UNIDENTIFIED FEMALE: Welcome to the SSAC Evolution of DNS Resolution Work Party teleconference on Thursday the 19th of January 2023. And back over to you, Russ.

RUSS MUNDY: Well, we're going to start by reviewing Andrew's homework.

ANDREW MCCONACHIE: Well, the dog ate it. I guess, I can't use that excuse, right? Dogs can't eat Google Docs.

BARRY LEIBA: Yeah, life has changed so much.

ANDREW MCCONACHIE: So my first assignment—

ROD RASMUSSEN: You can if you use Fido.

ANDREW MCCONACHIE: I used to use Fido. Oh, yeah. So I had to take the Importance of Referential Integrity document and put it somewhere in the document. And I made a choice to put it right before where we go into the perspective section. The thing about this new Section 6 or what we used

to call the Importance of Referential Integrity, is it's a whole flow on its own which kind of starts off with talking about the problem and then goes into trends and new stuff and then finally ends with some of the consequences of that. So it's kind of its own thing.

Me and Barry and Russ were talking before this. We certainly don't want to break it up, right? Well, for me at least, having it here makes sense but I do think this is a good lead in to the next thing we probably want to talk about which is, the overall flow of the document and how we want to organize it.

Let me just go up here and I'll put the link in the chat. But this is how the document is currently organized. One opinion I have is that Pressures to Evolve DNS Resolution is probably in the wrong place and we probably want to put that further down. Maybe between 7 and 8 or even between 8 and 9. If you have any opinions about the general flow of the document or if you have opinions about where I stuck that new section, please voice them.

RUSS MUNDY: Well, I think the reason it's the way it is now, that is, without Section 6, was the pressures to evolve was really a euphemism for why are alternate naming systems emerging? In other words, what features does the DNS not provide that these other people want? I think we need that before we talk about alternate naming systems because they're the result. Regardless of what we call it, I think those two do need to be interrelated.

ANDREW MCCONACHIE: Maybe we just want to rename Section 4 then. But why are we doing this or something?

RUSS MUNDY: Why are they doing that?

ANDREW MCCONACHIE: Why are they doing that? Exactly.

WARREN KUMARI: I think there are actually a number of reasons that people are doing that and it also depends on what exactly you mean by that.

RUSS MUNDY: Well, and that's what the section says.

BARRY LEIBA: The idea is, what order do we put things in the document as we lay this out. I mean, I think this new section, the Importance of Referential Integrity kind of touches on a lot of the other things that we were going to talk about first. So how do we weave it all together so it's coherent.

WARREN KUMARI: I mean, it kind of to me feels like Section 6 goes just after Section 2 or even kind of somewhat an executive summary type thing. Why are you bothering to read this document? Why does anybody care? Seems like—

RUSS MUNDY: So it's more like an introduction.

WARREN KUMARI: Yeah. It's not written very introductory already.

RUSS MUNDY: All right. So Rod.

ROD RASMUSSEN: Yeah. This becomes a framing question, right? Do we want to say, here's all the stuff that's going on and then after that say, well here's a distillation of why these—the motivations for why these things are going on or do we want to have a motivation for writing this paper upfront. That's really I think the choice here and you can kind of do both if you want and pick a little bit of this stuff out of the pressures and move it up. I see why you're uncomfortable, Andrew. What order do you actually put this in and you need a narrative flow settled as to how you want it. That's what you're looking for it sounds like.

ANDREW MCCONACHIE: Ideally, we can put what is currently Section 6 right now, just plop it somewhere. I don't want to break it up. Warren, is that a new hand?

WARREN KUMARI: No. I think it was when I [inaudible] to figure [inaudible] about why people want to evolve the DNS, so old hand.

RUSS MUNDY: Right. The discussion at hand isn't that, it's more of how to weave that into the rest of the document. What is the narrative flow of the document today is what we're looking to talk about.

ANDREW MCCONACHIE: So I heard one suggestion from Warren to put it up in the introduction. Does anyone hate that idea?

GEOFF HUSTON: This is Geoff. I think I do and I think there's kind of a little bit more that needs to actually happen around the kind of discussion in Section 4 on the Pressures to Evolve DNS Resolution. Part of the issue is, is that the naming system intertwined with, if you will, the metacode on how to resolve the name. The naming space is not flat. It's not a dictionary. The naming space is actually an ordered collection of nodes intended to drive a resolver through a distributor database.

