of a sign master file? The delegation information there to be signed, some people think although it can be later verified, that it should be signed in advance and exactly how you want to do this, there's lots of different opportunities out there. I think this is one of those things that people talked a lot about, not just for the root but also for globalization. The example people come up with is once upon a time CA.gov for the state of California was redelegated by the government people on the basis of the fact that there were two computers out of hundreds of thousands in the state of California that had pornography. Shocking. But the point there was that when CA.gov got deallocated, a lot of California internet stopped working. That was partially because they didn't have predelegation and the correct set of servers set up so that they could be independent of what was going on in the root for any activity that stayed strictly within CA.gov. You can think of various recommendations that can be used for strictly operational things to teach people to protect the use of principles. Next slide, please?

So, the idea behind shared control is we imagine a way to have sort of a workflow language that allows people different rights in the same zone. So in the case of the root, could we split control to avoid the whole single authority issue? We'll have to figure out how to put it back together again. The thing people often might say is -- Okay, let's let every TLD owner be able to change their server information without having to go through an elaborate collaboration process. For other zones, since this is not just -- this idea is not just applicable to the root -- we can think about coordinating DNS, coordinating forward and reverse domains in some sort of standardized way. There's history here already. Some people have proposed the idea of having multiple signatures on the root and furthermore I guess if you had at least three of them you could even vote.

And so if ten signatures say the data is good, it's good. I guess if you had three different sets of people signing data, they could vote. There's two different DNS proposals about restricting this just to coordinating DNS signing information which to some of us seemed like a pity. If you're going to go through the trouble of finding a mechanism, then some of them is a little bit more generic, it might be useful in a much larger set of things. Next slide, please?

Let's talk a little bit more about the ideas in the roots. Can we go to slide 15, please? Let's talk about what happens today in the root. Today in the root, ccTLD and GTLDs, they request changes to their part of the root data. ICANN changes viable business objects, in other words programs, as well as humans. It asks the Department of Commerce for an okay and then sends the certified request off to VeriSign and VeriSign does some additional checks and sends it out. There's already this three party process in place today. Next slide, please?

So, one idea -- this is just one version. You can imagine many different versions of this as there are different governments around the world. CcTLD get to create a request for change in a journal that's viewable by all other ccTLD. Other ccTLD vote yes or no and if you get a majority then the change takes place two hours after you get the majority. If there's no protect, presumably you have a way to say yes or no, then it happens 24 hours later regardless. You do this by contextual algorithm that's applied to request the domain individually. Of course you'd have to figure out what you want in the way of rules. But this kind of technology is not all that complicated and could be implemented to get away from the idea of priorities that can freeze or control the process. Next slide, please?

Some of the implications here, distributing generation of root zones. I think before bit coin became popular some people would've argued that's impossible. I think today it's clear that it is. There's already a main coin for example. That would change the rules for more than just changing the data. There's yet another one of the perhaps research directions.

Which keys would you use? Single? Multiple? What protocols? I've been talking about journaling the data and moving it and so forth. You can imagine doing this in DNS or PPT or some other basic language. I think there's also a good question about what parameters do you want? I think you probably want to be able to have some kind of vote. You probably want to have some delay in the system so you can allow humans in the loop for at least some things. There's been discussion there about which small changes should be made unilaterally, which ones not?

I think we'd like to come up with a technical set of parameters that would allow workflow to be defined for the process. I think it would be delightful if we could do this, particularly because I can see the policy people who would have to figure out the exact workflow would be tied in knots for a long time. So, I think being able to do this kind of thing and get the issue off the table for one control for the whole system would be a very healthy thing for us to do in the DNS. Of course you're going to have to argue if whether the complexity is actually worth it and how do you debug it? I think one of the other questions is how to plug the global root zone -- I'm sorry, the system gets compromised and stripped, how do you ever restart it? Next slide, please?

