ICANN 55 - Marrakech
Board with TEG
9 MAR 2016

Introduction

[*This is an edited transcript of the Technical Experts Group, 9 March 2016. If the editing has introduced errors or lack of clarity, that is due to the editor, not the original speakers.*]

>>STEVE CROCKER: So thank everyone for coming to the Marrakech technical experts group.

>>DAVID CONRAD: One of the topics that I was going to raise in any other business but will, instead, raise now is that I would, with the acceptance of the TLG, like to propose the technical -- the office of the CTO's research agenda into the TEG for your input.

We are looking for input as to appropriate topics for the TEG to explore. We will be sending email to the TEG mailing list and would appreciate, if you are agreeable, for you all to provide input. If you don't think that's a good idea, please do send me email directly. Preferably before the beginning of -- wait. What is this month? Mid-March. March 15th. And we can -- we can discuss it on the mailing list.

But my initial plan at this point is to try to use you all as a sounding board for the research agenda that we are going to be pursuing within the office of the CTO group.

Reserving Special Use Names and RFC 6761

>>ALAIN DURAND: Good afternoon. For those who don't know, I'm Alain Durand. I work in the office of the CTO at ICANN. We're going to talk today about an Internet draft that was published yesterday on issues on the special name RFC 6761. We're going to talk about problem space, we're not going to talk about solution space, but some of the issues that came up with this.

So as a refresher, 6761, as the IETF likes to refer to documents by their number, is a document that allows the IETF to reserve special use names.

It can be any name, including but not limited to TLDs.

It has been used twice since its creation. The first time was to reserve .LOCAL. That was done for the Apple Rendezvous protocol. And more recently for the .ONION, which is used by the Tor protocol.

The last reservation for .ONION has been quite controversial. There have been thousands of emails that have been exchanged on the IETF mailing list, and this essentially does reveal that there are issues in this RFC 6761 process.

So a number of us came to write a document to describe a little bit those issues.

So there are four authors of this document at this moment. Peter Koch, Joe Abley, Warren Kumari, and myself.

So the first set of issues that we uncover were architectural issues.

The name space overall is not just about the DNS, the domain name system. It's actually larger than that, and it has never been really formalized. There are now a couple of documents that try to address this. And there are new name resolution technologies that are coming up that want to use completely different protocols and technologies and sometimes even way to represent names and they want to share some of the same name space. They want to use the reserved top-level domain as a signal for application to use a different treatment of those requests, and so that's what is done with .LOCAL, where it's a suite that says, "Don't use regular DNS, use something called MDNS, or multicast, and this is -- resolves somewhat differently." Or .ONION, where it says, "Do not send any requests at all on the DNS. This is going to be the Tor protocol and it is resolved completely differently."

There are many other candidates to use this process. It's not just about those two. I don't know if we are now in the tens or in the hundreds but there are quite a lot of them that would like to use this.

For example, the new name system, there's all kinds of names that are built on chain blocks and things like that.

Now, reserving a top-level domain to do this is one of the ways to do that, but that's not the only one. There could be other ways.

So if you look at the URI, there is a left part, a middle part, and a right part.

So today, the left part is http, colon, and people could use http, colon -- http, dash, their favorite name, as something that will signal that this is different or they could use, as ONION is doing, the rightmost selector to say if you put ONION in there, that it means that something special needs to happen, or you could event something new like a middle selector. Instead of saying http, colon, slash, slash, we could have http, colon, slash, selector, slash, and then the rest.

So this is really a tussle between application and network.

A contest between application and network locus of control

So if you talk to application people, they say, "Well, we don't want to change the UI format. This is way too complex. The shortest way for us to deploy new application is actually to use the rightmost label as a switch to say we want to do something different.

Talk to network people, sometimes -- not always, but sometimes, they worry about bad precedent in that space. If you remember in the email days, you remember we had .UUCP, .BITNET, that created a lot of confusion, so there are worries that we will be repeating the same thing.

So we now have this process, 6761, that has been run twice, and if you look at running it as running code, we are seeing that there is a lack of clarity and there are questions that are a little bit ambiguous.

For example, in RFC 6761, it asks -- offers a candidate to this special treatment to answer seven questions on how this is going to be used, but those questions are not really enough to evaluate the technical merit of the candidate in there.

For example, in the case of .ONION, they say, "Okay, reserve this, reserve this, reserve this," but why should this be reserved was never really answered by those seven questions.

Now, if we set aside the technical aspect of is it an interesting protocol for which to reserve a name, there is a second question about what name should be reserved.

And IETF has no process besides the IESG approval to go and evaluate names.

ICANN has a process to review gTLD applications to check for some geographical issues, some intellectual property rights issues, some string similarities. I mean, there is a very, very long list of those things, because folks have been -- who are dealing with names thought that those were important issues to consider.

In the case of the IETF, a decision is made by the IESG, and the way to deal with people who may not be happy with a decision is for an appeal process.

So there is -- this is at the intersection of IETF and ICANN. We now have two processes to deal with those rightmost labels or top-level domains, depending on how we want to call those things.

IETF Process and ICANN Process Not Coordinated

We can simplify this by saying we have a positive delegation and a negative reservation.

So positive delegation means that it's an ICANN gTLD that follows all the rules that are in the gTLD program or a negative reservation that IETF says this is something technical, we don't want it to be delegated, first we are going to reserve it.

Well, the issue is, there is no coordination between those two processes.

And one thing to note is right now the current round of gTLD applications is closed at ICANN, but even if it were not closed, in the current round there was no possibility to go and reserve names. All names will have to be delegated.

So if we think about the next round of gTLD, that might be something to keep in mind while designing the rules for applicants.