And oddly enough, if you try and muck around with the resolution, yet keep the same namespace, it becomes almost an impossible act, doesn't it? That in essence, once you think about different forms of resolution that don't use a distributor database, you invariably don't see why you're keeping your hierarchical name system with effectively no points that act as metacode to allow you to resolve.

And so, when you start thinking about evolving DNS resolution, you immediately get led to the Alternative Naming Systems in 5. And it's not just I hate ICANN or I hate the status quo, it's more if you want to change that algorithm of resolving a name, [inaudible] the context, you invariably have to reconsider the syntax of names and the naming structure because right now the DNS is intertwined between the namespace and the resolution mechanism.

And I actually think the order 4, 5, 6 actually works for me because if you kind of added a bit more of this commentary about changing resolution, invariably makes you reconsider what the distributor database is. But when you reconsider what the distributor database is, you [inaudible] to Section 5. And it's not just more TLDs. That's just bullshit, right? The issue is, can flat namespaces work? Can hash namespaces work? Are there other forms of resolution? And the answer is yes, of course there are.

Now, why should we not consider that Section 6? We're stuck where we're stuck because in some ways, this is not evolution. Once you move away from that space, it's kind of impossible to reconcile the two. So that's what I would defend what we're laying that down 3, 4, 5, 6, is actually being a logical commentary.

RUSS MUNDY: So you think the flow is about right as it is?

GEOFF HUSTON: Right. Because 4 and 5 are not kind of filled, and actually, don't quite explore both the specifics and the generalities, 6 looks like a bit of a Martian. But as I said, I would certainly write 4 were I to do it and I don't think I have [inaudible] right now to do so, is to actually talk about the almost tightly bound relationship between the structure of the name and the structure of the distributor database which actually forms the structure of resolution.

And pressures to evolve that resolution for the following sets of reasons and there are a whole bunch of reasons, the outcome of that is that you actually need to reconsider the distributor database and the name system. Full stop, end of Section 4, start of Section 5. Because once you start to consider that, there's the first set which are kind of dull. The pseudo TLDs which is more of the same.

And there's a second set of names which are—let's call them bit coin names or hash names or other forms of name system that use an entirely different resolution system. Therefore, I'm actually reliant on the existing DNS naming systems. And this implication in 5, the Alternative Naming Systems actually lead you or result from—it's a very tied relationship, to alternative forms of resolution.

And that then says, well, why don't we just explore this? Why don't we let a thousand flowers bloom? And then you go ahead to Section 6. The reason is, chaos would be completely unhelpful to everyone. That's why it works for me. This becomes a very conservative document in the end because it's really arguing the status quo because it's kind of the antithesis of let's explore all these spaces because they're fun. It's kind of explore at will but try to deploy it and you're going to confuse

everyone and fatally wound the system. That's all it says. I've said all I can say. I've used my quota or my connection is broken. You're all talking and I can't hear you.

ANDREW MCCONACHIE: I hear you. It sounds like then—I mean, we should leave it like this now. Of course, things could always change in the future. But for now it sounds like we're on the right path.

GEOFF HUSTON: Well, I think the glue that binds it together forms a story without surprising the reader. Because quite frankly, Section 4, if you just outline the motives, then it sounds thoroughly reasonable. I want more of what's good. I want less of what's bad. Great. But think about what you're wanting and why resolution works the way it does in its relation to the namespace.

And the inevitable conclusion in my mind of Section 4 is, once you move away from the current distributor database structure, you invariably find yourself considering alternative naming systems because the naming system, the model of naming is inextricably bound to the distributor database which is inextricably bound to resolution. So if we consider alternative resolution, we are I suppose automatically [inaudible] forced to consider alternative naming systems and these will be discussed in the next section.

ROD RASMUSSEN: I think too, Andrew, when we flushed out Section 4, I think there's a little bit of an introductory paragraph or so talking about—because you're transitioning from traditional DNS to pressures to evolve so just talking a little bit about DNS having this evolution that we have seen over the last 20 years or whatever this number is. And there are many reasons for it and here's some of them and we'll [inaudible]. So then you provide that story arch. I think that works fine.

ANDREW MCCONACHIE: Yeah, we can draw upon SAC109, the DoH-DoT document, which is the original idea with this document was that it was going to be a continuation of that document. But that also has a flow which has this traditional DNS resolution and then DoH and DoT stuff.