I think one of the questions is can we break the ossification. If so, which parts need bought? We kind of ran through all the different parts of the database. The abstract database in other words the domain and IRR structure and matching zone rules, what could you do there to make the system more useful? I think one of the things, if IRR types would be useful, being able to have metadata in the DNS is an idea that's been hanging around a long time. Either that or just the idea of using sub labels in place of archetypes might be a way to do it. But let's try to figure out how to break that logjam.

Query and other operations, there's an effort underway by VeriSign labs, Paul Hoffman has done some presentations about making the DNS IAPs synchronous as ECU views DNS SEC. These ideas are out there. We tried to put a little more information and will follow up on that but there's no reason why we have to reinvent the wheel if other people have done it. There's some issues about replication. I think at the end of the day there's a question about how to make the effort worth doing. I'm sure there's a lot of people who would like to sit around and think deep thoughts about the next generation of DNS and DNS2 but you have to figure out how it is that you're going to make the game worth the effort and it would seem V6 and IRR exchange takes a long time or even DNS SEC. But on the other hand the momentum you get from being the legacy might be useful. Next slide, please?

To wrap up, a few of the recommendations that are under consideration here and we can debate are putting more signed data out there for reserved data or reserved labels or other things under the IANA parameters, a study to define a way to categorize the ossification and perhaps break through. Prototypes of the open root publications and one of the shared zone control in other words new methods of replication or shared zone control. I think another thing a lot of the people on the panel agreed to and even talking about the collision problem is it would be good to run some test exercises or a fire drill ahead of the need.

That's what I had to say today. I see there's already a lot of discussion in a very small font that I couldn't read while I was reading the slides. I think we should open it up for discussion.

Alice Jansen: Thank you very much, Paul. We will not open the floor and ask that the lines be unmuted. We will refer to the hand space in the Adobe Connect room to create a queue. If you're not speaking, please mute your line to avoid echo. You may do so by pressing star-six. To unmute, press star-seven. If you're in the queue, please make sure that your computer speakers are muted before speaking to avoid echo. Thank you for your cooperation. Operator, please unmute the lines?

Operator: Listen-only mode is now off.

Alice Jansen: Thank you.

Paul Mockapetris: Alice, perhaps we should go through the chat room queue first?

Alice Jansen: Okay. That sounds great. The first question is from Donna. What is the benefit of distributing copies of the root zone to any center that wants one?

Paul Mockapetris: The general idea here is if you have your own copy of the root zone, when anonymous decides it's going to try and attach all the root server systems in the world, you can just sit back and watch it happen and not worry about it because you've got your own copy of that data. The design of the DNS from the very start was there's a way to do certain zone transfer function so you can plan ahead and distribute copies of things. Then there's the ability to query and get it in real-time. Obviously if you have a copy of the root zone or any other zone that you might need. You don't have to sorry about being able to access it over the network perhaps during a denial of service attack or some other kind of event. If you already have the data in your local system, you don't have to worry about denial of service on the network.

I think one of the things in the very early set of attacks on the root server system, I happened to be in California and then flew to Sweden and the people in California were under the impression that the Swedish root servers were down. I guess in Sweden they were under the impression that the servers were taken down and a different type of access types to the servers were down. You have to think in terms of access to the pipes as well

as the data. At any rate, if you have the data, somebody has to take it away from you actively. That's why.

Alice Jansen: Thank you, Paul. The next question is from Olivier Crepin-Leblond. Why does the US Department of Commerce need to give us the okay?

Paul Mockapetris: It's just the process that's been defined in the past. ICANN is currently supervised or observed in some sense by the Department of Commerce and there's various agreements between ICANN and the Department of Commerce about the way that ICANN will proceed. Certainly globalizing that supervision has been one of the political threads that's out there. If what you want to do is have somebody in charge of the root or some coalition in charge of the root, I personally think it would be better to globalize it rather than have the DSC there but that's a historical artifact of the way it's been done.

Alice Jansen: Thank you. We have an exchange between Bill and Geoff in the Adobe Connect chat. Geoff, would you like to say a couple of words?

