MARRAKECH – Policy Aspects of DNS over HTTPS (DoH), DNS over TLS (DoT) and Related Issues
Tuesday, June 25, 2019 – 15:15 to 16:45 WET
ICANN65 | Marrakech, Morocco

ALEJANDRA REYNOSO: Good afternoon, everyone. May I ask please to sit down? We're about to start in a minute. Thank you very much.

May I have the presentation ready, please? Thank you. While that is coming up, thank you, everyone, for joining us. We are going to have a high-interest topic on DNS over HTTPS and DNS over TLS. First, the agenda will be that I will introduce you to the session goals and introduce all of our panelists.

There will be a technical overview on the topic, then we will have questions and answers. After that, we will have potential deployment concerns, and after that, there'll be, again, another opportunity for questions and answers, and in the end, we will have a panel discussion on deployment considerations, and we are expecting for all of you to participate in this high-interest topic. For that, we will have roving mics. If you could please let yourselves be known where you are, you will see them with numbers. There's four, six, three, and I guess five is over there in the back. I can't see it, but it should be there. There it is.

Okay, so whenever you have a question, please raise your hand as high as you can so they can come to you and give you the

microphone. And if we could have the presentation ready, that will be awesome.

I will start by introducing myself. I'm Alejandra Reynoso from dot-GT ccTLD. I'm Vice Chair of the ccNSO. With me, I have Danny McPherson. Danny is the Executive Vice President and Chief Security Officer at Verisign, and Danny is also a member of SSAC.

Also, I have Peter Koch with me. Peter works for DENIC – the ccTLD manager for DE – currently as a policy advisor. Also, we have Barry Leiba. Barry Leiba is the Senior Standards Manager at Futurewei Technologies. Barry has worked on e-mail and related technologies since the early 1980s and currently focuses on the Internet of Things, messaging and collaboration on mobile platforms, security and privacy of Internet applications, and Internet standards development and deployment. Barry is also a member of SSAC.

We will have Alyssa Moore over there. I'm sorry, I'm not pointing at people. Alyssa is the Senior Policy and Advocacy Advisor at CIRA, the Canadian Internet Registration Authority, and the ccTLD manager for CA.

Barry and Alyssa will be our moderators for the questions and answers, and the rest of our panelists are, Tim April, a Principal Security Architect at Akamai Technologies focusing on security of

the DNS networking and incident response. Tim is also a member of SSAC.

Can we advance a couple of slides, please? We will have also Vittorio Bertola. Vittorio Bertola is the Head Of Policy And Innovation for Open-Xchange, the parent company of PowerDNS and has been discussing the policy consequences of encrypted DNS in several venues throughout the last year.

And finally, we have Michele Neylon. Michele is the Founder and CEO of Blacknight Solutions, an ICANN-accredited registrar. Michele is also a member of the GNSO council.

Can we move forward a couple of slides, please? I think we have some technical problems, as always, but it's not a difficult thing. I don't know if we can start with the technical overview, with other slides. Can we? For the sake of time.

We'll just have one more minute to see if we can get the slides up. If not, then we will not, but let's give it a couple of seconds. In the meantime, you can also download the slides from the schedule to follow the conversation. I think we need to start. So please, Danny, if you don't mind.

DANNY MCPHERSON: Sure. I'm going to get started without the slides, and it was a tough enough topic probably with the slides for many people, so this will be interesting.

I'm Danny McPherson, I'm a member of the SSAC, and in this context, I am presenting the SSAC slides. So the errors are the SSAC's, not mine. We'll go with that.

Anyway, we're going to speak about a topic today, I'm going to cover the technical side, and then peter and the panel are going to discuss some of the potential implications of it. But the general topic has got a lot of interest recently, it's about what we call DOH and DOT, it's DNS over HTTPS, and DNS over TLS.

And in a nutshell, these technologies are meant to provide confidentiality of DNS transactions. Traditional DNS didn't have any notion of confidentiality built in. It actually didn't have any notion of integrity built in either, but DNSSEC was kind of bolted on to provide integrity protection to the DNS, but that still leaves the DNS open to things like surveillance, eavesdropping, and potentially manipulating responses to various clients for an array of reasons. It could be for things like parental controls, state censorship or blocking access to a malicious site to protect the user.

Anyway, there's a broad array of motivators for why we want this, and when you get to slide six in that deck, when you get it up,

you'll see some of the why. So again, the crux here is that traditional DNS and having [a notion] of confidentiality and sort of in this geopolitical state of affairs we live in and various economic consequences for things, providing confidentiality of DNS transactions has various benefits, and also some offshoots.

So we're going to talk about the deployment models, not necessarily so much about those. Anyway, if you're following along, slide six talks about basically what I just covered.

So without confidentiality of DNS transactions, you leave yourself open to information leakage or disclosure attacks, so people can look at that information and surveil you or understand where you're going on the Internet and mine that information.

So the notion of both DOH and DOT is to provide on-the-wire encryption of what's occurring so there's an on-path attack or an observer, they won't see that information. So at the very primate level, that's what DOH and DOT are talking about here.

Slide seven in the deck basically provides a traditional DNS overview, and in that slide, basically if you think about the DNS, you've got a device, and on that device you have an application, and the application wants to resolve something in the DNS and it asks a local process on the device, like your web browser might ask your iPhone operating system or your laptop, it says, "Hey, how do I get to www.example.com?"

And then your device has a DNS resolver in it that'll speak to something either on the local network or out on the ISP network. So we can go forward probably two more slides, I think, to just go to the picture, the traditional DNS illustration, slide eight, if you would.

This is what I was illustrating a moment ago. You have applications that ultimately want to resolve something for a user or a process on a device, and the application speaks to the local operating system traditionally – again, the application could be an app on your iPhone or it could be a web browser on your laptop.

That application or that web browser would ask the local operating system where that destination name is on the Internet, and then that device would go out and it would speak either to something known as a local forwarder, or potentially to a recursive resolver.

Traditionally, those recursive resolvers were in an ISP network or a local network that the network access provider provided, but more and more commonly, they may be out in a cloud infrastructure, for example an OpenDNS or Google's recursive nameserver that provides it off in the Internet somewhere in the cloud infrastructure. So that's one of the changing architectural

[parameters.] It may not be local on the network where resolutions traditionally occurred.

That recursive nameserver would go out to the authoritative infrastructure, be it the root infrastructure, top-level domain infrastructure, the authoritative infrastructure, and resolve the name and then pass that information back to the application or to the stub resolver, which would give it to the application, and then the application would be able to connect to the destination on the Internet that is desired.

As you see from this diagram, all those transactions today, there's no notion of confidentiality of those transactions, so if there's an on-path observer in any of those places where you see a green arrow in this diagram, then they could potentially see what the user's trying to resolve, and it could be business competitive information, it could be security-related information, it could be sensitive content, it could be any of a broad array of things. So what DOH and DOT are talking about is some ways to protect that information. Let's go to the next slide.

So one of the two solutions is known as DOT. DOT is DNS over TLS. TLS is what we call transport layer security, and interestingly enough, most secure transactions on the Internet, if you've seen traditionally a padlock or you go to a financial or some other website where there's sensitive information, TLS is probably the

protocol that underpins that, that provides encryption at the network and transport layer, provides encryption or confidentiality of that information so that an on-path attacker can either manipulate it or observe the information at least.