So that's essentially the slides that I have here. If you are interested into this discussion, I simply suggest that you read the draft. If you can go back to my very first slide in the deck, you will see the name of the draft. Again, it has been published a day or two ago, and the discussion is happening in the IETF mailing list on the DNSOP working group. Thank you.

>>JONNE SOININEN: Yeah. Thank you. Jonne Soininen. I'm the IETF liaison to the ICANN board, just kind of like as an individual here commenting on a little thing that IETF doesn't really reserve TLDs. IETF, it reserves names for technical use. And they might or might not -- or they actually do not usually have much to do with the DNS.

So for instance, with -- with .ONION, it uses some completely other technology to resolve those names, so it's not based on DNS. As such, it's not a TLD. It looks like one, but it's not a TLD in the sense as we understand it in ICANN.

>>DAVID CONRAD: This is David. I'd actually argue that it is a -- I mean, since it is in the domain name tree, it simply is not a DNS implementation of the domain name, so it's a -- you can get into semantic arguments here, but I do take your point that it is a -- outside of the context of the DNS.

>>ERIKA MANN: I just wanted to be certain I understood you right on the last page. You made the comment that you practically recommended to consider maybe to foresee a requirement for special use in the new round for -- Is this correct? For special use names in the new TLD round?

>>ALAIN DURAND: I'm not at this point recommending that you do that. I'm recommending that it's time to think about --what it will be.

IAB Activities in this Area

>>WARREN KUMARI: Warren Kumari, IAB TLG person.

So I largely wanted to comment on what you had said, David, and also mention that the IAB and the IETF is looking at this as sort of a whole problem.

The IAB has a program which -- what's the -- names and identifier program, and there's going to be a BOF in the next IETF alternate resolution context and what these are both looking at is the fact there's sort of one global name space and the DNS exists in that, and there are other potential uses of the name space, and that there needs to be some coordination here so that there doesn't end up being conflict.

Patrik Faltstrom, chair of SSAC, the Security and Stability Advisory Committee of ICANN.

As it is -- as we are chartered to sort of keep our eyes on things like this, yes, we are keeping our eyes on this. We created a work party that is looking at the overall sort of name space-related issues.

And I can confirm that one of the things we have discovered so far is just like you said, David, terminology is kind of important here. On the other hand, we are very happy to see that other groups are working on this issue and just because we are looking at it doesn't imply we have to say something. So if other people are doing their job, then we don't have to. Thank you.

>>DAVID CONRAD: Thank you. Are there any other questions or comments or observations regarding the 6761bis problem space? If not, thank you very much, Alain. And I guess we will now go to the presentation by Marc Blanchet on label generation rules.

Actually I did have a question on the 6761bis stuff. I am aware that there were a number of additional requests for names to be put into the special use registry, things like .CANOE, .ZKEY (phonetic), .BIT, all of those.

What is the status of those? Warren?

>>WARREN KUMARI: So I cannot speak authoritively, obviously. But as far as I know, that's on hold while we're busy sort of trying to figure out the sort of larger metasolution to the problem, a. Better way to actually deal with this and understand the correct processing I think is what's happening now. And those are on hold.

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

Did you want to comment, Jonne?

>>JONNE SOININEN: Just to say that that is my understanding as well.

Label Generation Rules

>>MARC BLANCHET: Label generation rules, so this talk is about two things. One is a super fast primer for smart, technical people. And, second, is a few, you know, concluding remarks on issues or challenges or topics that might be interested in general for board, community, and the technical people.

So in one page, the LGR is essential for every script that a script exists, you would define the code point repetoire, which is a list of the code points you want to use; define the variants related to those code points, if they exist; and then define the rules applied to the whole labels. All this is to -- for now being done for defining non-ASCII labels in the root zone.

I will show you a few examples in a minute.

All those are defined in an XML language called label generation rules, which is actually also a working group item in the IETF called lager. And that domain specification has actually been pushed yesterday to the IESG for proposed standard.

So given that it's related to each script, so there is a team of experts for each script called generation panels that is being done here in ICANN. And they do their work

to the three topics at the top. Then they send their proposal, their work, which is actually the LGR itself, to a panel that integrates all the LGRs together. All the LGRs together defines what is permitted in the root zone. Obviously, that work could be used at other levels in the DNS industry, and we'll talk about this in a minute.

Examples of problems

So some examples. This example is not done by me but obviously some people knowing Arabic. So here's an example of code points that you -- as you see on the screen, some are with X saying, We are not using those. Some are okay, some different character sets, and you select a subset.

Then you define -- you find variants. So, for example, on the right, you have the variants shown in green and red. Those are not similar in a sense of look similar but actually real variants in the actual script which means they mean the same thing in the script itself.

Again, a good example in Arabic of the third bullet I've shown before, which is you need rules for the whole labels. For example, in this example, those two characters, those two code points, cannot be mixed in a single label because of Arabic writing rules.

So if you see the first two are the same code point used twice in the label. But if they are mixed, they are not possible.

Again, the XML specification has been recently sent yesterday for IESG for proposed standards.

It's actually been implemented multiple times by multiple people. So it's a pretty major specfication.

So these are examples of code points of an LGR. So you see that the code point 0620 is allowed in the LGR, therefore, it can be used. And you see the next one was defined 0622 and has four variants.

Disposition of LGR Analysis

And those have a type or what we used to call a disposition, which says
block,
allocated,
allocatable
blocked.
So what that means is the community -- the Arabic community says those variants should not be used except one, the one which we just marked allocatable. That

doesn't mean it will appear on the root zone. It just that it should be -- it's okay to use that variant.