Geoff Huston: Yes. I can suppose briefly summarize this issue. Before we had signing in the DNS, this whole issue of how do I know the data I've retrieved is good data, valid data? The way the DNS was so structured, what you did is you put inside the DNS where to retrieve the data from. The theory was if I got it from this nominated server, it must be good. And if I got it from somewhere else, I don't know. Part of the thought process about the root in any other zone is once you sign that data and validate the data, the location where you retrieve the data from is not that important. In other words it doesn't matter where you get this root zone from as long as you can validate its contents with the root key, that's trustable data. Really the issue swings around to distributing the root key, the public key is extremely important. But the root servers need not necessarily have a distinguished role in a world where everyone validates.

This discussion is around --Well, what do we do about this? Do we have enough tools that will allow anyone to set up a root zone server? I guess the answer is if you are willing to intrude yourselves into one of the in cache clouds, you certainly can. If you want to use one of the cached copies, you certainly can. But could we go a bit further than that and remove this reliance on a particular set of root servers and instead rely on the key itself as being one of the central instruments of saying -- I don't care where I got it from as long as I can validate this is authentic, current data, then I can trust it. That seems to be what we're discussing at this point. It's not that we will do X or we will recommend Y at this point. This is really just considering the possibilities of how can we leverage more recent things in the DNS such as the increased use of DNS and obtain potentially some degree of operational hardening of the root as a consequence. Thank you.

Alice Jansen: Thank you very much. We're not moving on to Ivan's comments. I have a question about dot-co. The subject of the webinar and the title itself is identifier technology. But the only thing I'm hearing about is next generation. Why is there no discussion of alternatives to DNS which has defined networks as QOR and why this limitation to this discussion of innovation?

Paul Mockapetris: Hi. This is Paul. I think I said that the research world for example has ideas about replacements. The question is how we build the bridge to that if we believe that's what's happening next. There's certainly things like that out there and various of these flat identifiers that are out there. The question I guess is that when I think of the roadmap of what's going to happen in the next decade, pretty much a lot of the identifier stuff is going to be the bets have already been put down and people have already done the work about what we're going to see in the next few years. The only question is what's on the last five years of that timeframe. I think that being able to have more distributed control over the name space, DNS or something completely new, there are people who think that handles for example can replace all of this.

From the standpoint of the ICANN mandate, what should they do? So, I'm willing to listen to your suggestions. I don't necessarily see software defined networks as an alternative in the identifier space or for that matter but name coin and a bunch of these other research things I think are alternatives.

Unidentified Speaker: Also I don't think there's any attempt to limit discussion of other possibilities. These are the ones we've discussed in our panel but the point of meetings like this one is to open this up and make sure if there are good ideas out there that anyone can suggest that they get suggested.

Geoff Huston: Another thing I think is that these are -- with regards to the recommendations and things we're discussing here, these are things where we could be more specific. When you're talking about replacing the whole thing and moving to a whole new system, I'm afraid we've got more nebulous ideas.

Alice Jansen: Thank you, Paul. The only other question I see in the Adobe Connect chat room right now is from Olivier Crepin-Leblond. QR codes. Question mark.

Paul Mockapetris: Various barcode representations in the DNS are out there. There's a whole numbering space that we've managed at least in part by EBT global and the identifier tag and internet of things space that I think are going to evolve and certainly there's other encodings that are out there. I'm just not sure exactly what we would recommend ICANN would do in this space.

Geoff Huston: Let me follow up from that? The question is the DNS itself uses ASCI code as a human interpretable representation of that data. And in some ways the QR code system is an alternative form of encoding of the underlying data. The question is do the QR codes themselves represent an alternative identity space? In my mind I must admit in looking at QR codes, I look at that as an encoding vehicle as distinct from a truly distinct identity space. But others of course might have different views on that. But in looking at the sort of whether identities and identifiers, there's a distinction to made in that long-term vision as to the means of encoding those kinds of identities and the actual identity itself. Sometimes like QR codes, it's kind of the encoding that looks like an identity space but it isn't as distinct from other forms where ASCI representations have more mapping and human interpretable mapping to the underlying identity. Thanks.