ROD RASMUSSEN: Mm-hmm.

ANDREW MCCONACHIE: We even start off Section 4 with talking about DoH and DoT and just referencing 109.

WARREN KUMARI: DoH and DoT are still DNS resolution though, right?

GEOFF HUSTON: Right. And that was the point I was going to make, Warren.

WARREN KUMARI: Yeah, exactly.

GEOFF HUSTON: Because the pressures to evolve are sometimes cosmetic in the current structure, DoH and DoT, and sometimes fundamental hash titles.

BARRY LEIBA: Yeah.

WARREN KUMARI: I mean, that actually raises the—the title is Pressures to Evolve DNS Resolution. Is it actually more Pressures to Evolve Name Resolution?

GEOFF HUSTON: Well, that's the [inaudible] that we're trying to put forward, Warren, that says, because names and resolution become so inextricably bound, once you sort of talk about evolution of resolution, you're inevitably led. So you can start by thinking it's DNS but maybe the section actually leads you to the conclusion that once you start thinking more radically than just cosmetic changes in the channel encryption, once you start to think more radically, you actually start to question what is the DNS per se versus what is a naming system because that's where this leads.

BARRY LEIBA: Well, and we did decide to change the title of the document accordingly.

RUSS MUNDY: Is this in line with that change? I don't remember.

BARRY LEIBA: Yes.

RUSS MUNDY: Okay. So I should just accept this then. So maybe then we want to call Section Internet Name Resolution—Pressures to Evolve Internet Name Resolution?

ANDREW MCCONACHIE: Yeah, that's what I was trying to say.

RUSS MUNDY: Right, right. Okay. We don't even want to say domain name. We want to be even vaguer than that.

BARRY LEIBA: I think so, yeah. I think name resolution is—internet name resolution.

RUSS MUNDY: Okay.

WARREN KUMARI: When I started on this, a lot of the discussion in Section 4 is specifically about DNS name resolution.

BARRY LEIBA: Yes.

ANDREW MCCONACHIE: It makes sense to start with that.

WARREN KUMARI: Yes, but true, I guess.

GEOFF HUSTON: If you're talking about the pressures for evolution, then I would actually argue the pressures lie in the shortcomings of the current system. And this is certainly a list of shortcomings.

ANDREW MCCONACHIE: Right.

ROD RASMUSSEN: Oh, or at least perceived shortcomings.

GEOFF HUSTON: Yeah, yeah.

RUSS MUNDY: Yeah, and some of them are related to the governance of the system.

ROD RASMUSSEN: Can't control my people. I can't make enough money, mm-hmm.

GEOFF HUSTON: It's not my DNS.

WARREN KUMARI: I mean, there are people who really do believe that a resolution system should not have a central point of control which we don't really talk about here. There seems to be—the bottom line is, fairly dismissive of alternative naming systems or at least the justification for them. And clearly, there are a bunch of folk who are just out there to make money, gold rush. But there are a bunch of folk who truly do seem to believe that the current system is stacked wrongly. And we're not really covering that. We're ignoring a fairly large justification from a set of fairly smart people.

ROD RASMUSSEN: Let the [inaudible].

WARREN KUMARI: Yeah. I mean, that's also somewhat—

ROD RASMUSSEN: [inaudible].

WARREN KUMARI: People like Christian Grothoff and Appelbaum and folk. They truly do believe—and Zooko, right? They all truly do believe that the current system should not be centralized in any way. And they have reasonable reasons for believing that. We might not agree with them, but—if you look at GNS, it's not a ploy to make money here.

GEOFF HUSTON: I think in terms of pressures, Warren, I'd certainly agree with you. That this centralized system that offers uniqueness as an outcome is not the only way to get uniqueness. And why else do we have a centralized system is actually a really good question. And can you offer assured uniqueness without centrality? And if the answer is yes, then that's a legitimate answer and a legitimate pressure. And so, I agree with the folk going, why do we have it this way? It's actually a good question.

WARREN KUMARI: And yeah, I like the ideological concerns that Andrew put in there. That does seem like a—

RUSS MUNDY: You're talking about a difference in opinion and etiology. And we might want to have a subsection here where we just talk about the fact that I guess, the DNS could—promotes a hierarchical etiology in a sense, right?