This example here is an example of a rule that says essentially what I just shown two slides ago which says two code points cannot be used at the same time in the same label. So it essentially reflects that rule, the linguistic rule. And you apply it. And say -- at the end, it says "action" which says "disposition invalid" which means you cannot use the combination of those two code points in the same string or in the same label.

The way you decide if a code point should be used in the LGR should be obviously protocol conformance. So it should be an IDNA PVALID. That code point or that character should be in use which means a modern use. In the integration work, we actually verify that all those code points are actually in use. So the generation panels have to provide evidence that those code points are in use. And there's other criteria that I'm passing around.

The last one is we want the least number of allocatable variants because it could create a combinatory explosion. So that's the topic I want everybody to be aware of in the sense that if you have multiple variants and they are all allocatable, then if a single label as multiple code points and each code point has multiple allocatable variants, then you just create a combinational explosion.

So you may then go to a point where a single label may have a thousand variants. Do you want to do -- the DNS wants to have 1,001 entry for this same code label, right? So challenges here.

Allocatable and Variants of Allocatable

>>MARC BLANCHET: So impacts of this work is the LGR will be used for determining eligibility of a label for the root zone in all its allocatable -- and I invented -- varianted labels. So it could be also used for other levels in the DNS tree and also to determine if a specific label is valid in general in the whole ecosystem. So I may want to verify if a label is valid in the sense of -- on the language side of it or the script side of it in a registration system and other places and application maybe. I'm not saying we should or we should not. I'm just saying it could be used at other places.

That's really my last slide which are kind of topics to think about as possible impacts. One is what if a single applied for a label of the 10, 100, 1,000 allocatable varianted labels, right? What do we do here? So that the first probably answer is to minimize the number of allocatable variant labels or allocatable variants in the process of creating the scripts for the LGR for every script. But then we may not be able to restrict as much as possible.

We do have also cross-script variants where a variant is actually used in other scripts. So in this context, you have communities that are impacted. So this is more difficult also.

I was saying that there are rules related to scripts, for example, the illustration I made of -- two code points that cannot be at the same time in the label. Well, how much do we go into the actual kind of grammar of the language or the script about this, right? Do we need to go further too much because, in turn, DNS, there's no grammar. You can use whatever label A, B, C, D, E. That doesn't mean anything, right? It's not conformed to a grammar or anything.

But obviously in languages with ideograms, then you have restrictions or you could make them. But how much do we go there?

Finally, another topic of interest is the fact that recently -- well, it's getting old now, but some time ago, IAB made a statement about the fact that there's potential issues about our Unicode encoding related to the IDNA assumptions. And the result -- and result of that was IDNA tables were freezed for Unicode 6.3, which means right now -- and all this work is currently being based on Unicode 6.3 and its related IDNA tables, which means that any code point that is defined after 6.3 is not being considered. So it may impact people that have, you know -- that are communities that have scripts that are being encoded after 6.3. Food for thought.

>>RINALIA ABDUL RAHIM: Marc, I think what would be useful for the board members around the room is if you could indicate clearly whose problem those interesting topics are and who gets to solve it.

>>MARC BLANCHET: So last one, Unicode IDNA Table 6.3 -- and it's my personal comments -- I think it should be IETF and Unicode people.

>>MARC BLANCHET: So IETF Unicode, grammar, and cross-script variants, so this is in the current process of LGR and generation panel and integration panel. However, those are kind of difficult topics where we may -- may end up not being able to agree -- right? -- in the sense that at some point in time there's an arbitrary line where you have to -- to do.

>>RINALIA ABDUL RAHIM: So for the third bullet, it's the integration panel that makes that decision?

>>MARC BLANCHET: I guess by default, because we -- we have the duty to accept or refuse an LGR. I would not like us, me being a member of the integration panel, not to do this because it's -- it's up to the language people and script people to do that.

The first one, I think it's everyone, which is -- at some point in time it may impact ICANN in the delegation of TLDs because if -- if we end up having 1,000 allocatable labels, then, you know, I think we all agree it's an interesting problem.

>>PATRIK FALTSTROM: Patrik Faltstrom, chair of SSAC. I will speak in multiple roles here, as I have multiple roles regarding internationalized domain names.

First of all, yes, this whole thing might be my fault because I wrote the document, the original standards, so you can beat me up afterwards.

Secondly, let's start with the Unicode Version 6.3 issue, and here I speak as the liaison from the IETF to Unicode Consortium.

The problem there is that there is one character that was added to Unicode 7.0 that might create issues regarding backward compatibility. I'm saying "might." Because this is -- this is where IETF has not reached consensus on approving the new versions of Unicode later than 6.3.

So that is, from my perspective, an IETF problem.

But unfortunately, there are not enough people with script experience -- and in this case it's the Arabic script -- which have been able to inject information, data, enough in the IETF for the IETF to sort of draw any conclusion there.

Let me then change hats and say that -- and speak as the chair of SSAC.

We in SSAC are concerned over the fact that not only do we have these issues, but for example, specifically the last one, that we have not moved forward with Unicode -- with the Unicode version, and the question is whether issues like the ones that is just presented here, by themselves, create a security and stability issue, the harmonization or non-harmonization of various IDN things, why we are not moving forward.

The work party that we just created inside SSAC, we have decided just because there is a lack of people with -- with sort of knowledge about specifically cross-script issues and cross-language issues, we will run that work party more open than we normally do and hope that we will be able to get enough people in the same room, as Marc was suggesting, so that we actually can do a proper evaluation of the current situation.

So I hope that might, for example, get ICANN and IETF together enough energy to actually move forward here.