So in the DOT model – let's just go on to the next slide, I'm just going to illustrate this from this slide. Basically, in the DOT model, what's traditionally happening – again, both DOH and DOT can be deployed, an array of different techniques. This is still being developed in both the standards and the operational community. But DNS over TLS is traditionally envisioned in that a local system, for example your iPhone or your laptop, might have a systemwide setting that says I'm going to use this resolver and the infrastructure to resolve things in the DNS, and every application on that device would use that setting and go out to the infrastructure and do that.

So what you see here for example is a web browser would do the same thing it's traditionally done. It would ask the local stub resolver, it would say, "Hey, I need to get to example.com on the Internet. Would you resolve that?"

The stub resolver now however instead of sending a clear text would send it in an encrypted channel effectively to either afford or some place in the cloud infrastructure, or the ISP network to

resolve the information. And that's basically what you see with the red arrows here.

So it would be encrypted, so an on-path observer or an attacker couldn't potentially manipulate that information or at least observe what's going on. And then further along in the authoritative infrastructure here, there's really not a lot of consideration today for where a DOH or DOT solution may fit with the root or a TLD, or maybe a second-level domain authoritative infrastructure. But some of that's still being fleshed out.

We're going to contrast this in a moment to DOH, which basically moves the encryption level back a bit. So let's go on to the next slide and talk about DOH for a moment.

I'm moving too fast. Sorry. Okay, so basically, instead of using TLS for DNS transport, what DOH actually does is, you know what? I have a lot of web applications on my device or a lot of web traffic and a lot of software built around HTTP transactions in this operating system or on this device, so instead of using TLS natively in the infrastructure, I'm actually going to encode DNS responses in web queries, in HTTP queries, and then interestingly enough, put that over TLS, which is what HTTPS runs over, and then transmit it on the network.

This provides a lot of hooks for an application to either directly transact with the resolution infrastructure and completely

bypass a local stub resolver, or it may use a stub resolver run on the operating system. So let's just go on to the next slide and we'll illustrate this for everyone.

Basically, what you see – and again, it's just one deployment model. It may vary. But what could happen is my browser may use one recursive nameserver with DOH in the infrastructure, whereas another application may use a local one or it may use my system resolver.

Now, when this gets interesting is that if something breaks in that scenario and different applications using different DNS, it may be complicated to understand what's occurring.

The other thing that happens is ISPs traditionally may use DNS queries as a control point in the infrastructure, or an enterprise may use DNS queries as a control point, and they may not want encrypted transactions going directly from an application across some perimeter or boundary in the infrastructure, because they may lose visibility to security or parental controls or other things they have in that infrastructure.

So the crux of this here is that rather than using the stub resolver in the case of DOH, which was more traditionally envisioned that an application would directly query the Internet infrastructure, the resolution infrastructure to get a DNS response and bypass everything, both in the operating system and potentially at a local

network service provider. That's basically what we're illustrating here.

Okay, let's go on to the next slide. So now one of the other things that's interesting to point out is that if you're looking at this from a control point or an eavesdropping perspective or a surveillance perspective, then DOH effectively co-mingles your DNS resolution traffic with other HTTP traffic on the network. So it makes it much more difficult to potentially surveil or to eavesdrop on, or even filter, so you'd have to crack open all of that HTTP traffic to do something technically from a control point perspective with DOH-related DNS resolution queries.

DOT again is a systemwide setting, but you would also potentially have to do that, and then the last thing I think that I'll point out on this slide is that the deployment models you see for both DOH and DOT here, they could be mixed and those arrows could be flipped. It's a matter of what the application wants to enable, what the system administrator wants to enable and so forth.

A system like a stub operating system on a device might actually use DOH rather than using DOT, or maybe use traditional DNS. So I think this is still being hashed out.

Not as much thoughts given today on the authoritative infrastructure, so [inaudible] net, gov, edu, jobs, whatever the TLD is, or even the second-level domain potentially, and there's

some other techniques such as QNAME anonymization that provides some privacy protections there, although there are some offshoots with various aspects of that as well.

So I think I've tried to catch us up, speaking a little quickly as Michele kept reminding me, but we're going to pause here for a moment and see if people had questions on this before we go to Peter's deployment considerations piece of the deck. So if you have questions, you can ask those now about what we've said so far for the panel, or you can hold those and let Peter talk about some of the other implications of DOH and DOT, and then ask your questions.

ALEJANDRA REYNOSO: Thank you, Danny, for such a great and speedy explanation. We have one question here. Thank you very much.

NIGEL CASSIMIRE: Yes. Good afternoon. My name is Nigel Cassimire from the Caribbean Telecommunications Union. This is new to me, so I'm trying to understand the problem you're trying to solve with this whole thing. So, is it an attempt to make the DNS more secure? And how does this technique compare with DNSSEC for example?

MICHELE NEYLON: Okay, so now we're fighting over who gets to speak. So I'm just taking the microphone. So DNSSEC is often something that's talked about a lot within ICANN circles as if it's a magic bullet to fix all the problems with the DNS. It isn't. DNSSEC is a way of you saying when you go to your bank.whatever that you're actually going to your bank and that somebody hasn't got in the way and inserted something in the middle.

So that kind of attack is what you'd call DNS poisoning, which has been an issue in some places in the past. So DNSSEC fixes that. With the DOT, DOH, it's trying to both add a level of privacy and a level of security, but there are issues with both. In terms of the privacy, you do get that, and I think some of us will probably be talking further on about how it can have negative implications on some of the security side of things. But what it's doing essentially is moving those DNS queries, the way that the lookups are done by the devices – that's your laptop or your phone or your iPad, or something else – moving it away from the traditional DNS to piggybacking on other protocols.

So in the case of DOH? It's just seen as an ordinary web request. Bear in mind I'm not as techy as him, so he'll probably then correct me, but that's kind of a simple way of looking at it.

ALEJANDRA REYNOSO: Thank you, Michele. We have a remote participant. Please.

ARIEL LIANG: There are two remote participant questions. The first one is from Mohammed Yousif. Does DNS over TLS cause any performance degradation as to the time of resolving a query? And then we'll read a second question next.

MICHELE NEYLON: It depends on how the stub resolver is implemented. It can impose additional penalty or round trip to the resolver at the initial connection setup time, but if your stub is configured to persist the connection over time, it can be – the [amortized] cost is about the same as regular DNS.

ARIEL LIANG: There's a second question from remote. It's from Yazid Akanho from Benin. Slide six, technologies such as QNAME minimization may also be effective at preserving user privacy. How can all resolvers implement this?

DANNY MCPHERSON: We didn't want to speak much about QNAME minimization, but it's a pretty lightweight technique. Traditionally, DNS was very verbose, and if I wanted to resolve something in DNS, I would give a fully qualified domain name, so

internalsecretserver.foo.verisign.com, and I would ask every authoritative server in the path the entire question, when the reality is root only needed to tell me how to get to the next level of the hierarchy, so to dot-com.

So when I ask the root a question, I don't need to tell it everything I need, I just need to ask it how to get to dot-com, and then dot.com would tell me how to get to verisign.com and verisign.com would tell me how to get to internal secret server.

So effectively, you don't disclose the entire name of what it is that you intend to resolve, so you're minimizing this. It's a privacy-preserving name resolution function, it's very lightweight, and it is deployed and implemented already in most recursive nameserver implementations in varying ways and it provides some measurable attack surface minimization from a privacy perspective.

ALEJANDRA REYNOSO: Thank you very much. I will ask for everyone to stay as much on topic as possible. We know there are very related concepts around security and the Internet, but right now, we should try to focus on DOT and DOH so that the panel can move further along. I'll have two more questions. There's four, and there's five there, and then we pass to the next presenter. Thank you.

