KATHY SCHNITT: Welcome to the SSAC Evolution of the DNS Resolution Work Party Teleconference on Thursday, the 21st of April 2022. Barry, I'll hand it back over to you.

BARRY LEIBA: Okay. Thanks. All right. Also, unfortunately, we really didn't get input in the document during the week in between, which we're still hoping we can convince people to do. We got some edits from Russ and me but not much from others. So please, over the next week or two, we'll have to decide. Yeah, I will be out next week. But please, over the next week, flesh things out in the document as we add things from today's discussion. I guess what we're going to do is go back to Sections 3 and 4 and see if there's more text we want to put in, if there are more items we want to put in, if there are things we want to take out, what we think about what we've written so far. Andrew, do you have a plan of attack for this? Or should we just try to open it up?

RUSS HOUSLEY: Well, I think we should do one other thing, in addition to that, Barry, and that is, if somebody wants to grab the pen and write a few sentences instead of one bullet, please let us know so we don't all grab the same bullet.

BARRY LEIBA: Good point. And yes, we have in the chat the pointer to this document. So, Andrew, the question was, do we just want to open it up for discussion, or do you have some things you'd like to say?

ANDREW MCCONACHIE: No, nothing to say. I mean, I think we just start at the top of Section 3, if I can get my windows in the right places. Or actually, 4 is where we have all the interesting stuff, the new stuff. So let's just start in 4. We can comment on the comments, and then add things as they come up. Barry, can you kind of walk through the comments while I type in what people say?

BARRY LEIBA: Sure. All right, so we're moving down to 4.

RUSS HOUSLEY: Start with the title, right, because Andrew made a comment there. I wondered whether we should drop the word "technologies" and "evolution" and talk just about what the domain name resolution mechanisms are without claiming whether they're forward moving or backward moving.

BARRY LEIBA: True.

SUZANNE WOOLF: We're doubling down on the nonjudgmental

RUSS HOUSLEY: Yes. That was exactly my intent. Yes.

BARRY LEIBA: I like that. So then, what do we want the title to be?

RUSS HOUSLEY: I put it in the comment there, Domain Name Resolution Mechanisms.

BARRY LEIBA: Okay.

RUSS HOUSLEY: Because the first thing we talk about in the section is why they all have a syntax of domain name.

BARRY LEIBA: Yes.

RUSS HOUSLEY: Right. Okay.

BARRY LEIBA: Okay. I'll buy that. Anybody think that's not right?

UNIDENTIFIED FEMALE: I think it makes sense to me.

SUZANNE WOOLF: Yeah. I think there needs to be text explaining that this is actually an area where people think of it as pretty static and lately it isn't, but not in the title. I'm agreeing with Russ is doubling down on that.

RUSS HOUSLEY: Are you saying there should be another bullet below applications want to consume?

SUZANNE WOOLF: Yeah. I guess that's Andrew that just typed.

ANDREW MCCONACHIE: No, I didn't.

SUZANNE WOOLF: Okay. Yeah.

ANDREW MCCONACHIE: Making a note of what you said. So we don't lose it.

SUZANNE WOOLF: Yeah. It's a one sentence bridge between. We're just talking in a nonjudgmental way about resolution mechanisms. And oh, by the way,

this might be something you're interested in because it is change, whether you think it's evolution in the forward direction or backwards, it is a fact that there's some recent—we'll be polite and say innovations in how that stuff works.

BARRY LEIBA: All right. Good, good.

SUZANNE WOOLF: I think there's a bridge. I mean, we don't want to abandon the idea entirely, just because we don't want to put it in the title is kind of what I'm thinking.

ANDREW MCCONACHIE: Change the word this to resolution mechanisms.

BARRY LEIBA: I was just making a note. I just didn't want to lose Suzanne's thought.

SUZANNE WOOLF: Thanks.

BARRY LEIBA: The bullet after that I think sets us off on one of the key points that because the syntax of domain names is built into so many things, we have that kind of locked in. But how you get that domain name into use can change, and that's where we are with these different mechanisms.