The current LGR work that you presented looks very much like picking code points and selecting them individually. The IDNA 2008 standard is, compared to older

versions, not selecting individual code points just because -- it's, instead, an algorithm. And that is just because the IDNA 2008 standard would be independent of Unicode version, which means you can take a newer version and apply the IDNA 2008 standard and everything is fine.

If it is the case that IETF quickly -- which we know is about the same speed as a glacier or something -- is approving new versions of Unicode, will that -- do that imply that the LGR work has to restart for newer versions of Unicode?

>>MARC BLANCHET: It depends on what you mean by "restart."

>>PATRIK FALTSTROM: Do you want me to clarify?

Do you have to reassemble the panels to reevaluate the new added code points?

>>MARC BLANCHET: Okay. So what happens is the way with -- the procedure that was defined for this work is that the -- you always be conservative in -- so there is a cost for having every signal code point to be defined in the LGR because as soon as it's permitted, then it's almost impossible to remove.

So the corollary of this is that it's always -- we -- it's -- by fault, it means that you can add new code points later in the process without any issue.

So the LGR is more constrained, so it will be easier to add new code points because of the way the procedure is done.

Then the question is about the specifics of what is envisioned.

For example, if Unicode 7 brings additional, say, for example, Arabic code points, but then the Arabic LGR has been done, then obviously that means, you know, Arabic experts to look at this and then, you know, some process to update the LGR.

How is this applied now?

>>GEOFFREY HUSTON: Geoff Huston, SSAC member.

I have a very quick question to you, Marc, and I was a little bit unsure from your presentation whether these LGRs applied in general as Internet standards or specifically limited it to a certain set of applicability such as the root zone.

If it's the latter, if it's a limited applicability, what would be the issues in trying to seek an Internet standard on LGRs that would apply across all zones in the DNS?

>>MARC BLANCHET: So what IETF does is just the actual XML language, right? So that's -- that's not the standard. And the current work is for the root zone. But there

is parallel work for a second level, and obviously it's not the same because you don't have the same constraints. As we know, the root zone doesn't have the same, you know, constraints in the second level.

So the mechanics could be used, but then we also know that the DNS tree is all -- you know, every anchor of the DNS tree has its own policy, right?

So -- so the groundwork could be done, but then it can be applied in different ways, depending on the policy at each anchor of the tree, right?

>>RAM MOHAN: Thank you. I wanted to respond to Rinalia on one of these things.

That very first bullet point that's there, a question of what to do if a single applied-for label has N number of allocatable variants, that is something that is sooner or later going to land in the ICANN board's hands, and that's part of the reason why this whole topic of variants and what to do with it, it's a pretty complex thing and you have to sort out the various parts of the tree before ICANN can actually go out and tell the community, "Yes, you may have variants," because it has impacts in multiple places.

Existing LGRs as IDN Tables

>>KIM DAVIES: Thanks. Actually, this, I think, partly addresses Geoff's point. There's actually, more or less, at least 960 LGRs out there today. We just call them IDN tables. So the format that is coming out of the IETF now is creating a universal format to express them, but it's really the next generation of IDN tables which anyone uses IDNs in new gTLDs today is mandated to provide to ICANN and we list them on the IANA Web site.

So the root LGR work is to create it for the root, and I think given the expertise that's going into it, we're hoping that will be a baseline that a lot of TLDs and other ones will reuse, but today it's quite fragmented. There's a lot of different rules. And hopefully this effort will sort of start to bring them together.

>>RAED ALFAYEZ: Yes. My name is Raed Alfayez. I just want to raise an issue. Variants is there a long time ago. For example, for us, we see a TLD like "Black Friday" or "Accountant" has more than 1,000 variants, because if you consider uppercase and lowercase, they are variants. And these variants are working from the DNS point of view, so nothing need to be done at registry, at application. Nothing at all.

This is a problem. So if you say -- put magic numbers like 1,000 variants, what shall we do, 100 variants, what shall we do, variants -- we know all the IDN (indiscernible) so it's not something clear and that we are dealing with complex scripts, so in our script there are more than 50 languages inside it. We need to have

actually solutions. Not only to be afraid and just say don't have variants. Variants is normal. Believe me, it's normal. Like uppercase and lowercase in English, it's normal. Variants is normal in our side.

So we want from the community, the technical, the ICANN, try to find solution to make them, instead of going -- asking the registrant to register a domain name and enable three or four variants for reachability, and this is part of the core of the rules of the Internet, the domain name should be reachable all over the world. The problem that apply at least to Arabic variant names using Arabic keyboard, if I go to Internet cafe in Pakistan or in Iran, I will not be able to reach it because they are different, yeah, like the example you show. This is just a note we need to enforce or ask the community to find solution to automate the domain name hosting, so no need to do it in every place or have something like a solution or a standard to how to make, you know, multiple variants and automatic mail to be hosted like the same -- the original domain name. Thank you.

>>DAVID CONRAD: Yeah. I think that is a summary of some of the challenges that have been driving the IDN program in the DNS for, oh, what, about two decades now.

So it is a known problem and we're -- everyone with -- not everyone, but people who are focused on this stuff have been working on for some time.

And with that, I want to go over to Howard Ben Benn and a presentation on ETSI NGP.

ETSI Next Generation Protocols

>>HOWARD BENN: Thank you. Let's wait till the slides come up.

Okay. So for those of you who don't know myself and Francisco, we're representing the European Telecommunications Standards Institute here, so we look after all of the standards for the mobile industry and we work very closely with an organization called 3GPP, which actually generates most of the standards.

So as the slides are popping up, so ETSI is an interesting organization because a few years ago, we had a look at the -- the architecture within ETSI, and so what we have what we call TBs, technical bodies. They're the main core of the way that we write our standards.