FRED BAKER: I was a little bit surprised by your response about the difference between the TLS and running with DNSSEC, because they secure different things. DNSSEC secures the content, the actual resource record, where TLS secures the channel.

By comparison, you might think in terms of pipes and water. Let's imagine that I've got a wonderful pipe that's armor plated, this is now the absolute world's best pipe, and upstream from it, I have a lake that I have filled with poison. When that delivers through this armor plated, very wonderful pipe, it's still delivering poison.

So securing the content gets rid of the issue of the poison. Now, I'm not going to speak against TLS. Having a good channel is a good thing too. But DNSSEC becomes very important in terms of ensuring that the name actually carries the thing you think you're trying to get.

DANNY MCPHERSON: I'll just reply to that. I think that's a great point, Fred. I think that even if you have DOH or DOT fully deployed in the ecosystem, you're still going to want both DNSSEC and QNAME anonymization to provide added protection. They address very different things, so that's a good point.

ALEJANDRA REYNOSO: Thank you. Number five?

JIM PRENDERGAST: Yes. Hi. I'm admittedly not an IETF guy. I know there's plenty of you in the room who are. Danny, as you were listing through the benefits, you also talk about some of the things that this may break. How did this get approved as a standard if there are some of these things happening that are unintended consequences?

UNIDENTIFIED MALE: [You should ask the IETF.]

UNIDENTIFIED MALE: It's a fair question, Jim.

DANNY MCPHERSON: Okay. I'm sorry. I think the technology is there and the ecosystem will adapt to figure out the right deployment models. I don't think that anyone in the ecosystem from browser vendors to operating system vendors to recursive nameserver operators or authoritative infrastructure operators want anything to break.

Interestingly enough, this does cause some deployment challenges where now if you're an ISP and you have no visibility to a user's web browser DNS traffic and they're using a cloud service somewhere else to resolve DNS, and they call you to fix a

problem for them with DNS, you may not be able to do that. Or if you implement parental controls with DNS, you may not be able to do that.

So I think the ecosystem is going to have to adjust to that, and that's why I think it's really important with both DOH and DOT to realize that the deployment models will vary and adapt, I think, as market and the dynamics there dictate what's sort of optimal, what works and what doesn't work.

ALEJANDRA REYNOSO: Yes. Thank you very much. In summary, we will have now the potential development concerns regarding DOH and DOT by Peter.

PETER KOCH: Yes. Thank you, Alejandra. My name is Peter Koch. As I was introduced, I work for DENIC as a Senior Policy Advisor, and I'm one of the ccNSO-appointed co-conspirators in this working group.

So I was invited to talk about potential deployment concerns. And the unofficial subtitle indeed, Jim, was the protocol is innocent and then things happen. This is probably what is going to address some of the concerns.

The first part will be a bit technical, so we have two standards that more or less address the same – oh, yeah, we need, sorry –

ALEJANDRA REYNOSO: [inaudible] slide, please. Thank you very much.

PETER KOCH: Yeah, that's the one. Okay, I'll do it from there. So we have two standards that in slightly different technical manner address the same issue of confidentiality of DNS traffic going back and forth.

And just to remember, people, of the why of this, there was a guy named Snowden a couple of years ago, and what he discovered, or at least shared, was that DNS traffic can be a source of intelligence, can be used to identify people or can be used to identify an action that people engage in, like visiting websites.

But let's not only focus on websites the DNS is used for. Every other service as well. So that is one of the motivating factors where the IETF – and I'm not speaking for them, but they have published documents about that – declared pervasive monitoring as a threat that will be mitigated by a couple of protocols, and these attempts are actually addressing that very issue by responding to pervasive monitoring with pervasive encryption.

So this is about encrypting the DNS traffic on the fly. It has some other aspects on the technical side that we'll get to later, but just to add, it's not only the state actors that do this, it's also other pieces in the puzzle that may have an interest to look into the DNS traffic going back and forth, because while the DNS information is mostly public, the fact that somebody is asking a particular name at a particular point in time is very likely not public and is valuable information.

That said, we have these two competing standards, and that's quite easy to describe. It just describes how the one part, the resolver ,communicates with the other part, which is in this case the so-called DOH resolver. Looks like a webserver from the outside, but instead of delivering websites, delivers DNS responses.

What is not solved so far is, how does the user, how does the web browser in that case, get the information, who to ask? Usually, this is an information that is delivered by the operating system in those cases where we have this, which is the case for most laptops and smartphones that you use. There is an operating in there and the name resolution is deeply buried in the operating system, and is usually today done consistently for the whole box that you have in your hand or in front of you on the table.

And these things might be going to change. So the IETF or the developers are still working on automatic configuration, how to find this DOH resolver, and there are also initiatives underway to give users more choice and enable them to manually configure the DNS service, but this is still in the making.

So there was already a web browser, a vendor for their web browser enabled DOH, which is DNS over HTTP, treating DNS name resolution a bit like the web, and that contains a hard coded URL – this is the identifier, the web address that you hopefully know from your web browser when you visit webpages – and to hardcode that information for DOH. So all of the web browsers from that vendor at that point in time would have used a particular DOH resolver. Again, there would not have been a single instance, there's Anycast and everything.

And that would override the information that is given by the operating system. So the application would now choose to use a different DNS resolution path than what the rest of the operating system does. That might face some challenges, and as it say on the slides, can interfere with some network managers' security policies where people try to mitigate access to certain information, mostly websites or phishing sites, you name it, by intercepting DNS traffic, and that would then, as somebody already suggested, no longer work. Next slide, please.

And then of course, the interesting question, why are there two standards? I'm trying not to dive into the technical details, but the DOT – the DNSO over TLS, over transport security, is kind of the maybe more engineering-like thing where we think of the network in layers and so on, but it had at least one problem which his you need a certain information, another hole punched into your firewall, to get to that information, whereas everybody is letting everybody to any webserver these days.

So the DOH traffic looks more or less like access to a website and cannot be separated from that, as it says on a future slide. Or on this slide, actually. Sorry.

So anybody cannot block access to this DOH-based name resolution service without at the same time blocking access to important webservers. That's the trick behind it, so to speak.

And research is still underway to add this feature to the DNS over TLS, like the DOT approach. Of course, on the technical side, there are some gory details, but they're not really part of the topic for today.

As a consequence, network managers may no longer be able to block name resolution, because at the same time, they would likely block access to websites, or to popular search engines for that matter. That can – may or may not – interfere with some regulatory requirements in some jurisdictions where ISPs are

ordered to block resolution of certain domain names, and when I'm reading that, I'm not saying that these blocking mechanisms are very effective, but they might be regulatory requirements nonetheless. And as I said, that can never be perfect because it can easily be circumvented by configuring your own resolver, using a VPN or running your own resolver on your own system.

[Masquerading] the DNS queries in the stream of the web traffic may then help users to get around DNS-based filtering, and you may call that censorship, which is when the filtering is imposed onto the user by a third party, or it could be blocking malware, which is usually a thing that the user subscribes to or that is allegedly enacted in the interest of the user. Next slide, please.

So a bit of the bigger picture, because again, the protocol is [innocent,] but then weird things happen. DNS over HTTP does not prescribe a particular deployment model. Any enterprise could run a DOH resolver and point their web browsers to that, and then act as before. However, in the discussion so far, we can observe a certain development model that really leans towards concentration and consolidation, as in web browser vendors cooperating with DNS resolution providers – and I'll refer to that in the next bullet item – and then pointing all their web browser customers, their web users, to the resolution services of a particular provider which gives the provider a lot of insight which

also gives the users probably a bit of stability, but which fosters concentration.