All right. Any comments? What I'm looking for here is perhaps more bullets or sub bullets that will eventually turn into paragraphs of text. If we can throw in some more thoughts, then it'll be easier for presumably, Andrew, who's got to flesh this out to write some text that we can then review or for somebody to grab a section and start fleshing it out offline.

UNIDENTIFIED FEMALE: I'm thinking, because one of the things we want to do is tell a story. What are we trying to explain with this paragraph? What is the issue at hand?

RUSS HOUSLEY: I think at the high level, what we're saying is, people want domain names because that's what applications consume regardless of how the actual resolution mechanism behind it works, and then talk about the plethora of things beyond DNS that have come about to resolve those names. Why? The why is Section 3, right? Or the pressures that cause people to do something other than DNS.

BARRY LEIBA: Yes.

UNIDENTIFIED FEMALE: So we're saying that resolution can happen through non-DNS protocol mechanisms?

RUSS HOUSLEY: We're saying it does. We're not saying it's good.

UNIDENTIFIED FEMALE: No qualification on good or bad, right?

BARRY LEIBA: Yeah. I think we're saying a few things. We're saying that there are a number of changes to how this resolution is being done. Some of it is outside DNS, some of it is within DNS but outside of the traditional method where one resolver was used by everything. Your machine used one resolver all the time, and now different applications may use different resolvers, so that's one thing. We have resolution of names outside of DNS is another. We have things like .eth and .onion, and whatever. So that's all listed there.

JACQUES LATOUR: Barry, did we say that explicitly somewhere that the multiple resolver potentially available instead of just a single stub?

BARRY LEIBA: Application-based DNS resolution sub bullets under that talk about that. Different DNS resolution in different apps, different apps choosing different resolvers, different apps choosing different mechanisms, that kind of thing.

JACQUES LATOUR: Yeah, but just in multiple resolvers on one host, right, instead of a single stub?

RUSS HOUSLEY: Right. I think his point is that there could actually be a search order. Look at this one. If it doesn't get you an answer, look at the next one and so on.

SUZANNE WOOLF: Yeah. Which is already the case.

RUSS HOUSLEY: Yeah, but the bullets don't say that.

SUZANNE WOOLF: I'm agreeing that it needs to be included.

RUSS HOUSLEY: Yeah.

UNIDENTIFIED FEMALE: Is application-based DNS resolution basically quick? Or is that just one example?

BARRY LEIBA: That's one example. It's DoH, DoT, DoQ, whatever.

UNIDENTIFIED FEMALE: All right. So those are also. Okay.

RUSS HOUSLEY: Well, they don't see in S based ones.

BARRY LEIBA: Right. Also keeping in mind that the last bullet there that says this has always been going on to some extent, but DoH made it easier that you don't need DoH to have applications do their own resolution. It's just happening more now. Comments are disappearing. I love it. Okay.

RUSS HOUSLEY: Are people happy now down to the remaining purple text? If so, I want Barry to explain what he means by the Yahoo with DMARC.

BARRY LEIBA: Okay. Well, let's get to that. I'll just jump down to that. "Maybe they don't" part is what I'm saying. I'm not saying this needs to be in the document. I'm just giving an example of maybe the deployers don't care. The idea of squatting on TLDs or using app-specific TLDs because that's what makes my app work nicely, well, that creates problems if there are name collisions, and maybe the app deployers care about that and want to avoid name collisions, or maybe they don't care about it and just want what works best for their app. And the comment about Yahoo with DMARC was exactly that that Yahoo was misusing DMARC

and didn't care because it solved a problem they had and they didn't care what that did to mailing lists.

RUSS HOUSLEY: Oh, you're talking about the mailing list thing? That was the piece I didn't get. All right. Thank you. You could mark my comment resolved.

BARRY LEIBA: Yahoo basically broke mailing lists for any—

RUSS HOUSLEY: For everybody else. Yeah.