Alice Jansen: Paul, I see Olivier's hand inside the Adobe Connect room.

Olivier Crepin-Leblond: Thanks very much. I thought I'd expand on this a little bit. Actually the QR codes question mark wasn't a question in itself but rather to expand on the data question. But you touched a moment ago on the internet of things and I recall not so long ago -- I'm afraid I haven't followed up with it, but I recall an organization in France trying to put together a rival root with regards to the internet of things and a different naming system, et cetera, but effectively thinking the internet of things will be overlaid over the other internet using different numbering and naming systems. Has your committee been thinking about the integration of the internet of things with the current systems? And also would the resource allocator side of things -- I haven't got the exact language for it, I just thought of it right now, but effectively when you've got http: or FTP: or other things, at the moment we're really looking 99% of the time at HTTP. Have you thought about other possible rival types of naming that could work on this?

Paul Mockapetris: I think there's a couple of questions in there. Let me try and answer one. How about the internet of things. I happen to be in Paris and been here for a little over four years now. And there's a project that I was actually involved in called the Wings project. I don't know if that's the project you're talking about with regards to the RFID tags? Are you talking about ETT global?

Olivier Crepin-Leblond: I think it was ETT global. It was led at the time by GS1. It's vague in my memory.

Paul Mockapetris: Yes. Okay. I'll just give you a little travelogue of my voyages through RFID tag land. I got a call one day from some people at MIT who said -- Hey, we were interested in trying to figure out how to unify all these RFID tag systems because there's multiple competing number spaces out there. To simplify it a little bit, the number five in the US Military's RFID tagging system might be an M15 where in the RFID tag that people were prototyping for consumer products, it might be a box of Kleenex while in the one that's being done for the airplane industry it might be landing gear. They sort of said, what we should do is try to unify all of these.

Obviously the thing to do was they defined this 96 bit tag space which would have as the first few bits an identity code that would allow you in legacy systems and another huge numbering space out there for future expansion, outside this 96 bit space. Furthermore, what they said is you could allocate this side of the 96 bit space with the same ideas we use for allocating address system, sub netting in IPD4 and D6. At any rate they said we'd need to have a system where we'd have identifiers and we could distribute control and so forth and so on. Is there any reason why we can't consider the DNS? I started thinking. That's a great idea. We worked a little bit on the idea and they came up with a scheme. This is the MIT auto ID lab. That eventually got handed off to ETT Global whose headquarter I think are in Paris. They're the people who originally did the barcodes. And in the process of doing the standardization, they decided to exchange six levels of hierarchy because it was argued that the sub netting is sort of variable sub netting scheme was too complicated and could never be made to work despite the fact that some people including myself were saying it works in practice on the internet. How can you say it won't work?

At any rate it got resolved into a more constrained system. Some people think that was just a commercially based decision. Be that as it may, that system, we're embedding RFID for basically putting an RFID tree in the DNS that's created. There's a research project to follow up on that called Wings. They identified the fact that -- and it was not strictly a French project. There was a German component to it and I think a couple of other countries. It was a big project with several different countries in it. One of the things they did was they though you couldn't use DNS for onus because of the issue of the single control of the root. And they wouldn't listen to ideas about what you could do is you could have trust anchors for different country parts and indeed the numbering space and RFID tags didn't have a great deal of country-based RFID embedded in it or not totally. So, they went off and designed their own systems. Now, there are other people who claim this whole approach is wrong and are doing other things for organizing RFID data.