The DNS name resolution traditionally, over the last 30-odd years, used to be highly decentralized. That is, in ISPs or even at your laptop, or thinking 30 years ago on your mainframe, or you name it. However, DNS name resolution as a service has evolved over time, and these are the so-called quads, like the 1.1.1.1, the 8.8.8.8, you've seen that probably on a picture painted to a wall in certain countries where people were faced with DNS blocking and then circumvented this by going to one of these resolution providers. And there are others that use the same figure in every position. It's just a matter of curiosity and ease to use.

But these aspects in addition to the choice of resolution path per application rather than per system, or per the enterprise, or even per the ISP where the ISP gives their resolution choice to their customer, definitely leads to increased concentration of resolvers that are increasingly large. And with large, we mean that the population behind that resolver is growing and growing, which means that obviously, the weight – and that might also mean the policy weight – of that particular resolver operator can be expected to grow and become more and more important. Next slide, please.

Okay. So as we said, DOH and DOT both provide privacy on [the wire] and we talked about the reasons. However, the resolvers – and that is either the resolver at your ISP or the resolvers that are provided by these big resolution services – do see users' requests at a different level of detail.

For some reason, sometimes there's a particular information added to the question so that users can get tailored responses, because at the last bullet item – I'll forward here for a second – some technical scenarios depend on the fact that not everybody gets the same response when asking for the same question, so-called content delivery networks use the DNS often to direct users to the closest content delivery system to lower the latency and give users quicker responses.

So privacy is not only addressed by the encryption on the wire, but also very much by the DNS resolver policy, as in the resolving operator telling you or promising you being accountable or whatever, what happens to the query data that [that sees.] You might have circumvented the ISP or the state actor as somebody interested in your data, but probably doesn't help very much if then the resolver operator steps in. Next slide.

Didn't work. So for policy question for the DOH resolvers, interesting things to discuss and open to discussion, how should they be selected? And we had that on a previous slide, what's the

technical means, how do I as a user decide which to use and if I have decided for one, how do I configure that into my application, my system or what? And how are the operators of these DOH resolvers held accountable for what they promise and for what they do? Because again, they might be required to disclose information by the request of whatever entity, in most cases hopefully law enforcement.

And who determines which policies are acceptable – which is another discussion where one vendor has said, "Okay, we understand that there are concerns in the community that there's a single vendor that we cooperate with. We might be open to have cooperations with other vendors, but we would like them to adhere to certain resolver policies, and that means that they do and do not do certain things with the user data." Next slide, please.

So even bigger picture, because one of the questions of course would be, why are we talking about this in the ICANN context? Now, assume a group of cooperating DOH resolution providers. The group will be small, not so spread out as it was in the early days. And further assume that there is a certain application, a service that is used dominantly on the Internet, like the web, as we all do.

And there is sometimes an interest in additional name resolution path. And there was discussion in the IETF, and also ICANN, about the dot-onion TLD, which is not really a TLD but which is now reserved. But it's a different resolution path needed for something that somehow fits into the namespace.

So practically speaking, when we have this group of cooperation resolution providers -and they have a big population behind them – who would really be in a position to decide whether or not to open new segments of the namespace? How would that look like, and what would it mean for ICANN's role with regard to the DNS root zone when there is a shift in power, a shift in – what, simply people voting with their feet or their web browsers. Next slide, please, and that should be it.

Okay, so conclusions. Preliminary conclusions, probably. We've learned that some of the deployments of DOH and DOT might impact the traditional control points in the resolution. The ISP, the enterprise can intercept DNS queries and send different responses. For delivering ads, but also for mitigating malware and botnets.

Standardization of DOH and DOT, of the resolvers in the application and how to select them, is still underway. There's no final conclusion. For registry and registrar operators, currently there appears to be little impact. However, again, there could be

someone knocking at the registrar or registry's door and saying, "I'm one of these resolution providers. Why don't you just give me a copy of your DNS data? I'm really happy to deliver that to your users even faster."

And it's too early to say, of course, what the impact on users will be. And as we already heard, this is completely orthogonal to DNSSEC and other privacy mechanisms like QNAME minimization. The need for those has not changed. That should be it. Next slide, please. Okay.

ALEJANDRA REYNOSO: Thank you very much, Peter.

PETER KOCH: Thank you.

ALEJANDRA REYNOSO: Any questions from the audience? We have time for a couple of questions. I have number five.

WARREN KUMARI: Hi. Warren Kumari, I work for Google. I'm not on the Chrome team, but I'm relaying some stuff from them. So Peter, in your presentation, you provided one sort of deployment model by one of the browsers. That's not the only deployment model. What

Chrome is planning on doing is that it's going to offer DOH to its users in two ways.

If the user's system resolver already supports DOH, Chrome will simply upgrade to use that. If you're already using your ISP's resolvers, and it turns out that they support DOH, it'll just do DOH over those.

If the users want, they can choose a different resolver. That's very much the same as what currently happens. If users aren't happy with their ISP's resolver, they can choose a different one.

Whatever the case, Chrome is not going to change what the users selects without them opting in. And also, we don't require that users actually choose something like Google public DNS.

What all of this means is the existing protections that people have for stuff like malware, if they do enterprise-type DNS, all of those sorts of things continue to work in the same way.

So I think what's sort of worth realizing is the way that this gets deployed is the important thing, not actually what the transport itself is. Not sure if you have any thoughts on that.

PETER KOCH: Yeah. Thanks, Warren. We deliberately did not have names on the slides, and I hope I didn't mention any. So thank you for doing

that in that particular case. And yes, it's a valuable addition to the scenario of deployments model. We did not strive to be exhaustive there, and I think the slide said that multiple models are discussed, and the one you showed obviously is part of them.

For the other part, I would like to defer to the panel later not to preempt that discussion, and maybe we can focus on immediate questions on things I did in an incomprehensible way.

ALEJANDRA REYNOSO: Thank you, Peter. We have one remote participant, and then we will go with three and I will go with four, and that will end the queue for this part.

ARIEL LIANG: There's a question from Dirk Jumpertz. DOH is already being abused as an attack vector to insert malicious content in webpages through the use of crafted [TXT] records. Are you aware of this? This is extremely hard to block as it uses a trusted channel to attack [on THHP] as combined with DNS.

Doesn't this make DOH a threat rather than a blessing?

PETER KOCH: That's interesting information, Dirk. Personally I wasn't aware of that. Maybe the rest of the panel is. I'd like to defer that question to the panel session, and maybe we can keep that in mind.

ALEJANDRA REYNOSO: Will do. Number three?

EDUARDO DIAZ: I have a question. A potential user problem is if I download an application, and that application by [loading it] is using their own resolvers without me knowing it, so it has a very – can affect a user very easily without knowing what's happening in the background. Thank you.

PETER KOCH: Yes. Thank you for that remark. One aspect that wasn't mentioned is that in theory – and as we learned, theory immediately becomes practice – different application could deliver different results, so that the domain name system looks differently from a web browser and, say, for a mail application or for your VOIP phone, because depending on which resolution path you choose, some of the domains might be blocked and others path through, or even you might be sent to point A there and to point B elsewhere. So that could indeed be a user

experience. But this is the very beginning of how the end user might be actually expected. Thank you.

ALEJANDRA REYNOSO: Thank you. Number four.

[REMMY NWEKE:] Thank you. My name is [Nweke Remmy]. I'm from Nigeria representing NCUC, Noncommercial User Constituency. My concern is on one of the points on the slide, it is too early to say what the impacts of DOH and DOT might have on the user, but at least we can still try and look at the potential impact, the negative impact these might have users.