BARRY LEIBA: Yeah. Whenever a mailing list had Yahoo subscribers, and they didn't care. We discussed this with them. I don't mean to bash Yahoo and belabor that. But the point is that when an app developer, app deployer, or whatever it is, find something that they can use that benefits them, maybe they don't care what collateral damage they have because they made their app better. That may be something we want to fold into the discussion as we write it.

ANDREW MCCONACHIE: Is this the subsection where we would want to talk about SAC113? That's the .internal document. I'm kind of wondering where that might … For application-based DNS, I'm imagining there's going to be like a

paragraph just summarizing SAC109. Do we want to do something similar for SAC113 down here?

BARRY LEIBA: If this stays in the document, then yes, this is where we do that. Saying that this whole thing of using app-specific TLDs is why we wrote SAC113.

RUSS HOUSLEY: Maybe also add .internal to the list out there.

BARRY LEIBA: Yeah. Maybe, maybe not.

RUSS HOUSLEY: Because that's what supposed—well, whatever.

BARRY LEIBA: If the ICANN Board decides to follow our advice and create something, we don't know that it's going to be .internal.

RUSS HOUSLEY: Fair.

ANDREW MCCONACHIE: Well, also these are about different protocols. So I'm kind of wondering—because SAC113 is specifically about use with the DNS protocol.

WARREN KUMARI: Yes, .internal was specifically in [inaudible]. SAC113 is specifically only for use within the DNS protocol. There's another document in DNS OP, which is OP TLD, which is supposed to be the same thing but specifically for a non-DNS protocols. So there is another.

ANDREW MCCONACHIE: So maybe this isn't the place to include something on SAC113. Go ahead.

BARRY LEIBA: The point of this bullet, at least a part of this bullet, is the idea of an app grabbing a TLD and using it as its own thing. That is what .internal is meant to be to be used for, that while they should use a sub domain of .internal instead.

WARREN KUMARI: You keep using this word TLV.

BARRY LEIBA: Something that looks like a TLD. Kind of like .mail or .corp is what I'm talking about. It may be that we just don't want to go there. Maybe

we've covered that enough and don't need to have that in this document. I was throwing—

WARREN KUMARI: I think that we really should. It's just I want to make sure that we're using the same terms.

BARRY LEIBA: Yes.

SUZANNE WOOLF: Yeah. We came to some form of consensus on that question, Warren, in 109, I think, with possible references to the RFC where we also thought that out. So it's entirely possible that there's already agreed upon terms and we just have to refresh our respective aging memories on what they were.

WARREN KUMARI: SAC109? That's the one on DNS over HTTPS and TLS.

SUZANNE WOOLF: Okay. You're right. Sorry, the .internal.

BARRY LEIBA: Okay. That's 113.

WARREN KUMARI: 113, okay. Cool.

SUZANNE WOOLF: I like Jacques's suggestion in the chat.

BARRY LEIBA: [Inaudible] to watch. Top-level ghost.

WARREN KUMARI: I mean, pseudo TLD is what we'd been using in the—

SUZANNE WOOLF: Yeah. We discovered in the IETF context that that upsets some people, but SSAC can decide to not care.

ANDREW MCCONACHIE: The SSAC use a different term in 113. It wasn't pseudo, and we had a discussion on the list about it, not using pseudo. What is the SSAC saying?

WARREN KUMARI: What alt TLD says is this.

ANDREW MCCONACHIE: The SSAC does private use TLDs.

WARREN KUMARI: Sure. But that's in the DNS, because then it is an actual. So pseudo TLD is the name that shows up on the right but it's not registered in the global DNS. So .onion is a pseudo TLD, .com is a real TLD, .internal is probably a real TLD.

SUZANNE WOOLF: Yeah. One of the things about alt is that we didn't agree on whether it was there or just that we're committing to the absence of it as a real one. Or reserve portion of the namespace that is not in fact a delegation.