GEOFF HUSTON: Right, and the motivation amongst others. But one of the motivations is assured uniqueness. My name not your name.

WARREN KUMARI: Which I mean, some of the more extreme folk like some of the Grothoff type people, don't actually necessarily believe that uniqueness is the correct thing. They think that a name should be—what's the quote? When I use a word, it means exactly what I expect it to mean, nothing more and nothing less.

BARRY LEIBA: Yes, Humpty Dumpty.

WARREN KUMARI: They think that you as a user should be able to say what dub-dub-dub dot Facebook dot com resolves to. When I use the word Humpty Dumpty said and [inaudible], it means just what I choose it to mean, neither more nor less. Actually, you know what, this would be a fairly nice quote to have in the document somewhere.

GEOFF HUSTON: I think you're right.

WARREN KUMARI: Right? The question is, whether you can make words mean so many different things. The question is, that Humpty Dumpty, which is to be the [inaudible]? That's all. I mean, that actually covers a fair bit of—

RUSS MUNDY: You want different answers for different questionnaires.

WARREN KUMARI: Potentially. The GNS folk believe that you as the person resolving the name should be able to say what the name should resolve to. I think that they're often crazy pants land some of the time.

ANDREW MCCONACHIE: It reads to me like a rejection of objectivity. I mean, it's just—if we take that Humpty Dumpty quote as the example, it's like, why have dictionaries?

GEOFF HUSTON: That's the point in the tension. All of us use words and names the way we intend them to mean. We communicate when we're careful to either use those words in a way that we think others apply the dame meaning or we spend a lot of time saying more and more words to explain that context. So what we intended it to mean is conveyed in all it's wonder to the listener.

Now computers don't have that luxury is part of the problem. It is using a network. So on human terms or on computer terms is actually the

tension. But oh my God, are we talking about this in an SSAC meeting? Yeah, right. [inaudible] BO1, come on down.

ANDREW MCCONACHIE: When you compile a program and you use the same symbol in two different places, guess what the compiler does? I mean, you may mean that this symbol means two different things in each place, but no, the compiler has a very clear—I've never met a compiler that let's you just use this symbol again and have it mean something else.

WARREN KUMARI: Yeah, but not in the same context.

GEOFF HUSTON: Now you're saying scope.

RUSS MUNDY: Right. So maybe scope is the right word for that but that's—you're talking about that within one scope. Here we're talking about using a name in a web browser or an email client or a telephone app or whatever and there may be a reason that it should give you different answers in those different contexts. So it's not the same.

ANDREW MCCONACHIE: So the GNS people then imagine using the same name and different applications in the GNS naming system, in the GNU naming system?

WARREN KUMARI: Name resolution and zone dissemination is based on the principles of a pet name system where users can assign local names to zones. You as Andrew can decide what you think ICANN.org should resolve to. And that is your right. And the fact that you have decided that ICANN.org should take you to Facebook and that will confuse everybody else in the world. It's a feature not a bug.

BARRY LEIBA: And I'll look at a real-world case where I agree with that. For a long time, I put in my host's file, Facebook.com goes to 127001 because I just wanted it to—I didn't want any Facebook link to work, and that worked for me for a very long time. It was fine. Why shouldn't I be able to do that? And there are obviously different ways to do that. Fudging it in a host's file is one way but there are other ways and other reasons I might want to do it.

WARREN KUMARI: Yeah. And I mean, put the blanket in the trunk mean completely different things in the US versus in the UK.

BARRY LEIBA: Well, or even in the US in different context. Am I packing to leave the house by train in which case I mean a box trunk or am I packing the car in which case I mean the boot of the car, so yeah.

GEOFF HUSTON: Where's the elephant?

BARRY LEIBA: Where's the elephant? I don't think I would—

GEOFF HUSTON: You're talking about [inaudible] the room.

ANDREW MCCONACHIE: And it is the elephant in the room.

GEOFF HUSTON: Yeah.

BARRY LEIBA: So that takes us to the old Groucho Marx joke. I once shot an elephant in my pajamas. How he got in my pajamas I'll never know. Human language is fraught with ambiguities as we have in that what's now Section 6. This is why we don't want these ambiguities when we're doing name resolution. And now we're looking toward that as something that people want.

GEOFF HUSTON: There is no doubt people want it. Whether it's actually deliverable or not is really the thing here.

BARRY LEIBA: Yes.