Another thing I would like us to clarify is what are the countermeasure that we could use on these negative impacts of DOT or DOH, and also, what are the responsibility of the user, cost implication, not to the technical side now but on the user. Thank you.

PETER KOCH: Okay. I think that's a contribution rather than an immediate question which can feed into the discussion. We do have two more slides before the panel opens. So I don't see any more questions.

ALEJANDRA REYNOSO: No, we will go to the panel now. So we can move a couple of slides, please. Thank you very much. So now the panelists will answer these questions that you see in front of you. I will read them all out.

Do you foresee any impact from deploying DOH and/or DOT on your operations?

Are there any issues with DOH/DOT that fall within ICANN's mission?

How do you think DOH should be implemented in applications such as web browsers?

What concerns do you have about DOH and/or DOT?

So we will start with Tim.

TIM APRIL: It would take forever to go through all of those questions, but the one that sticks out most in my mind, being from a security background, is what concerns do I have about DOH and DOT, and that mostly applies to the end users and how their perception of the namespace may change.

Basically, if the first mile from the browser or the application up through the resolver is using DOH or DOT, you get the privacy of

the channel there, but you don't necessarily have a guarantee that the connections going forward from that resolver to the authorities have any sort of protection at all.

So if you're concerned about leaking data through a communication channel, that may happen beyond the resolver, and in some cases may be attributable to the end users as well.

There's also the debugging problem of if you're using – depending on the implementation of the – specifically in DOH, if your application is using a DOH resolver without your knowledge, you may attribute some resolution problem to your ISP's resolver and call your ISP, and they're going to have no idea what's going on, and it's just going to be a long debugging problem that is opaque to the end user unless they have a strong technical background and know what to go look for.

I'll let others keep going.

ALEJANDRA REYNOSO: Okay. Vittorio? If possible.

VITTORIO BERTOLA: I think I have several things to say. Starting from the first question, as a software vendor and DNS service provider for some of the biggest ISPs of course, we have an impact in terms of

implementing the new protocol and making it work in the real world, in [platforms that serve several ] million DNS queries per second, but that's not the real problem.

We are more concerned as software companies and open source company with the problem of openness of the Internet and the impact that this could have in the shape of the DNS resolution market and service in general.

So I think that the real issue here is not with the encryption, so it's not about the transit through an encrypted connection, which is fine for the privacy. The added move of the DNS and changing it from a network service, something which is provided as part of the network service by your operating system, like the TCP IP stack, to an application service, something which is directly managed by each application.

This opens the way to a number of issues. Part of them are related to potential confusion as we were saying, different applications going different ways. But the most concerning one is about the fact that the application market, especially if we take the web, which is by far the most widely used application, is much more concentrated than the network market.

Currently, if you want to put together 95% of the world's DNS queries, you have to put together the top 1000 DNS resolvers. In

the web, that's just for companies, basically. All of them, by the way, are in the same jurisdiction in the same country.

So in terms of potential policy impact, especially in terms of jurisdiction, sovereignty and all these issues, this changes a lot, because we all know that for governments, I think several constituencies are affected. One is the ISPs, but I think in this context, maybe it's [worth ] speaking about governments and end users.

For governments, the impact, the issue is really about losing control of the DNS resolution, and especially for countries that have decided to use these either to provide additional services like parental control, or also to apply any kind of content control and filtering towards their citizens can see.

In the end, what happens is that web browsers could just start using these global platforms, and all this would go away and all the control would shift somewhere else which is not under the jurisdiction of the country. So this is why at least the British government has been putting some attention, and I expect several more governments to do this.

And for users, this is a potential issue with choice, because if the application starts deciding – either just sending DNS queries to whatever party they want, but also even just limiting the choice of saying, "We now [inaudible] are the people who decide who

can run a resolver, and there's just this list of 10 resolvers globally that we have accredited, and everything else we will not allow in it."

Then they become gatekeepers and they decide the policies for the DNS resolution. So in the end, this depends on policy, so the last message I want to leave is that this really depends on the deployment model, but to agree on a deployment model, there is the need for some shared policy, either bottom-up or so that the application people don't just go and do whatever they want but there's a shared understanding of what is going to happen.

ALEJANDRA REYNOSO: Thank you, Vittorio. Michele?

MICHELE NEYLON: Thanks. I think the questions that we're looking at here are far from simple. They're the kind of questions that everybody loves, the hard ones. And I think some of this stuff is quite theoretical and academic, whereas at the moment, it's very early days. DOH, DOT, up until recently, were hypothetical. Now they're becoming reality.

And what is that reality? How is that going to impact us? The second question. ICANN's mission could be impacted in some ways with this if you end up in a situation where the public

identifiers are no longer public. You're now ending up potentially in a situation where a much smaller number of DNS resolver operators are now deciding what is on the DNS, what people can actually reach, what they can get to. So I think that has some potential impact.

My own company is a hosting provider, a registrar, and we also run ISP services. People don't understand what a domain name is. They don't understand the difference between a domain name and a browser, they don't understand the difference between a search engine and a browser address bar.

So when somebody says, "Oh, the user can choose to change which service they use for that," That might be true if you're talking to a bunch of hardcore geeks. How many people in this room run their own nameserver? Okay. And I look around the room and I know these are all the hardcore geeks.

How many of you run your own mail servers? It's the same group, by the way, mostly. Now, would any of you honestly, hand on heart, say that you're a typical Internet user? Okay.

This is the thing. It's like to say that there is choice is not entirely true. And ultimately, the kind of policy and technical aspects of this, it is opening up pandora's box in some respects. But why did we get here? If you go back to the presentation from Danny and you look at the comparisons of the different technologies, the

question to ask is, why did this happen? How did this – where did this come from?

And the reality is that in the world we live in now, privacy and security are things that people are concerned about. If you're not concerned about privacy and security, where have you been for the last few years? The DNS was too public in many respects. It had a bunch of interesting issues.

So now we have a potential path to fixing some of those issues, which opens up a new set of issues, which of course will keep some of us in jobs for the rest of our lives.

From an operational perspective, I'm not sure how on earth I am going to explain to some of my customers why certain things don't work, because it's bad enough now when they ring you and they tell you that they're having problems with Outlook when they're actually not using Outlook because they think that Outlook is the only e-mail client or that Firefox is the only browser.

So I think there's some interesting operational issues we're going to have to deal with, and if you look at what's going to happen over the next few months as this becomes live in a small set of browsers and potentially other applications, you're going to have those security concerns, you're going to have – people are already finding new and interesting ways to exploit the new technology.

They're using TXT records in DNS to spread malware. I saw a presentation on that a couple of weeks ago and I was going, "Wow, that is incredibly scary, but why the hell didn't I think of that?" Sorry, I am joking sort of.

But I think it's something we're all going to have to look at very closely. I personally have a lot of concerns about the idea of handing off that control, that decision. I look at it in terms of my own office network, will we be able to protect our own staff from malware and various other types of attacks? Do we have the technology to do this?

And I suspect the answer is no. But is this a fundamentally bad thing? I think the answer is no it's not, but it's going to have to evolve.

ALEJANDRA REYNOSO: Thank you very much, Michele. And now, again, the floor is open for questions for panelists. I have number six, and number five after. Okay, please, number six.

MILTON MUELLER: Hello. The word "consolidation and concentration" arises in many discussions of DOH. Not really DOT as far as I understand. But somebody who deals with economic analysis, these have very specific meanings. Concentration and consolidation are bad

because they may convey monopoly power, pricing power to the supplier.