But what we also realized was that sometimes people going out to industry forum to write pre-standards, in particular, and we thought it would be a good idea if we had a process through which we could write these industry forum standards within ETSI to try and get more buy-in from ETSI members.

So we formed what was called an ISG.

So an ISG is kind of part of ETSI but kind of isn't at the same time. Which is interesting when we come to this topic.

So NGP stands for "next-generation protocols."

We started looking at 5G probably about three years ago, so I think I gave a presentation here a while back on some of the early work that was going on. That work has been progressing across the industry.

The pace of that work is increasing all the time.

What we're trying to do is we're trying to base it on fact-based requirements, so one of the things that we did is we went out to the 5G Innovation Center, which is based at the University of Surrey, and that's a -- basically a hub within the U.K. for doing 5G research. It's got -- all the large operators are in there, but also broadcasters, the BBC, are in there, and also some other Internet players.

So if we go to the next slide.

So when we started looking at this whole area, we looked at where is the Internet traffic that's being carried on mobile, what's good, and what's bad.

One of the things that became apparent very, very quickly was when we look at the packets of data that get transmitted across the radio, there is a lot of overhead. And this is a problem within fixed networks, but it's a little bit cheaper in some circumstances to put an extra 10-gig fiber into a network. When we start having to upgrade thousands of radios across the network, it gets a lot more expensive. Spectrum is also a massive cost to mobile operators. So the need for additional spectrum to carry more traffic is putting a great burden on the operators.

So when we dug down into what that traffic actually is, it was quite interesting.

So, first of all, the thing that was apparent was that our old friend TCP/IP, it's all designed to run on fixed networks originally and also wireless networks later on. Generally these networks were considered to be reasonably low speed, were unreliable transmission.

Our modern networks aren't like that. So the way that we handle things like errors within the network has changed over the years. So we now run hybrid ARQ. We have extremely reliable transmission based on the radio. And so we solve a lot of the transmission errors at the radio layer.

So by the time we get out to the network layers, the reliability is actually very good, so greater than ten to the 6, ten to the 7 in most cases. So we end up with very, very lower error rates.

Of course, we're faced with increasing demands on the spectrum. More and more access of the Internet is now of mobile phones. We often talk about numbers. So I know that the Internet community likes to use this 3 billion, or whatever the number is today, of people accessing the Internet. We're up to 8 billion mobile subscriptions at the moment. So about 6 billion of those are human beings, and there's about 2 billion of other things connected to the network today. Massive numbers.

So one of the things that we started to think about was what if we started again. Now, this is a very brave move. So one of the reasons I wanted to present here is that I think this is a good place to share this information and, basically, allow people to kind of come back and say, you know, this will never happen, which I think was the reaction of most people. Or maybe we can see this as an opportunity.

If we do have the chance to start again, maybe the best place to start this was not in the IETF because they've obviously addressed these issues over a number of years. I think the issues with TCP/IP over -- a lot of the other protocols as well over mobile networks has been addressed many, many times in the IETF.

I don't think many people have been brave enough to say, What happens if we got rid of everything?

What Happens if We Get Rid of Everything

So, again, what we started doing was looking at some of the use cases around this, what's driving this. So for 5G, there's a lot of drivers. We often talk about the verticals. Within the telecommunications industry in general, we do deal with a more holistic view of services. So the IETF did a great job of designing the protocols. Don't necessarily bring in the whole end-to-end use case in the way that we do in mobile networks today. And so we're increasing this. The way we handle Internet of Things, the way that we handle high-speed mobile broadband, we look at holistic solutions across the whole network.

So here's just a highlight of some of the things that go on within our network. So as we know today, it would have been nice if mobile networks would have been able to use an I.P.-based protocol, mobile I.P.-based protocol. But it wasn't to be. They weren't up to the job when we first started mobile networks.

So we had to add GTP on top, that's the GPR tunnelling protocol. We had to modify that over time. And we have a GTP-u tunnel at the moment. But we have to have this tunnel because that's the way we handle mobility within the network.

So when you move from one cell tower to another cell tower, you keep the same I.P. address, but that's magically all tunneled through this GTP protocol which actually does all the mobility for you.

On top of that, we tend to run some IPSec tunnels as well. So when you start looking at putting tunnels on top of the tunnels, the overhead just goes up and up and up.

This is what a real life network looks like. So as I say, we're trying to base this on a fact-based analysis. So we went along to our friends EE in the U.K. They've just been bought by BT. So there's an EEBT now network. And they started looking at all the boxes in the network. So this isn't the kind of diagram that standards bodies normally look at because this has got more to do with protocols per se, just what what the boxes are in the network.

It's safe to say that the operators would like to have fewer boxes, not have to add the boxes to add things like security, which should be built in from day one.

Next, this goes through some of the various stacks that we run in various parts of the network. Again, packet zip-up events, protocol stacks all the time, adding complexity.

So we were talking a bit before about where we are going with 5G. We're looking at increased number of services we're providing, especially to verticals. So to some extent, the mobile networks are starting to do this today. We are having more and more modifications to our standards to handle IoT, smart city, smart buildings, gigabytes per second. We are getting higher and higher data rates all the time.

And what we're doing is we're modifying the radio networks to be far better at being adaptable to these individual services. So when you go along to whoever your service provider is they will be able to provide a tailored service for what you're doing.

Today, if you want to send eight bits of data for a little temperature sensor that's connected to your network, you have to go through a wide variety of hoops. You have to set up PTP contacts. You have to do all sorts of stuff before you can actually send that one packet of data.