GEOFF HUSTON: Because the last 20 years of computer evolution has actually been all about me and you and everyone else. It hasn't been about the computer. And it's kind of, how far can we go in what you want as a human versus what the technology can produce using ultimately deterministic tools.

WARREN KUMARI: I think that the issue here is that humans don't use language in a particularly deterministic manner. You said deterministic tools. And I think that that's a fair bit of the thing. If you look at—if you watched Star Trek, there is a lot of interactions where people are talking to the computer and what is being expressed is certainly not anything that is precise enough that machines that we use today would be reasonable.

BARRY LEIBA: I will question that. Lots of people today say, Hey Siri, Hello Google, Alexa, and then they blurt out—

WARREN KUMARI: [inaudible] by devices.

BARRY LEIBA: And they expect that something reasonable will happen. It may not be exactly what they had been looking for but computers are getting better and better at sorting that out.

RUSS MUNDY: But how is that different than what you type in the bar at the top of your browser, right?

WARREN KUMARI: Yeah. Well, I think that's—

RUSS MUNDY: Because it's no longer a domain name.

WARREN KUMARI: Yeah. So I think that that's actually a really good point that we don't initially cover particularly well in here that the naming systems we use currently are designed to create a unique one-to-one mapping. And that is very different to, as you said, what you type in to the search bar, which is where people used to type domain names, right? If you and I type things into the search bar, we're likely to get fairly different results based upon a bunch of heuristics. Maybe one of the findings of this paper is that for a lot of things, the right answer isn't what the user actually typed but more what the user wanted. And I don't really know how you get to there.

GEOFF HUSTON: I think though, Warren, you're now crossing the beams because one of the conventional use of a network is a passive instrument that allows—well, let's go all the way. Human uses to communicate. Whereas the other mode which is the mode that dominates today, is the network is the other side of the conversation. I talk to Siri, I don't talk to you. And oddly enough, when I talk to you, the requirements around the symbols, languages, context, et cetera, oddly enough are different than when I talk to my computer, whatever that is.

WARREN KUMARI: So everyone who's using Chrome, if you open a new window and just click in the URL bar, what I'm going to call the URL bar, it will show you the most recent set of things you went to. The most recent like 6 or 7 or 8. How many people actually have a domain name in that top line things?

BARRY LEIBA: This is something that I've said a while ago that increasingly domain names are not relevant to end users the way they used to be. They're relevant to protocols and stuff behind the scenes.

WARREN KUMARI: I mean, there are specific targeted places where domain names are still important to users but that's doing odd things like configuring their mail client or similar. Although using many mail clients these days they automatically figure this out from various discovery mechanisms, so even that is going away.

GEOFF HUSTON: So to bring this argument or [inaudible] discussion back to the document, are you arguing, Warren, that there are pressures to evolve the DNS to reflect this form of usage? Because it's true what you're saying, search bar language, all that kind of stuff is much vaguer than the DNS itself. Are you saying there are pressures that want to evolve DNS and name resolution along those lines?

WARREN KUMARI: I think that there are pressures to evolve the DNS in technical manners to add things like extra privacy, security, cryptography. There are pressures to evolve name resolution for ideological things. Blockchain, etc. There are pressures to evolve name resolution for profit but also in a lot of cases, there are pressures to simply ignore name resolution as a thing. And there are a bunch of examples there, right? The fact that users don't actually interact with any sort of naming system very much anymore unless you want to call a search a naming system.

Examples of this are things like people search for things. People use apps on their phones. They don't use a browser, they open the Facebook app or the Spotify app. And so, users increasingly don't have to deal with resolution systems at all. And it feels like the document is talking a huge amount about evolving the name system itself and the resolution and all of that. But we're not really touching the elephant in the room or golden cow or however you want to use the analogy. That name resolution as a thing is becoming less important as exposed to the user. It's still important for systems. And yeah, we do have some

discussion on it but maybe we actually have more than I remembered. It looks like more words than I remembered it being.

ANDREW MCCONACHIE: If we need to beef up this section, I hear what you're saying, Warren. And if we need to beef up that talk, I would recommend it here in 7.1.

WARREN KUMARI: The top search things on Google number 1 was Facebook, 2 was YouTube, 3 was Amazon, 4 was Wordle, 5 was Google, 6 was Gmail, right? All of those are domain names. Maybe not Gmail but the top 5.

RUSS MUNDY: They just didn't want to type the rest of it.