ANDREW MCCONACHIE: Does anyone have a problem with the fact that I reorganized this into these two categories? Names intended to be used with the DNS protocols and names not intended to be used with the DNS protocol?

SUZANNE WOOLF: No, that's fine.

BARRY LEIBA: That looks good.

RUSS HOUSLEY: What if it's both? That's what happened with Barry's "don't care," right, about the impact.

BARRY LEIBA: Right. That's what happens when someone uses one that then becomes delegated, right, which is the whole NCAP thing is talking about.

SUZANNE WOOLF: Okay. At this point, I hope we're not rewriting our advice and maybe just put that in the bibliography.

BARRY LEIBA: Yes.

RUSS HOUSLEY: Yes. I hope that's the case, too. I just don't want this to read like it's a binary choice, or that it's static, it won't change over time.

BARRY LEIBA: Well, let's see what happens as we flesh out the text. We can decide to keep this or not and we can decide how to spin it. But I certainly agree that we should be pointing to our earlier recommendations rather than rewriting them.

I'm just having a scan through SAC113. The private use TLD is what we were using throughout it, and the only thing I can see is ad hoc usage of a TLD is the phrase that I see. Anyway.

SUZANNE WOOLF: Yeah. We can put that aside for the moment for further development.

BARRY LEIBA: Yes. How about moving down to the use of DNS as a general purpose database? Any thoughts we want to add here? Certainly one aspect of this is using TXT records for many, many, many, many different things.

GEOFF HUSTON: Well, I wanted to go there because in some ways it affects caching and performance. Because if you just use TXT records, the text query gets the lot and it caches the lot. Once you start splitting it up into resource records, you're actually delineating the queries and altering the caching properties.

BARRY LEIBA: Yes.

GEOFF HUSTON: Now, I'm not saying good or bad, but there's a big thing about should resource records be used like confetti or sparingly? And that gets into this argument pretty squarely.

SUZANNE WOOLF: Clarifying question. Did you mean any rather than text, because text isn't RR.

GEOFF HUSTON: What I mean is this is perception that we just pile all the data into text, including all those silly tokens about authentication, look at bbc.co.uk, look for the TXT record and you see what I mean. And I really do mean

the TXT record because it's so much easier than getting your own resource record, isn't it? That was meant to be sarcastic.

SUZANNE WOOLF: I was going to say. That conversation—

GEOFF HUSTON: I know. But that's what the trend is.

BARRY LEIBA: Yeah. I completely got where Geoff was going with this. This minting new RRs to do this versus doing TXT records and the implications of each.

SUZANNE WOOLF: Okay. I think I was convinced and confused by overreading a specific phrase used. So carry on. Sorry for the disruption. Well, a little sorry. But thanks for clarifying.

UNIDENTIFIED FEMALE: I find it interesting. Actually my day job, I came across somebody who wanted me to authorize using a TXT record to authenticate them to be part of a marketplace. And I'm like, "What? This is authorizing that you own the domain. Okay." So I was going to look into that a little bit more because I don't—

GEOFF HUSTON: You should dig the TXT record of bbc.co.uk just to illustrate what's going on.

UNIDENTIFIED FEMALE: Okay.

SUZANNE WOOLF: There's all sorts of scary things going on.

GEOFF HUSTON: There is and should have they all been separate resource records, I don't know.

WARREN KUMARI: I mean, it is much easier just to say, "Stick in a TXT record because anybody can do that." Whereas, if it's a different record type, a lot of people cannot actually enter anything other than a TXT record.

GEOFF HUSTON: The answer is, of course, 1608 octets and we get into PDP help, da, da, da, da, da, da.

WARREN KUMARI: Yes. But the people who run bbc.co.uk presumably don't care.

GEOFF HUSTON: They don't care how fast the DNS resolution of bbc.co.uk is. Okay, fine. I think they do. I just don't think they know what they're doing. But it's in the bullet points, were they?

BARRY LEIBA: Geoff, are there other things we should put in these bullets we're hacking out right now?