My understanding is that most of these DNS services that people are using now that are not concentrated, that are distributed, they're not paying for it at all; is that right? And is the concern that this concentration will lead to some form of monopoly pricing with DNS services, or is it some other concern? Could you specify more accurately what that concern is and how the overall market for Internet services would be affected?

ALEJANDRA REYNOSO: Anyone?

VITTORIO BERTOLA: I don't think it's a pricing concern, because today, you get your DNS from your ISP and it's part of the Internet access service you buy. It's more in terms of concentrating information and control.

So for example, this is a protocol, by [inaudible] promote privacy, but if in the end like 60% of the world is using the same resolver, yeah, that resolver will get to see the information about the browser [attituded] of 60% of the world, so potentially, that will be a big loss for privacy in the end.

ALEJANDRA REYNOSO: Michele?

MICHELE NEYLON: Thanks. I think, Milton, as Vittorio says, it's not to do with price, it's more to do with that the Internet works because it's distributed. It's a network of networks. Every ISP can set up their own resolvers, every network, we can all have our own resolvers on that. if you concentrate it, then you lose that stability, that resiliency. There's potential for that resiliency to go away.

And also you have the issue where there's a crazy amount of data in DNS traffic, not just of what's there but what's not there. So what people are trying to reach actually doesn't exist. That is worth a lot of money.

ALEJANDRA REYNOSO: Thank you very much. Now five.

UNIDENTIFIED MALE: There's a confusion which needs clarity, is when you talk about DNS at the browser level, that's the user level, and who applies the policies then? How do we bring clarity at the policy level? That's the question. Yeah.

TIM APRIL: So I think you're asking who gets to mandate or define the policy that is implemented in the browser. Is that what you're asking? That's up entirely to the browser maker and any of their users that provide them feedback that they actually choose to implement. There's no policy mechanism to enforce them to do anything. It's their software, they can do whatever they want, essentially.

ALEJANDRA REYNOSO: Thank you. Number three.

EDUARDO DIAZ: Thank you so much. Is it possible that if I become a big company with a big resolver, then I can start selling or offering top-level domains without going to ICANN? And then the other resolvers can contact me if they don't find their root, right? Can that be done? Is it possible for that to happen?

TIM APRIL: It is technically possible. There's nothing prohibiting that.

DANNY MCPHERSON: I don't think DOH or DOT changes that at all.

TIM APRIL: That's very similar to how dot-onion was stood up.

ALEJANDRA REYNOSO: Thank you. Number four.

FRED BAKER: The thing I find myself concerned about is Internet routing. The case I'm talking about, you're probably very familiar with. It's a national entity, but I'll try to avoid mentioning the name.

What the issue is it's often also an enterprise issue. Companies will impose information security models, and they'll do it in part by denying access to certain sets of names in one way or another.

Now, the entity that I'm thinking about, people in that place decided to start using the Google resolver, and that got circumvented, that characteristic got circumvented by the company in question hijacking the route to the Google resolver.

And at the point where a security solution becomes hijacking routing, as a routing person I get really worried. I'd be interested in your comments on that.

ALEJANDRA REYNOSO: Danny?

DANNY MCPHERSON: I'll just say, yeah, the routing system is a web of trust, and the way it works is you refer to it as routing by rumor, and you choose to believe what someone tells you and propagate it, or not, and there's no central authority today. There are some techniques like RPKI and other things that anyone that's operating any critical infrastructure services ought to be using those and other techniques to better secure the routing system. But I agree, I think the underlying routing system is probably one of the largest security concerns on the internet today certainly, and every service is captive to that until we button it up a bit more.

TIM APRIL: And there's also the case of – this is a great chance to say that people should be considering using DNSSEC and [DANE] to do [inaudible] for any of the resolvers they're using there so that when the device or the resolver, or whichever piece is actually making the request, tries to contact the server, it can validate its certificate through the DNS and using DNSSEC actually validate that it's actually the server it thinks it's talking to, where if you're in an area that has access to a key that is trusted by your [cert store,] then you can't just rely on an X.509 certificate checking through your chain of trust, as you can't really trust it at that point.

ALEJANDRA REYNOSO: Thank you very much. Number six.

MARK SVANCAREK: On the questions to the audience, what are your concerns, a concern that comes up a lot is the concentration of the DNS providers. And why is that not the question on here as opposed to these protocols? If the most popular browser by default – and this is my understanding, is by default – will go to 8.8.8.8, you have this gigantic concentration already regardless of these new protocols. Why is that not the issue that we're most concerned about? This just simply accelerates that trend.

So I would think that that would be one of the bullets on here in addition to these protocol questions. Thank you.

TIM APRIL: I'll beat Warren to the punch. Warren was saying that Chrome is not going to select 8.8.8.8 by default. It's only something that can be selected. There's another browser –

MARK SVANCAREK: [Warren said that it was only for DOH.] [inaudible]

ALEJANDRA REYNOSO: Please use the microphone.

TIM APRIL: Warren can correct me if –

MARK SVANCAREK: Sorry if I misrepresented Warren.

TIM APRIL: Warren can correct me if I'm wrong on this, but I believe the way that Chrome is planning to implement this is by default if your configured resolver, so your system resolver, accepts DOH, it will use that as the resolution pack. If it doesn't, it will fall back to the system resolver and then the user can select to use DOH to whichever resolver it chooses that supports it. So you could select 8.8.8.8. It may be a preconfigured option in the dropdown that you can select, but it's not going to be on by default in Chrome.

VITTORIO BERTOLA: If I may add something a little more general, yes, you're right – many of the issues in terms of security, privacy, sovereignty that [inaudible] we have been talking about already exist today when a user goes there and enters one of the for example 8.8.8.8 servers rather than the default one they get from their network.

The point is that this is really making this the default, so it's making it much easier for the browser to basically switch the

people from the local resolver to the biggest resolver, and in the topics – I agree that any kind of concentration is a concern already. And I'm all happy that Google is saying that they're not adopting this kind of deployment model now, but of course, what happens in five, ten years or whatever?

ALEJANDRA REYNOSO: Thank you very much. Number five.

ROBERTO GAETANO. Thank you very much. I'm old enough to have witnessed the times when network software was a bunch of proprietary solutions that were doing a little bit of everything in different parts. And that then we gave birth to an architecture that was in seven layers, with a transport layer, the physical layer and so on, and all those seven.

And the result was the possibility of having open software and to have solutions that could for each layer compete to each other. And now with this sort of approach of DOH and DOT, aren't we going back to the proprietary solution, limiting the possibility of having competing solutions and going back to the thing in the [60s ]that was called improperly – especially for me as an Italian – spaghetti code. Thank you.

DANNY MCPHERSON: Yeah, I think that's a fair point. I think that this is still using the TCP IP stack and the layering model, it's just going over the top name resolution rather than using the stub resolver in the local system.

Now, certainly, if you see diffusion in resolution paths on the local system or they circumvent that altogether, then that has implications on the user and on the network operator and on the infrastructure, and there are various parties that may benefit and various parties that may lose from that in certain ways.

So on one hand, I see your point. On the other, I don't think that it's diverging from that, I just think that if you own the end user applications and you have the ability to tether directly to the name resolution infrastructure you want, then you can kind of see both sides of that transaction, and it may influence network operators, as Warren and the Google folks pointed out, to upgrade their resolvers to support these new capabilities, and other times, a user may be more captive to that name resolution infrastructure and that service provider than they actually realize, and that can be problematic.

So I think that's a fair point from that perspective.