WARREN KUMARI: Yes, and I think that that's an important point, right?

RUSS MUNDY: What I find funny about that is you've added a round trip, right?

WARREN KUMARI: Yes, the round-trip time for typing Facebook and pushing enter is tiny. A round-trip time at least for me for typing Facebook.com is significantly longer because my typing sucks. There's also—

RUSS MUNDY: The point is that the browser vendors are doing everything that they can to make fewer round trips and then users are adding them.

BARRY LEIBA: So where's the cause and where's the effect? I think they want to decrease the number of round trips that they generate because users are adding them.

RUSS MUNDY: Maybe so.

WARREN KUMARI: Well yeah, I mean, it's also—there's fairly good evidence that user's attention span is really short and every extra couple of milliseconds makes them likely to wander off and get bored.

GEOFF HUSTON: If I understand the last 10 minutes or so, what appears to be the theme here is, names are being increasingly used in context which are—I'm thinking of the word inimical to the DNS. It's opposed in some ways to the deterministic unique form of DNS. We use names and we're probably used to using names these days in approximate senses like in search terms in similar where it doesn't have the precision and uniqueness it used to. And that actually creates pressures in name resolution be it by a search engine or other mechanisms to actually accommodate and allow for such imprecision and apply more context like a search engine in order to arrive at something that suits what the

original entrant, the querier had intended. Is that where you're going here? Some of these things are simply just not DNS things but there is pressure to have names used in such a sense.

WARREN KUMARI: Yes, I think so although it's possible that you were assuming I had a lot more coherent thought behind my shouty rant. Where I'm heading is necessarily from here to the exact point that I'm making. I think it's more within the ICANN world/ICANN bubble, there is this belief that names are the most important thing. It's a huge industry. There's billions of dollars that slosh back and forth and perhaps it's worth reminding people, don't put all of your eggs in this whole DNS basket or naming system basket because at some point, I think it is going to occur to people that spending however much money you will spend on that app or however much money you will spend on business.com is not necessarily a reasonable investment.

I don't remember business.com sold for $350 million in 2007 money. Is that actually what names are worth now or if somebody's actually interested in whatever business.com was hawking, would they just type whatever it is into the search engine?

RUSS MUNDY: I'm the wrong person to ask that because I am completely mystified why people are paying so much for domain names today.

GEOFF HUSTON: But doesn't that level of investment—

WARREN KUMARI: And I still don't know what they sell.

GEOFF HUSTON: But doesn't that level of investment in existing infrastructure [inaudible] pressure to make business.com relevant to other forms of usage of name. In other words, doesn't that become a pressure to evolve oddly enough—devolve the DNS to suit these other places where people are using names because I spent 350 million bucks on this name dammit and I don't want to see it being casually discarded by some stupid hashtag algorithm, hashtag business. So there is some pressure to actually preserve that investment to make the DNS relevant even though the DNS doesn't want to go there.

WARREN KUMARI: Yeah. I mean, I think there is a fair bit of pressure from within ICANN trying to make the DNS system more relevant than it is or keep more focus on names than is necessarily realistic for them. Here we go, actually, I don't know if this is going to be entertaining to people. I wrote this down when I was in KL. I don't know if anybody can read it but there's a bunch of columns with things. This column—wow, doing this backwards and upside down is tricky. This is how many times I saw a domain name in ASCII. This column is how many times I saw something that ended in .ml.

So I saw 242 names that ended in ASCII. I saw three names that ended .my because we're in Malaysia. I saw maybe one IDN version of it but

the actual interesting thing is I saw at least a hundred instances of whatsapp being the contact method that was being suggested. And this is just URLs in the front of shops, etc. So there were still some more domain names than whatsapp references. But I wasn't even counting hashtags. And I think if I had done that, if you add whatsapp and hashtags together, I am fairly sure it would have surpassed any of the domain name DNS lookups.

So I don't really know if we should say something like when somebody uses a hashtag, there is no real expectation of uniqueness there, is there? There's nobody that you go to and like, I want to use the hashtag [inaudible].

RUSS MUNDY: It's ad hoc.

WARREN KUMARI: Yeah, and the people are just using and using them.