GEOFF HUSTON: There is this entire issue of fuzzy search and canonicalization. By fuzzy search, I suppose, what our meaning is there was a strong desire to turn the DNS into a search engine shared by everyone except Google, and there was a strong desire by Google to stop this at the root. The canonicalization gets right to the heart of IDNs, where instead of trying to push all the potential variants into the DNS in all potential ways of doing it, you try and canonicalize the query, and then the query mechanism only uses canonicalized name forms but they're intended to refer to all potential variants. This becomes an endless source of fun and confusion.

BARRY LEIBA: Can you try to type in a bullet or two that covers that?

GEOFF HUSTON: Yes.

BARRY LEIBA: Anonymous—

WARREN KUMARI: Yeah. I must admit, I didn't parse a lot of that. But I'm also trying to figure out why an interface refuses to auto negotiate.

BARRY LEIBA: I like anonymous Quokka. Look up what a quokka is.

WARREN KUMARI: Who's anonymous Quokka? It's a four-legged kind of like a zebra, I believe.

BARRY LEIBA: It is a short-tailed scrub wallaby.

GEOFF HUSTON: That's right. It's Australian. What's your problem? Quokkas are good.

WARREN KUMARI: No, it's not.

SUZANNE WOOLF: The problem with Australia is that it's trying to kill you.

GEOFF HUSTON: Yeah, that's right. It's probably carnivorous.

SUZANNE WOOLF: It's all carnivorous, if not actually poisonous.

WARREN KUMARI: I guess that I misheard what we were talking about because I heard quokka. Oh, now I see it. Q-U—okay. I couldn't actually see that showing up before.

RUSS HOUSLEY: He's talking about that strange Google Doc thing.

BARRY LEIBA: Yes. I always find it odd that by the cursor, it says Anonymous Quokka, but on the right side, it knows it's Geoff Huston. Why doesn't it say Geoff Huston by the cursor?

GEOFF HUSTON: Talk to Google.

BARRY LEIBA: Hey, Google.

WARREN KUMARI: That's weird.

SUZANNE WOOLF: I don't know. But I found these results on Search.

WARREN KUMARI: Oh, ha-ha. Yeah, that is odd.

SUZANNE WOOLF: I saw it on Google. It must be true.

BARRY LEIBA: Okay. This is good. Geoff is putting a whole bunch of bullets in here. I like that. Thank you.

UNIDENTIFIED FEMALE: Did we capture the CDN? So we're using DNS to figure out where to have the data sent to local caches, or whatever that was?

BARRY LEIBA: I don't think that's in any of the bullets here.

WARREN KUMARI: Maybe you're trying to say EDNS Client Subnet?

UNIDENTIFIED FEMALE: I don't know because I don't remember what that is. But I remember—was it Akamai and some other CDNs, like they would actually put some

location. They were utilizing the DNS to find out where the user was located so that they could provide the caching information closer to where the user was.

GEOFF HUSTON: Yeah. I've got that lower down.

UNIDENTIFIED FEMALE: Okay. So it is EDNS Client Subnet. Okay.

GEOFF HUSTON: Well, there may be others. That's the one that we all had the fits over. You know, variations on a theme are endless.

UNIDENTIFIED FEMALE: Yeah. They work because I remember I spent a long time looking at passive DNS data for years. Going on, look, that's interesting. Oh, look, that's interesting. Oh look, they got credit cards in here. Wow, we should tell them.

BARRY LEIBA: All right. There's a bullet about IoT. Shall we discuss that a bit? The first thing is, at the moment, IoT devices are not generally in the DNS. Yes?

JACQUES LATOUR: For example, CIRA, we're working on the IoT registry, where every IoT device that have eSIM or a secure element, so we would generate a

unique ID and assert and have a matching DNS record with [inaudible] TLSA, which means billion. If we ever get to a billion IoT device, it could be billions of identities that are in the DNS. Is that a good thing? I don't know yet.

GEOFF HUSTON: Well, it's scary. If the IEEE had had done Ethernet MAC addresses in the age of the DNS, DNS domain. And it's the same kind of thing.