I think one of the more -- the two things I think illustrate some of the political problems here is the way you got barcodes is North Americans came up with a ten digit system, people from Europe came along and said we'd like a chunk of that space, will you give it to us? And the people in the US said -- Yes. Here's the bill. They said -- No thanks. We don't want to pay you for numbering. Instead of paying you for numbering, we'll create a larger space for ourselves and encapsulate you inside of it. People from the airplane industry who wanted to get a chunk of that space in order to register airplane parts went to ETT Global and said we'd like to pay this. They said, no thanks. We'll create our own numbering space. One of the attracting and repulsive forces here is that you'd like to be part of a coordinated space for some reason but as an alternative to paying you create your own. It could be interesting to see how it all works out. The history suggests it will take awhile for people to figure out the allocation. I'm sorry to go on so long. At any rate, there's a lot of history here and if you like I can point you at it.

Alice Jansen: Paul, Anne-Marie would like to add something.

Anne-Marie Eklund Löwinder: Yes. I just wanted to add to that that we had an experiment in Sweden together with GS1 and we actually made a practical demonstration to share all the possibilities with the federated onus service and by using DNS to refer to services connected to EPC codes and from what I remember the whole -- it was sort of very clear indication of the difficulties we had, not only by choosing such different cultures which the GS1 and the onus represented together with the more open minded and not so formal DNS part of it. To be honest I think that demonstration of the directory service with onus and the extended packaging, we just handed it over to GS1 and they could do whatever they liked with it. As far as I know it never went further than that actually.

Paul Mockapetris: The internet of things has been a very popular research topic. I think when I talk to -- I think there's still some difficulties regarding the hardware implementation and the intellectual property rules that were going on and EPC global were peculiar because people said our intellectual property is in the hardware, the software should be free and open. But the hardware should have intellectual property. Where software people said -- no, no. The software is the crown jewel. The hardware is a commodity. The compromised result in EPC Global was that each working group if you will for those of you who are familiar with the ITF, each working group was its own private closed society for intellectual property and they had a separate committee I believe it was that in any case one committee wanted to pass something to another it had to be vetted for the right intellectual property controls before it could go from one group to another. So, I don't know. The commercial world there is very strange. The other thing is that the hardware hasn't always worked as well as people would like. There's not just one RFID code and transponder type and standard. Some of them are powered, some are not. Some of them are supposed to operate over miles so that you can read all the tags off the container ship while it's still out at port. There's an awful lot of complexity there and there's a lot of intellectual property. Anne-Marie? It looks like she's having a little bit -- lost her phone line? Is there anyone else who wants to put anything in the chat room? We're caught up?

Alice Jansen: In the meantime, Olivier has his hand up in the Adobe Connect room.

Olivier Crepin-Leblond: Thank you very much. I've got another question. I'm sorry for raising -- not really sorry for raising the question. It's a thing that's been puzzling me forever, since I've been dealing with DNS actually. We -- the record which is indicated by IN and which I know is also in some cases can be chaos is the only other one in the record class. Is there any -- some people say -- You know, you change the record class and have another DNS that can rival the DNS and the IANA DNS. Is there any worth to these validations?

Paul Mockapetris: There was a bunch of -- in the original specifications there was a bunch of things that were there for future use and expansion. The original plot behind the class system was that you might have exactly parallel name spaces controlled by different parties and this whole distribution of control thing was in there from the start and class was thought of as one of the ways that might be useful for doing it. Similarly, if you had asked me back then, the idea that DNS queries would always have exactly one thing, one key in the query, I would've said no. We'd have new kinds of things and selectors and so forth and newer naming systems. At any rate I think the track record is that the internet class is the only one that's worked out. There was some other stuff but I think a lot of people think it would be too hard to get the existing space to support a new class and defining the rules might be too hard. It might be, I think those are the practical difficulties.

Paul Vixie: This is Paul Vixie. I'd like to weight in also. A lot of what Paul Mockapetris designed was much better than we knew, certainly better than I knew at the time I was working on it in the first couple of decades. You really can't have more than one query even though there is a field indicating how many queries there are. If you say there's one other than there are it won't work. Similarly, you can't really use classes other than IANA. I know MIT with Hesiod tried to do that but we had some trouble interpreting the scriptures in the years before I knew Dr. Mockapetris personally. So, we pretty much broke all of that. Now, those ideas were good. And if the code had been a little bit more passive and had

said -- I don't understand what to do with it so I'm going to ignore it instead of -- Gee, I don't understand what you've said, so I'm going to reject it. Then we could explore on top of the platform the world has today. But that's not the situation we find ourselves in and in fact it would be easier to roll out a new protocol on a new port number than to make classes or multiple queries work on the existing port number. I apologize for that but that's our situation.