PETER KOCH: Roberto, both you and the previous speaker avoided to say that this is kind of silo building and going back, of course, in the bigger picture. But this is a trend that has not immediately to do with the standardization of these protocols. And as I said, they are probably very innocent. But it's following a general trend.

Most of you have lots of apps on their smartphones, and there have been long discussions about what that means for standardization and then also of course for the use of central infrastructure, because when I use apps that just call home, what do I [inaudible] standards for other than the HTTPs level and everything is done on there, I can do by myself?

That's part of a bigger picture. It's not the only thing, but of course, it's another trend, and this is concerning the very infrastructure that ICANN deals with, which is one of the major reasons to bring this to the audience here.

ALEJANDRA REYNOSO: Thank you. Number three.

JÖRG SCHWEIGER I take it that the conclusion has been drawn that whether or not DOH is good or evil depends on the deployment model, but I'm wondering if that is really true. And if only the user would have choice, then it would be beneficial to use DOH. But take into

consideration that the user downloads an app, then the resolution path will be buried deep within the application.

So there currently is no choice, and if that app store would belong to a major player, then certainly, there's no choice. So, is it really all about the deployment model?

VITTORIO BERTOLA: This has also been a hot discussion at the IETF, because of course, there's been some discussion on how far this is an issue of the protocol and how far this is an issue of the way people use it.

I think the most important thing is anyway that we understand if and how there can be a discussion that involves all the stakeholders on the proper deployment model. Because in the end, if for example the applications were required to let the user choose or even to use the default that the user has configured in the device, in the operating system, as their default, and if they did this as a rule, then most of the problems would at least start going away.

But the point is, how can we have that discussion? Since there are very few people from browser makers here, and maybe the same companies but not the right people that make the browsers. So how can we engage these people in a policy discussion?

ALEJANDRA REYNOSO: Thank you. We have two remote questions, please.

ARIEL LIANG: The first remote question is from Christopher Wilkinson. Concentration has been a global issue for 20 years. Root servers, name servers, [ISPs,] DNS, etc.

Why would we be going now in the opposite direction? Where will these resolvers be located [cause insecurity?]

VITTORIO BERTOLA: One point I would like to make is that the resolver platforms maybe can be distributed and you can have a server in each country. But if the company still is in a specific place of business in a specific jurisdiction, they will always be subject to that. So I agree with the concerns that Christopher has made.

DANNY MCPHERSON: I would just add that I think the root server system, and maybe some registries are probably the most widely distributed resolution and Internet service systems in the world today both geographically and from a resolution perspective. So I do think that if – [done] some work in the past on what we called hypergiants where 20 or so Internet entities made up about 80% of all Internet traffic and destinations, and certainly, if those are

the entities that are operating this and ISPs and other people don't opt to protect the confidentiality of resolution data, then those entities may see more traffic, and that'll certainly cause jurisdictional and other issues. But I think that natural economics and capitalism are going to help address and t ease that out over time. This is a very nascent technology we're talking about. So I think that it's got a ways to go yet.

ALEJANDRA REYNOSO: Thank you. Second question?

ARIEL LIANG: Second question is from Mike Bagley. Doesn't DOH allow better bypassing of DNS-based security systems and also stop adblockers? Isn't that increasing security risk?

MICHELE NEYLON: Short answer, yes.

DANNY MCPHERSON: I'll pile on to that. I made a point of comingling in my slides, and what I meant by that is if your DNS resolution occurs in the same protocol to the same destinations, and it's at the application layer where web traffic occurs, then anyone that wants to manipulate

that in some manner would certainly have to do a bit more work to tease it out and find what it is that you want to manipulate.

And quite frankly, that is one of the issues here, is that some people manipulate DNS responses today, and if you're a browser vendor or an in-system or an app operator and you can keep people from manipulating responses, then you can measurably impact economics of things.

So I think that there are going to be winners and losers in that as well, and so I think that security systems are going to have to step up, and you may actually either block these protocols wholesale in an enterprise, which is probably what many enterprises will do, or you're going to want to proxy them. You're probably not going to allow these things to resolve natively over the top in very controlled environments, or even particularly from a sovereignty perspective, and that could be problematic for the ecosystem.

ALEJANDRA REYNOSO: Number four.

KAVOUSS ARASTEH: Thank you very much. I have to make some comments instead of asking questions. Michele, thank you very much. You said that how many of us understand what is DNS and how it works. I can't

have answer to that, because we could not have any statistic, I could not talk on behalf of anybody. So that is your statement.

Then you said that, are we concerned about security? The answer is yes.

Are we concerned about privacy? The answer is yes.

Are we [inaudible] technology? The answer is yes.

But for some of us – not many of us – these are the new issues. We have to digest that. We have to understand that. Before answering any of these questions, we need to see how it works and whether it replies or responds to the issue of security and privacy. So it is living questions or topics, we have to follow that, and it's difficult to answer any of these questions, even the question number two which is directly related to ICANN mission. Maybe we have more question to add to that one, or maybe [just] that one. Anyway, we need time. Thank you very much.

MICHELE NEYLON: Thanks, Kavouss. For once, we're actually in agreement. This doesn't happen that often. I think the thing with something like this is it is very new, and I think several of us have made reference to it being new, nascent technology.

The thing, I think, which a lot of us have been trying to do is trying to encourage people in various parts of the ecosystem to start asking those questions, to ask the simple questions, the more complex questions, the really hard questions just at kind of a theoretical level, then also talking to the companies that actually are already deploying these technologies.

And they're going to say, “Oh, it’s fine, it’s great. What we’re doing is for the greater good.” But unless you actually put them under the microscope, then you don’t know that it’s going to be that way forever. Something which might start out as being innocent could become something else. Or maybe it will remain innocent.

So I think it is something that we’ll need to look at, and engaging here with some of you in the room, maybe engage outside this room, and start continuing that conversation, because this has been something that’s been discussed in IETF and some of the tech circles going back, what, three, four years? Maybe longer. It started off with a simple, “How can we make the DNS more private?” And then it morphed and morphed. But a lot of the people in this room who aren't in the IETF space, the hardcore geek space, weren’t really looking at it, and now it’s becoming a reality, and I think it’s high time we start having those conversations.

ALEJANDRA REYNOSO: Thank you very much. Number five.

ANDY BATES: Hi. Thank you. Andy Bates from the Global Cyber Alliance. We're one of the cofounders of 9.9.9.9, so I find it a refreshing debate around consolidation. I guess the question to the panel is that I think we don't want to users just to stay with normal DNS, so whether you use any of the quads or any of the solutions, I think the point is that that's giving the users genuine protection from cybercrime.

So I think the throwing the question at this, that, would you like consolidation or cybercrime? There aren't any real other choices. But I welcome your opinion, please.

VITTORIO BERTOLA: Encrypting the path is a positive thing I think we should be doing. The advice we'd give to operators [inaudible] is to deploy DOH. At the same time, if you solve some of the problems in terms of privacy and security, but then create other privacy and security issues which can be even bigger, then you have not really made an advance.

So I think the positive way out of this is to have a shared understanding of what's happening and a shared policy on how

to do it so that you maximize the positives and address the negatives.

ALEJANDRA REYNOSO: Thank you very much. Number three?

WOLFGANG KLEINWAECHTER: Thank you very much. This is the GAC room here, and you mentioned already one of the slides, this will have some implication for national regulatory frameworks. Do you see a role for governments here, or did you already get some comments from law enforcement agencies?

MICHELE NEYLON: Wolfgang, none of us – we're not in charge of this. We're a bunch of people who were asked to talk about this for a variety of different reasons, but if you want to ask that question, don't ask us I think is probably the best thing to say.