JACQUES LATOUR: Yeah. Well, for sure.

BARRY LEIBA: Yes. When we start getting every light bulb on the planet in the DNS and so on.

GEOFF HUSTON: Right. How else would you implement the equivalent of a MAC address today? And the answer is, as always, the DNS.

JACQUES LATOUR: And [inaudible], that working group is facilitating that.

GEOFF HUSTON: Well, there are a whole bunch of problems the DNS facilitates and one is the gap between manufacture and deployment. I manufacture a device to be generically placed somewhere but I need to put some credentials

up to allow the deployment to, if you will, securely or in some way with some certainty handshake with the device and locate it somewhere that we all know it's this device and it's here. The DNS is an ideal way to map manufacture the deployment. It's just the solution to a whole bunch of problems that were previously just used registries. And, yes, it's always the DNS.

WARREN KUMARI: I still don't entirely understand the whole IoT DNS tie way. It's not just firmware on device knows what to connect to. And when it does that, it's already got its built in IEEE, whatever they call their whatever UID thing, and it bootstraps from that. It seems so much simpler but—

GEOFF HUSTON: I don't think we need to give the marks out of 10 for what they're doing, Warren. Someone is doing it.

WARREN KUMARI: But very, very few IoT devices are using the DNS at all other than for looking up the name of the service to connect to.

BARRY LEIBA: Right. For that today.

WARREN KUMARI: Yes.

BARRY LEIBA: The presumption is that that may change.

GEOFF HUSTON: It is a whole bunch of plug and play, like every single IoT device—or not every single—but a whole bunch require me to jump through hoops with a Wi-Fi-based station installed on the device, yada yada, yada. This is crap. It's banging the rocks together. I want to power on the device and have it figure out what it's meant to do is the naive user demand, and bootstrapping that is going to be a fascinating problem.

BARRY LEIBA: Yes.

WARREN KUMARI: But that part is the bootstrapping of the device to the Wi-Fi network. That's before you have any I can talk to anything in order to get DNS or anything like that. Your issue is a layer below where IP is and where the DNS is.

RUSS HOUSLEY: That's true, Warren. So what's the point?

WARREN KUMARI: I'm trying to understand how it's the—and IoT, therefore, they'll all show up in the DNS. I mean, other than the [inaudible] stuff.

RUSS HOUSLEY: It's your area's working groups that are doing this.

WARREN KUMARI: They're doing what? That's part of what I—

RUSS HOUSLEY: The [Netcomp] stuff, right? It reaches back to the manufacturer who redirects it to who bought it and so on. But you're right, that's all after somebody enters the SSID and password.

WARREN KUMARI: Yes. And has pretty much nothing to do with DNS is what I'm still trying to—

JACQUES LATOUR: What we're working on is before the device is on the Wi-Fi, it's a mobile IoT device, 5G, whatever. It's mobile, it's got an eSIM. And then we push a cert with a subject name that matches a domain name that as the matching cert to prove the identity. So before it's on the Wi-Fi, we send a cert, the cert has a matching DNS record to prove authenticity. And then they can be connected through radius through the Wi-Fi with that cert.

WARREN KUMARI: Yes. I see that they got stuff.

GEOFF HUSTON: I'll buy that. I'll make millions of baby webcams on that model, yeah. I don't think we need to be accurate and accurately foretell the future. But if we're looking at the points of pressure, I think that's as valid a point of pressure on the DNS as any of these others.

BARRY LEIBA: Can we start poking on some of the bullets in Section 5, the implications of all of this? Have a scan of that section and comment as necessary.