Paul Mockapetris: I think when you opine about doing some sort of significant revision, you're probably thinking about a new port number or a new transport vehicle on top of something or other. You probably would want to have some backward compatibility in order to consent people to upgrade. It always seems to me with ITT4 and D6 I was chair of the ITF when the current D6 stuff was selected in a beauty contest. You have a defined upgrade from D4 to D6 whereas the DS tried a different approach which was to have a thousand little modifications proposed and some of them adopted into the code base over the years. So, the DNS versions were like a real number rather than an integer of six. So, those two things have interesting consequences out there. But I think if you wanted to you could go to a new port number and perhaps a new packet software. If only we'd scaled the NTU by three orders of magnitude like we scaled the link rate we could argue about whether it's three or five. But let's say it's at least three orders of magnitude. A lot of the problems we had about getting DNS adopted would disappear. There's also been a bunch of people who decided that for your own good we will in our software implementations basically restrict what you can do with the protocol. I'm always amused reading various discussions about whether the characters allowed in the domain names at various times and even today there's a lot of confusion about that where the original spectrum is very clear that the arbitrary bites with the case mappings and some ideas. We assumed there might be other classes with simple binary or other parts of the domain space which would have binary matched rules. But that's not what happened. Somehow or another DNS turned into a seven bit protocol. But that's history. Let's talk more about the future if there are future questions?

Unidentified Speaker: We're looking for future recommendations as well, right?

Paul Mockapetris: Sure. Anne-Marie. I see you're back. Do you want to follow up?

Anne-Marie Lowinder: Yes. I'm back. The only thing I'm thinking of is I'm not convinced that DNS is that bad. I mean, it works. And I think that a lot of rules are implemented and even though we've been talking to a couple of people who started to discuss about using HTTP for the alternative root system, we did have a meeting that scrutinized the idea. To be honest, they haven't been thinking much about what they're trying to do, what problem they're going to solve with that because you don't solve diffused problems with diffused technical solutions. They're technicians trying to come up with better solutions for things we haven't even thought about how to implement. I'd like to start with what can we do with what we already have to make it as good as possible because we haven't really done that yet. I mean, as opposed to the DNS we have right now, how we can do better, improvements we can do. So, before we talk too much about replacing DNS with something that we even don't know how it works, I'd prefer to be more self-critical and just see what we can do with what we have. It's not only about identifying things. That's probably somewhat easy. The hard time is when you're trying to find different entities that you have identified in one way or another in a very smooth, efficient, and quick and trustworthy way.

Alice Jansen: Paul, I see Olivier has his hand up.

Olivier Crepin-Leblond: Thanks very much, Alice. I have another question totally unrelated to these previous ones which I guess is technologically related but at the same time policy related. We are speaking about the expansion of fields in DNS, et cetera, and of course the string of DNS has brought a lot of new record types in the DNS. And I've also noted some record types that were originally designed for something that's changed over time. The TSC record

type for example was originally for text -- non-machine readable data, human readable stuff, but it now carries stuff like center for policy framework, SBS, DTIN, et cetera. And while there has been some attempt to clean this up a little bit, like for example the DTIN -- the FTS record, it doesn't seem to have been taken up. So, we're seeing a lot of things being stacked up. Do you see any opportunity to come in there with record types and sort of in a more managed manner than it is at the moment or do you see any threats to this as time goes with just so many more people working on new types of possible record types, thus ending up with the DNS that's used for something completely different from what its original intent was.