Vittorio will now of course contradict me, but that's his role on this panel. Peter will contradict me.

PETER KOCH: For the first time ever, Michele. So yeah, excellent question. I think there are some aspects to this whole blocking debate, and there

are governments who believe in DNS blocking and that it would prevent the dedicated consumer of some content to reach it.

We know that these methods are easily circumvented. However, on the other side, where the domain name resolution is used to prevent the accidental access to content or whatever, talking about malware, phishing and so on and so forth, or [inaudible] contacting botnet command and control systems, that might work, but nothing says that the resolution providers – some of them already today offer some certain services, as in DNS protection, or I'm not sure that it's branded, but it's the DNS firewalls and so on and so forth. They are around in the wild. You can actually go there.

And then to go back to Milton's question, yes, and some of them charge money, the others charge data, which is a different topic. But some of them charge money for you to get to their resolution service that actually then has blacklists for known malware and phishing sites. And there's no reason why this would not be deployed by at least some of the providers that we were talking about. So in that sense, not all the stories about the regulation and that it can be easily circumvented should be taken for granted. There are some difficulties in there and some details.

And if you believe in DNS blocking, you might believe in DNS blocking even with DNS over HTTPS.

VITTORIO BERTOLA: I just wanted to add one thing. Personally, I don't particularly like DNS blocking, but I think it's really important that the decision on whether to have blocked content or not is taken democratically by each country with its own Internet community, and is not taken by the Internet companies and browser makers together. So I think it's really a matter of authority in that sense. And that's what irritated me, because some of the DOH proponents went out and made interviews saying "We're going to save the world from censorship and any kind of content control is censorship even in democratic countries."

This is really something that, as a European citizen, really irritated me. And in terms of other governments – I only know of the British governments, but if there are other governments that are dealing with this, they're welcome.

ALEJANDRA REYNOSO: Thank you very much. We go with number four, and then we close the queue with remote participants.

SÉBASTIEN BACHOLLET: Thank you. I will speak in French because we have the interpretation tools, and we have qualified interpreters in the room. So I am an individual end user, I am a member of ALAC, and

there is a certain number of questions that I wanted to ask. But I already have the answer.

I do have two more questions: what is the choice of the end user in all of this? Isn't there a risk that we find ourselves in the situation where we were a few years ago where we were with MSN, with CompuServe and so forth? Now it will be others. But someone will choose for us where resolving happens.

And my second question has to do with the fact that we are at ICANN. What are the possible consequences on naming, on root servers, and on the way that all of this will be managed in the future? Can we imagine that ICANN does not need to exist anymore? Because those resolvers, servers might decide that they add in their files new extensions, new names, or they might remove them? So block things or add things. Those questions are important, in my opinion.

And I agree with what you said earlier. It is important that we continue to work on those questions. It is too bad that standardization is done before we even were able to discuss these topics with all of the stakeholders. Thank you very much.

MICHELE NEYLON: Sébastien, thanks for the questions. I think your first question, we kind of covered already in some of the Q&A earlier on in the

session. Yes, you do have the potential of ending up where a couple of vendors are choosing what happens. And as I kind of touched on earlier on, you do have the opposite of blocking where you have the potential of addition. That is a risk.

But the protocols and standards are things that are developed by people within IETF and other standards bodies. They have discussions there. And you can follow those discussions. They are open.

Of course, the barrier, there is a technical barrier to entry. It's not for everyone. There are standards that impact our daily lives, and a lot of us wouldn't have a clue what they're talking about because it's not our area of expertise.

So people have been aware of these things and can have those discussions. I think that's why having these discussions now is valid. I don't know, anybody want to add anything else? Tim?

TIM APRIL: I was just going to add on that the technologies of DOH and DOT don't inherently make the – aren't the culprit in this case. It's up to the implementation that's really going to impact hos these decisions are made, where all this could have been done without DOH or DOT being proposed standards in the IETF. It could have been implemented by the browser vendors orthogonally.

And the main reason that it is now such a hot topic is because they are the proposed standards and there's so much discussion about enhanced – or adding this privacy mechanism to the first mile of DNS requests.

I'm sure that if roadblocks are put in the way of this sort of deployment, clever people in the IETF will continue to find other ways around those roadblocks.

ALEJANDRA REYNOSO: I'm so sorry to interrupt you, but I think we need to take the last question that is remote because we are running over time. So sorry to interrupt. We can continue the conversation in the hallway afterwards. Please, remote.

ARIEL LIANG: The remote is actually a comment from Paul Hoffman. DOT and DOH are new protocols, but applications and operating systems have been able to do something identical to them for well over 20 years.

VITTORIO BERTOLA: He's one of the authors of the DOH standard. But the point is fair.

ALEJANDRA REYNOSO: Well, thank you very much for the comment. Now [I will to round up] and finish this high-interest topic, I will ask each of you to think very quickly of something that the audience should take away from this conversation. And you can start now. Thank you.

TIM APRIL: I'll start and take the easy one. Like I was saying a second ago, DOH and DOT are two proposed standards in the IETF that aren't adding any technical capabilities to the name system that weren't already possible through non-standardized methods.

And the big concern in my mind at least is that a lot of the conversations we've been having here depend a lot on the policies and implementation details of both of these protocols in [inaudible] applications or in resolvers and authorities as we go forward with these measures.

VITTORIO BERTOLA: My message is just to continue understanding, especially if this is the first time you stumble upon this discussion, there's lots of people that are happy to help, there's already material presentations on the web. You can find stuff. But then think of how [usual] stakeholder can g et engaged with the rest of the community and contribute to the discussion either at the IETF or

some policy places that are not determined yet but that could take up the more technical and policy-oriented issues.

PETER KOCH: Yeah, I want to say that the standards coming out of the IETF are probably instrumental to the developments we see, but they are not the root cause, and we should focus on the root cause and also really get to the bigger picture, what does this mean for the ICANN and ICANN environment and the future of the governance of the namespace?

DANNY MCPHERSON: Yeah. From an operational perspective, I think that there's some overhead, but there's also some benefit from a privacy and security perspective, and understanding where and how these are deployed is going to have implementation on that.

I think that with an SSAC hat on that SSAC is just beginning to consider this, and that we certainly welcome your feedback. We're still ingesting, and it's a moving target. And these are on a standards track in the IETF. They're not full standards yet, but they're certainly on the standards track, and [I think] understanding the implications to ICANN and the policy folks, the people who participate at ICANN in particular, is hopefully where

SSAC's advice will help to address or provide some more insights for people to consider as they do their jobs at ICANN. Thank you.

MICHELE NEYLON: Alejandra, you've done something very dangerous. You've given me the last word.

ALEJANDRA REYNOSO: Go ahead.

MICHELE NEYLON: Thanks. I think there's been some very interesting questions and comments from both people here in the room and others. I think for me personally, I had some feelings about the technologies before I came in here, and listening to some of the questions we've been asked and some of the comments, my own thinking on this is still evolving. And I think that to me means we're probably on the right track in terms of actually having that conversation.

So the message to the rest of you is if you look at the slide that's up on the screen at the moment, there's a couple of points there about where you can find out more at the upcoming IETF meeting. [I think it was] dnsprivacy.org I think has a lot of information on the underlying technologies. There's a lot of blog

posts out there from a lot of different companies, and other groups like CENTR I think has put out a paper on it recently. Take the time, read some more on this, and ask questions.

ALEJANDRA REYNOSO: Thank you very much. This session is adjourned. Big applause to our panelists.

**[END OF TRANSCRIPTION]**