GEOFF HUSTON: There is something there in that implication to ISPs that I think is being downplayed more than it should. The more complex the DNS gets, the more expensive it gets to operate. The DNS is a hidden cost. Users don't pay for queries. Domain name publishers don't pay to have their queries answered. In fact, it becomes a burden on the ISP, and the ISP increasingly goes not paying. When we talk about the wasteland of the DNS, no one fixes it. It's all really old, crusty stuff. What we really mean is there's no economics from the ISP to actually invest in it. So ISPs are actually anti-evolution as one of the biggest forces against it. I think in some ways, the whole issue about DNS encryption from stub to recursive is actually about jumping through the ISP's natural reluctance to invest in DNS because they're not getting paid for it.

So the folks who have money, the apps, the content folks, don't really want to just pay an ISP tax to get the ISP to lift up its act. It's much more effective to solve the problem directly. This is why you get a whole

bunch of folks in the DNS from people like Google and Akamai, etc., who are busy saying, "How can I avoid paying an ISP tax and get the DNS I really want?" That's a legitimate question, but the whole issue is the evolutionary changes change the economics of the DNS and some players don't want to play because it's just a cost.

BARRY LEIBA: The bullet just below where Andrew was just typing "inexplicable outcomes lead to" and I tried to put a couple of things in there. There are other effects from—

WARREN KUMARI: Once again, I mentioned the driver behind DoH and DoT and things like that came out of the IETF. Not out of Google or Akamai. It came out the IETF in response to the Snowden stuff, and the primary drivers were privacy and user security. Like censorship resistance, privacy was DoH/DoT stuff. I know that people keep saying Google wants this thingy and has been pushing it. I don't know how else to say, "Go look at where they started and who was doing the pushing."

SUZANNE WOOLF: Yeah. The question of who's doing the pushing is always interesting because it's not like these things spring fully formed from nowhere, but there are people whose main mission, particularly in the standards area and the software, is encrypt all the things. There are companies behind it. I won't be too cynical and say mostly hardware companies. But yeah,

the sort of pervasive message is no protocol is good enough until it's been encrypted. And it's been encrypted—

WARREN KUMARI: Yeah. Thank you. I'm a bit tired of the whole narrative of Google wants to encrypt this stuff because they don't want the ISPs to monkey with it.

BARRY LEIBA: Right. So one of the points that I think we want to flesh out here is that confidentiality of the queries gives us some additional privacy. But encrypting everything, well, it's not encrypting everything. Other effects of deploying DoH are having a negative effect on privacy, and there's a balance. It's not black and white.

WARREN KUMARI: Negative effect on privacy?

BARRY LEIBA: Well, there are concerns about, for instance, the privacy issues of a large provider, not necessarily Google but somebody running a resolver, collecting aggregated information from what you're querying that they didn't use to have, and that kind of stuff. So there's protection against monitoring by external parties but you're trading that off for being monitored by an internal party, for instance, and that's some of the stuff that we might want to talk about here.

WARREN KUMARI: Okay. I'll buy that as the privacy thing. I want to make sure there wasn't going to be—and my ability to block little Johnny's browsing of sites I don't want to see. I wasn't sure if you kind of put that in privacy.

BARRY LEIBA: Well, no. That's a different issue of whether using the DNS for that purpose is a wise thing or not and how this affects it. In fact, DNS is used for filtering. Maybe having a reason to stop doing it that way and doing the filtering in other ways is a good thing. I don't know.

GEOFF HUSTON: I've put it in the chat I think my issues around this. I'm trying not to be judgmental in any way. But certainly apps have a legitimate desire to control the user experience of folks using their app. Of course, it's perfectly natural. And the more aspects of the behavior of that app they can directly drive, the more the app will behave in predictable ways that the app designers, providers, whatever, have a direct saying. This, to my mind, is an unintended benefit of the drive around encrypting a huge amount of the resolution process in the DNS. That's not we wanted to go to that target. But ah, this offers some abilities that we think are desirable. Is it worth it? I could either piggyback existing operational infrastructure or I could drive it myself. Pluses, minuses. Does that sound more acceptable to you, Warren? That's the point I was trying to make here.

WARREN KUMARI: Yes. That I'm fine with.