Paul Mockapetris: I'm sure other panelists are going to have something to say about this but I'll give you my perspective. I think the DTFTS thing illustrates a problem. Basically people used text records and then when they got that to work instead of wanting to go through the effort of creating a new type and changing the code, they said -- It works. We're done. We're out of here. I think the thing to do would be to try to figure out how to have a better path in the future for standardizing that data or perhaps just standardizing the idea the that we go to sub labels instead of IRR types. But think about a new scheme that would work better for future efforts. I think we're going to learn from it and persuade people to go back and clean things up is also kind of difficult. The other thing is while the ITF would like to believe it controls the standards out there, the standards have also been set by the IFC and Microsoft and so forth from the standpoint of what they put into the code they ship. It's not just those two organizations, it's those two I'm most familiar with. IFC at one time attempted to outlaw the underscore character in domain names and Microsoft made it compulsory which I thought was a delightful pairing. Address records are used as opposed to addresses. People will pump things through domain names, people have pumped video through DNS. Certainly there's not video data type. I think what we've got to do is try and realize that the DNS analogy I think is the DNS is always going to be a river and we want to make the channel down the middle of it as navigable and useful but recognize that people are always going to do strange things. I think the thing that surprises me is that the DNS works as well as it does given the amount of misconfiguration that's out there. I think people over the last 30 years, pretty much always found that DNS configurations are about 50% of defective, or there's at least 50% of zones that have some sort of small error in the set up and DNS makes that much more problematic.

James Seng: This is James Seng. I certainly follow the comments are the space. The use of text record, the TSC record is this kind of kitchen sink which gives you immediate gratification. I don't have any barriers. I can define stuff in TSC and I have this pseudo-barrage type without going through the rigmarole. The real problem of course is misinterpretation, some other group can interpret what's inside these text records with intention Y. I put stuff in there in a completely different motivation with intention Z. The folk in the other community see my text record and misinterpret my intentions. That's a problem. In the scheme of things it's a relatively minor problem. I think the real issue and problem in the space is when you put out different behavioral mechanisms. Streaming video in the DNS is one thing. Using the DNS as an IP tunneling protocol is perhaps going all the way where you're trying to completely eliminate all forms of classification and rely on immediacy of delivery and you observe it's one of the few things that moves from one end to the other if you morph it the right way could possibly be the DNS. Why don't we use it for IP as a tunneling mechanism? Once you start to play with the behavioral aspect of the DNS you start to think about the performance of the entire system. I have to admit one of the few things about the DNS that completely works and actually might have worked today is the fact that the churn rate of information in the DNS is low and this is leveraged by location. The reason why is it will collapse under the load, the reason why even modest attacks on the DNS has repercussions is because we rely on location to eliminate that load as close as possible to the querier. If you start employing behaviors in the DNS that in a large scale alter its behavior you start to make some inroads on the characteristics of the performance of the DNS and change that. And that may or may not

be a good thing. So, you know, the behavioral aspect is as much as issue as the interpretation of the resource record. Thanks.

Paul Mockapetris: Okay. I don't see anyone who has their hands up here.

Unidentified Speaker: Perhaps if anybody has an additional recommendations or issues to be considered, if they don't have them now but they come to mind after the call ends, we do have a mailing list. I don't recall what it is off hand, Alice. Can you remind people what it is?

Alice Jansen: ITIPanel@ICANN.org.

Unidentified Speaker: Thanks.

Paul Mockapetris: Yes. I think that if you're on the internet you ought to be able to find one of the panel members if you don't feel like talking to all of them. I'd appreciate any ideas, particularly ones that go off in a new direction. I think we're getting close to the time to wrap this all up. I don't see any last minute things. I'll thank you all for your questions. There's the ITIPanel@ICANN.org mailing list to post public comments and again, thanks for attending and being a great audience. Alice, do you have anything to say to close it up?

Alice Jansen: No. This concludes the webinar. Just a note that the recording, slides, transcript will be made available on the website on the announcement page and the panel's dedicated webpage. With that I think we can close the call. Thank you so much for joining.

Paul Mockapetris: Alright. Bye-bye. Thanks, everybody.