We want to make these networks more programmable. So we're looking at things like SDN and some other techniques. These are the wide variety of things.

And even go down to things to like self-driving cars, not really self-driving cars but cars where you are in the car talking to another car in a very secure manner.

>>JONNE SOININEN: Yeah, Jonne Soininen again and the IETF liaison to the ICANN board. But here kind of like looking more from the day-job angle where I work for Nokia, just a couple of things, however, that I wanted to point out in kind of like just to make sure that they come through correctly here and in the -- and are emphasized correctly, I was about to use a wrong kind of, like, description but let's try again.

3GPP is starting 5G work. And it is expected to start this month, I think in the next meeting. And to my knowledge at least, that work is not connected to this directly. They might use that as an input or might not. It is completely a separate thing.

This is work, like you said. But I want to emphasize this. This is an independent study group. It is -- people can join or cannot join. It's kind of like it can be researched. It can be standardization kind of like. And this is more maybe on the research angle of things.

I hope I described those things correctly, yeah.

>>HOWARD BENN: Yes, that is correct. So 3GPP SA2 are just starting to look at some of the architectural impacts of 5G. It's safe to say that some of the background discussions have actually been proposed into SA2. But you're quite right, this is an independent group.

>>PAUL WOUTERS: Paul, IETF liaison. I would also like to mention that the IETF is a really open organization, and we always welcome everybody to join us and help on standards. Specifically, there's some friction now where the 3GPP are making new standards and coming back to IETF and then wants IETF to adopt it. And then the IETF hasn't gone through its regular working group methods of seeing if they want to implement this or not.

And so, for instance, in the IPSec working group, we now just had -- a couple of days ago something came in on a method for an addition to the protocol that IETF would have done completely different. And now we're in this sort of dilemma where these two groups definitely need to work better together to extend the protocol because, otherwise, we are going to fracture these protocols.

>>HOWARD BENN: Yeah, I think that's an excellent point. And I think this is where organizations that are working similar areas at high levels just need to continue to work together. So I understand over the years, there has been a lot of friction between the different processes used in 3GPP and the IETF. And there's been certain cases where we found that the 3GPP specifications have been delayed by years by not having proposals go through the IETF in the speed that they have gone through 3GPP. I think that's something we can work together on.

I think it's safe to say for this particular area here there is no proposal at the moment for ETSI to write a standard in this area. I think that what we will be doing is just generating the input and then almost definitely be passing that to the IETF for them to go through their normal processes.

>>JONNE SOININEN: I just had to comment on what Howard said about the IETF delaying it through 3GPP protocols for years. The thing is that I've worked on both sides and sometimes really on both sides at the same time. Let's say that there have been things where the -- so to say the 3GPP community -- or 3GPP originated community in IETF and other parts of the IETF have not agreed on something and stuff like that. That is quite actually normal in standards that people don't agree, and then when people don't agree, something might not finish as quickly as one would think.

Mutually Agreed Norms for Routing Security (MANRS)

>>DAN YORK: Something like that. So I am also on the program committee for the TEG, which says something about the -- when you're drafted in there.

And I'm not here to talk about IP addresses or DNS. And I guess to Marc's point, I guess it shows that I moved from the technical side of the Internet Society into the communications group because I'm using -- I'm not using a white template, I don't have code, and I don't even have a network diagram in here. I'm actually trying to make it a little bit accessible. So we want to talk about how we can work together to make a more resilient routing system.

I want to put what I'm talking about in context. You know, we just spent a day -- many of us spent the last six hours before this in a DNSSEC workshop spending a lot of time looking at how we improve the level of trust in that area.

There's other work that happens at a lot of different levels to bring about a more trusted Internet. I'm going to talk about this issue around trust in the routing infrastructure.

To step back a little bit, kind of the first big incident that looked at this was back in 2008 when a misconfiguration in Pakistan wound up bringing about YouTube -- redirecting all interest in YouTube through Pakistan's network. Sucked it all in there. All gone into Pakistan for everything. And it was -- it was there.

You'd like to think that in the years since, we would have fixed this, but if you go to the next slide, you see that we're still dealing with this. And there are very real attacks, very real things. Some of them are legitimate attacks that are going on. Some of them are misconfigurations. Some of them are issues.

If you follow what's now called Dyn Research -- it used to be Renesys. If you follow them or if you follow BGPmon, they're reporting these things all the time. We're seeing these kind of redirects and things going on there.

We're also seeing large-scale massive DDOS attacks, distributed denial of service attacks, botnets, all of these things that are going on around.

So one of the questions is, how do we bring about this network, how do we do this. And the answer really is that what we see is we've got to work collaboratively. We've got to work in some way that works together.

And so this project that's been coming up is called MANRS and I'll talk a little bit more about it, but I want to give you one more data point.

I often ask people how many networks are actually participation in the routing infrastructure, and the answer is right here. It's 53,000 networks are actually part of the Internet's routing infrastructure.

Now, really, about 21- to 30,000 of those are more enterprise networks, end networks, those kind of things, and there's about 5- to 7,000 that are really in the core Internet routing elements.

Routing and BGP

They use this protocol called BGP. They use -- they use that to -- they advertise their routes to each other. They use -- they have a routing table. And they pick the best route, which is usually the shortest route.

So if I want to go, you know, to icann.org from wherever I am, my -- the routers look at the routing tables, figure out what's the shortest path to get there, and it all uses these autonomous system numbers, which, hey, look at that, we've got an IANA URL right on here. So managed by IANA. Okay. Let's go on.

The problem is that BGP is entirely based on trust. Every router trusts the information it gets from every other router. This is why when Pakistan, you know, broadcast its advertising that said "Hey, we're responsible for YouTube," a lot of people said, "Oh, hey, look, here's a YouTube route. Let's go this way," and it went into Pakistan's network.