GEOFF HUSTON: I don't see anyone invoking the regulator so I will. As soon as you start playing with users and delivering what is called inexplicable outcomes, various folks, in most national regimes, at least tweak their ears and look, when they think that the industry is playing fast and loose with consumers, they do get involved. And when outcomes are inexplicable—this is happening and I can't fix it—sooner or later, if the industry isn't fixing it, the regulator jumps in and says, "Whoa, you must do something better."

SUZANNE WOOLF: Yeah, that's fair. The regulator is always trying to fix the previous generation of problems.

GEOFF HUSTON: Well, of course, but ultimately, when you do create random outcomes in a system that's meant to be better than that, lots of other folks get involved because users, voters, whatever, get confused. They're being—yeah, you know.

SUZANNE WOOLF: Oh, yeah. We don't want people calling their MP, Congressman, local government just because they can't figure out how to find their favorite content on YouTube.

BARRY LEIBA: Or because when they go to YouTube, they're suddenly getting content from the Czech Republic instead of content from the United States because of some quirk in the DNS resolution that's going on.

RUSS HOUSLEY: I thought it was because they were using the VPN.

BARRY LEIBA: Some of the effects you get when you use VPNs, you could now get from using DNS resolvers that are located in odd places.

RUSS HOUSLEY: Sometimes on purpose and sometimes surprise, right?

BARRY LEIBA: Exactly.

RUSS HOUSLEY: And the naive user can't tell the difference.

BARRY LEIBA: Exactly.

WARREN KUMARI: That's basically what EDNS Client Subnet was supposed to help within some set of banners.

GEOFF HUSTON: I read through the IESG comments, Warren, and I never saw anything that said anything other than, "Oh my God. If we'd had our chance at not being profit a solution, we would not have done it this way." I kind of agree with the IESG of the day. The problem that we're trying to solve was real. The mechanism of solution was going, "Oh, did you really want to do that?"

WARREN KUMARI: Except nobody has been able to propose anything which solves the set of requirements which were—anyway, that's a separate—

SUZANNE WOOLF: Yeah, separate rant.

GEOFF HUSTON: Yeah. Separate rant. It's an evolutionary pressure, Warren, no matter which way you cut it.

WARREN KUMARI: Yep. Yep, yep.

BARRY LEIBA: All right. So I think what I'd like to leave people with, let's be more specific rather than saying, "Go edit some of this." Let's be more specific. These bullets that I've highlighted in particular, "Increased

complexity in DNS, many minor changes taken an aggregate and SSR impacts on the namespace," in between now and next week, can people have a look at those three in particular and beat them out a little bit? Put some more meat into those bullets or add more bullets, that sort of thing? What does everybody think? Can you find an hour to spend on that and be an Anonymous Quokka?

RUSS HOUSLEY: You said three but you highlighted four.

BARRY LEIBA: Yeah, because the loss of operator confidence is in the middle of it.

RUSS HOUSLEY: I see.

BARRY LEIBA: If somebody wants to do that one too, that's fine. But the three I said are the ones that I think we will be most interested in fleshing out right now. So we're coming up toward the top of the hour, and so I figure let's leave this as something for us to work on in between now and next week. Russ, are you good with running the solo next week because I'm out?

RUSS HOUSLEY: Yes, I don't have a problem.

BARRY LEIBA: Okay. Andrew, any wrap-up you want to do today?

ANDREW MCCONACHIE: No. I think the only action item is for work party members to comment on these four bullets.

BARRY LEIBA: Okay. Well, again, thanks, everybody. I'm glad we started when we did because nobody else came on after we started, that I could tell. So I'm glad we didn't wait for non-existent people. Yeah, I think I'm going to start doing that. We start at the top of the hour and we respect the time of the people who were there promptly. So thanks, everybody. Russ will talk to you next week, and I'll talk to you the week after.

SUZANNE WOOLF: Thanks guys.

RUSS HOUSLEY: Thank you all. This was a highly energetic meeting.

BARRY LEIBA: Yes. I'm very happy.

[END OF TRANSCRIPTION]