So it's all based on this trust. The -- the chain of trust goes across, you know, national borders, goes across all sorts of stuff. It's all there.

You know, these are some examples, and, you know, the best path says, "Oh, hey, this is short, let me go this way."

Another issue that we've had is if there's an issue, how do you contact the party, all right? And there's no WHOIS for any of this stuff. It's just how do you contact the person who's got that ASN. How do you go and do this.

There's a number of these technical things. They're called prefix hijacks, route leaks, IP spoofing, which brings about typically the DDOS attacks that we're seeing, these large-scale outages that go on there. All of these are issues.

There's also this collaboration issue. How do you reach somebody? If you know that Pakistan is broadcasting a route that sucks all of YouTube into its network, how do you reach the person at that network -- you know, on that network to go and say "Hey, stop broadcasting that route"?

A number of the groups within the IETF have been working on this. There's specifically one called the Secure Inter-Domain Routing, or SIDR, which is working on this. There are some tools such as the Resource Public Key Infrastructure, or RPKI, which winds up cryptographically signing the source, so that if Warren is sending me a packet from his router, it will have a signature on there that I can know that that really did come from Warren and that it had -- and that he's allowed to, you know, originate a route for that network.

There are some pieces like that.

There's another one called BGPsec which will go and verify that the information that came from Warren's router to me was not modified in transit. There's some pieces like that.

But one challenge is that in many cases, one challenge that you have is that making these fixes on your network doesn't necessarily help you.

There's two issues that we come into that bring a certain level of urgency to this. One is that as we are all focused on many efforts around how do we connect the unconnected and bring the remaining 4 billion people of the world on line. As we do that, we're expanding the number of these new networks that are out there. We're expanding the usage of this. We're bringing more and more people on line, and more things, which brings -- you know, just increases the attack surface of the routing infrastructure that's out there.

The other piece, of course, is what -- the whole world of IoT and bringing, you know, billions of more devices on line in some way. Again, you're increasing the larger attack surface.

The mutually agreed norms for routing security is something that came out of a project that came out of the network operators who were operating portions of the Internet.

In some cases, one example was Comcast in the United States was doing a lot of these things themselves and they said, "You know what? We can make some of these changes. We can stop allowing I -- you know, spoofed IP addresses out of our network. But in order for it to work successfully, we need everybody around us to do that. Everybody we peer with."

So this MANRS developed out of the operator community, and it defines four concrete actions.

Four Concrete Actions for Operator Community in MANRS

This is what has been agreed on as the first step in what are good manners for routing.

- One is filtering. You know, preventing the propagation of bad routing information out of your network.
- Another one is the anti-spoofing. Preventing spoofed IP addresses. If we could do this alone, we could wind up, you know, severely reducing the extent of the DDOS attacks that are happening out on the Internet in some way.
- Coordination. Again, looking at how do we provide a mechanism to reach people when there are problems. How do we go and do this.
- And then also global validation. Facilitating this.

It's more than just a document. It's -- it's -- it's really designed to be a community, designed to be a commitment to go forward.

Ultimately, the goal of the project that the folks are involved with is to look at how to turn this into a quality mark, to do something that you're -- if you're going to peer with another network, you want to find out are they a MANRS-compliant organization. You know, are they practicing good network hygiene, basically. Are they going to work that way.

It's not some magic firewall or some magic box. It's really a commitment of -- of -- to go to these different standards that are here.

In November 2014, this project launched with nine different companies that were there. It's now up to 14 -- 30 -- sorry, 37, I think, with 75-plus networks that are part of that, which sounds small when you look at the overall range, and it is, but when you look at some of the large networks that you have in here, you've got some

large ones like Comcast, Level 3, NTT, some of the other big players out there in the space.

You're starting to see some more participation around there. Of course the map looks great because you had a bunch of Russian networks that joined on and they obviously have big geographic space. We are seeing a few in Africa here that have joined on to be part of that.

The project has been going on with a launch back in November 2014. The ongoing initiatives, they're expanding the group of participants, building this community around there, and also one of the things that's been highlighted is the need to develop documentation and pieces to help people understand what are the precise steps they can do to take and go on with this.

So this is really all I wanted to bring up was to let you all know that this is something happening in the routing space, in the ASN space, to look at how do we go and build a more trusted routing layer.

It's at routingmanifesto.org or manrs.org, and those are there. It's something that people are welcome to join, organizations are welcome to join, who are part of the actual -- this is for the network operators, the people who are truly doing this and agree to sign up, and you can go there and see the organizations that are participating.

And I wanted to bring it here to let you know this was going on and also just, as we talk about how to build a more secure routing infrastructure, know that this initiative is out there and happening now.

I'm with the Internet Society. My colleague, Andrei Robachevsky, in Amsterdam is the driver of this from a facilitation point of view. We helped convene a couple of roundtables to try to look at this, and then the community kind of came together and we serve as helping facilitate it, but that's our primary role. It's the operators themselves who are actually implementing this and spreading the word. So thank you.

>>ERIKA MANN: Thank you so much. Very impressive, this presentation.

Tell me something. How relevant is this, actually? To create a trusted -- to create a trusted routing layer makes only sense if we -- the rest cannot be trusted, no? Is this the argument that in reality, there is a problem, a major problem, that you want to solve?

>>DAN YORK: Yeah. So I mean, the reality is that there -- that this is a major problem that we're -- we've tried -- if -- anybody familiar with BCP38? Okay. Right? So it's been around forever and nobody is implementing it because it's limited in a

certain space, but part of the reality is these technical measures have been around but there's lacking any kind of driver for some folks to go and implement these, because again, you have to go and do this in your own environment. It doesn't necessarily have an immediate business benefit to you.

The benefit is if you do that and the other people do that and the partners you work with, then you start to achieve that benefit.

So what we're trying -- what it's been trying to do is look at how do we build a community and a movement around that to make it so that it becomes something that over time we get to the point time where if you are an Internet-connected network, then you should be -- you know, this is something you should do. You should sign up to this, and you should become part of this. A quality mark, basically. That's part of that goal.

>>WARREN KUMARI: Warren Kumari. So, what this is trying to do is sort of change the cultural norms on the Internet. So if you are on a network, you need to be this tall to play. Otherwise, I'm not going to talk to you. So sort of a general shift of what's expected from a network.

It has also been driving other sets of discussions or at least it's created a place where people are having other discussions. There's a draft in the IETF by Jared Mauch, who some of you probably know, to suggest that router implementers turn on BGP filters by default. Currently when you configure, you know, BGP on a router, it automatically announces or receives anything. This is sort of a change so that the default will be to block everything unless you specifically configure a filter. So it's partly a sticker you can put on yourself to say you do good. It's also sort of a continuation of the discussion.

>>DAN YORK: I will also just add, too, the first stage you see is here. The group is looking at what are some of the next steps that could be taken to raise the bar even a little bit higher. Right now it's looking at getting people to sign on to the base level and then look at that as how do you start to ratchet that up to bring higher levels of that.

>>LARS-JOHAN LIMAN: That actually is -- I was going to suggest that: Could you add a requirement for IPv6 in there as well?

>>RUSS MUNDY: Russ Mundy. I wanted to just also ask if people realize that our current system of moving packets around the Internet, which is what the routing system is controlling, is really based fully on mutual trust. And when you look at Dan's numbers, literally tens of thousands of people -- and you have to trust them all to not do the wrong thing and they do do the wrong thing regularly.

And so this a series of steps forward ultimately ending with the type of thing that you see for DNS security where there's a way to cryptographically verify the correctness of the information that you're getting.

And so this is, as Warren said, an attempt to sort of move the thinking and the social way of doing things on the Internet forward from a "we trust everybody" to, "gee, I something that's trustworthy that I can make my decision on." I don't have to just say, "Oh, I can't do it without accepting it."

>>ASHLEY HEINEMAN: I was just curious. Are there measurements and metrics in place that are proving and showing that this is effective?

>>DAN YORK: Excellent question, and it's one of the items -- there was a MANRS session at the NANOG meeting just this past month where they were looking -- that's one of the pieces they would like to add next really, is look at how we can measure the impact this is really having on some of those pieces. And also there's been some discussion about how can you measure compliance because it is a voluntary type of thing as well. And so there's some discussion around some of that, too.

>>ADIEL AKPLOGAN: Yes, thank you. My question is about the scale of this. It starts from the core as the Internet of people who know each other and they spread the word and they can do it.

But if we extend these to all networks connected, how can we work together to help spread the word, get people to be involved beyond the usual people that get at the NOG meeting, at the ICANN meeting? Because that's where the biggest threat comes from.

>>DAN YORK: Excellent question. First of all, I would love to talk about are there ways that we could engage with ICANN in some of the ICANN spaces and some of those places.

Some of the outreach that's been happening right now has been through the network operator groups, the NOGs, going out to those kind of places.

We actually were just a few months back here in Africa talking to a number of different network operators and much positive interest just looking to get them through the steps of signing on.

We are looking at how do you reach -- right now a lot of the initial focus in the first year has been about how do you get to the core operators and validate some of the ideas that are around here and also now it's been identified, let's get some of the best current practices and some of the way to do this into a nice tutorial kind of form so you could bring it out to all of those stub -- the stub ASs, the ones that are for a single network, an enterprise, a company, and give them a cookie cutter, here's

what you need to go and do. Turn this on. Turn this on. Do this type of thing. And, boom, you can go. We're also -- so that's -- one step is looking at that.

Second step is how do we bring this to other conferences and events and things that are not part of the regular Internet circle of stuff. And we're starting to do some of that outreach.

And then the third part is we've also been trying to undertake -- and I would be curious to talk with anybody in this room. We have been trying to get out -- we started talking to industry analysts, the Gartners and the Forresters and those kind of elements. And we've had some very good initial discussions with them just talking about this kind of thing, looking to try to get to the CXO realm of things, to start to try to circulate this idea that this is, you know, a practice they need to look at, looking at how do we reach that enterprise space.

And that is something that I would love to hear from anyone in here if you've got ideas about how to reach into that space. We're definitely interested.
>>DAVID CONRAD: Moving right along, we actually are in the process of bringing in someone to provide administrative support for the TEG. With the small number of things that I have on my plate, it has proven to be somewhat challenging for me to try to arrange things and do things in a professional manner and for which I apologize.

I do want to explicitly thank the program committee who helped put together the agenda: Jim Galvin, Dan York, Kaveh, who I don't see, and Marc Blanchet. I thank you very much. I very much appreciate your effort.

>>DAN YORK: Well, I'll just put in the pitch since you didn't about the program for the next TEG meeting, whether it's in the B meeting or next meeting or whatever.

If you would be interested in presenting something like this, the general idea is that it is something that the board and staff would be interested in that would affect the future directions of ICANN or that people should know about in this regard.

We've generally said it should not be DNSSEC because there's six hours' worth of that before, and we don't need to have more of it here. But beyond that, we would be -- the program committee would be interested in what people would like to have, something short and aiming for something like this. Ten minutes or so around that.