ANDREW MCCONACHIE: My concern here, Warren is that we have very—I love your sample, your tiny sample size but I think we need more. I'm concerned that there aren't people doing more research in this area about how users are actually using names on the internet. And I think a while ago, I suggested that that might be something the SSAC might want to say. People should be looking into this kind of thing more. I mean, there's a big push for example to get IDNs in the next round of new gTLDs. Do people actually want them?

WARREN KUMARI: No, they do not.

ANDREW MCCONACHIE: Your question about do people in Malaysia prefer an IDN or an ASCII name, I mean, who knows? We have no idea, no one's looked into this.

WARREN KUMARI: Well, some people have looked into it but I don't think there's anything that I can cite. But from some [inaudible] stuff where I know people looked into it, it was clear that the populous being targeted like developing countries or whatever the correct term for that is that they want, developing economies. Didn't actually want the IDN names because in many cases there are using candy bar phones and changing input mechanism on that is really hard.

If they're not using candy bar phones, it's still really hard to change your input mechanism. And Apple had some—I believe it was Apple. Maybe it was Android release some stuff showing how much of the time people are using each localization version of the keyboard. And it's almost always people just stick with [inaudible]. Then there's also people who view ASCII names as being more prestigious for a bunch of reasons. And also, this is where IDNs might be used more commonly.

We're actually seeing that people are being introduced to the internet through generally phones and similar. And so, they view it as the mobile thingy or the internet is just Facebook. And so, there's also—I can't remember, I'm trying to find the stats. New gTLD stats I think it was

showed that the only set of names which are sort of grouping of names which are actually decreasing in registration is IDNs. Of course, pushing back against IDNs does very much end up with like, think of the children argument. I'll see if I can find the stats.

ANDREW MCCONACHIE: If there's anything we can say about that, if there's any kind of research we could cite, I think that would be really helpful.

WARREN KUMARI: Yeah, it's hard to not come off sounding like a—sorry, I'm kind of searching while I'm speaking at the same time.

GEOFF HUSTON: I still think the meta point here is, as well as enumerating the weaknesses in the current form of resolution, what I'm getting from this conversation is that there has been a wealth of use of names in context that aren't directly related to the DNS. They're not a domain name. And there has been an incentive on preserved investment to extend the current DNS model, to extend the relevance of DNS names into these other forms of use of name existence and inevitably that gets—the result is some pressure on the way names are resolved because those alternate systems don't correspond to the precision and determinism of the DNS. And that's a legitimate pressure to evolve. There's no doubt about that.

WARREN KUMARI: All right, I've wondered down the rat hole of trying to find citations.

ANDREW MCCONACHIE: No worry. I mean, if you can find some, then just drop them to me. I've looked, I just haven't been able to find much in that regard.

WARREN KUMARI: Yeah. I mean, this isn't what I've been looking for but Goldstein Report had stuff. What I'm trying to find is [inaudible].

ANDREW MCCONACHIE: So this wasn't really the way the call was meant to go to kind of flush out Section 4 more but I think it's been pretty useful.

RUSS MUNDY: Yes, I agree.

ANDREW MCCONACHIE: Barry, do you want to dive into the text in Section 7 now in our last seven minutes?

BARRY LEIBA: I guess, we can do that.

ANDREW MCCONACHIE: Okay.

BARRY LEIBA: You got any other ideas?

ANDREW MCCONACHIE: No, I just noticed that people were no longer talking about that, so I thought we might as well talk about something else.

BARRY LEIBA: Okay, let's do it.

ANDREW MCCONACHIE: Okay. So in the last five minutes of this call, take a look at this paragraph. This is in the introduction to the Perspectives on Ambiguous Internet Name Resolution where we have some discussion about the experiences of both users and software developers. We may add more perspectives in the future but just take a look at this paragraph and the introduction.

All right, I'll move on. Let's move on to the next paragraph. Feel free to chime in on either one of these paragraphs now at the top of the screen.

GEOFF HUSTON: That second paragraph.

RUSS MUNDY: Go ahead.

GEOFF HUSTON: Sorry. That second paragraph about DNS has little in the way of implicit context. Actually, skates over and explicitly ignores the bad idea that was EDNS zero client subnet where context was added into a query as an adornment of the query rather than an adornment of the name. And the whole idea was that the resolution service was meant to match the attribute of the query versus the database attribute for folk in this subnet use that domain. And so, there has been some efforts around adornment of the query such as client subnet. But in general, they themselves have extraordinary problems around privacy and accuracy.

ANDREW MCCONACHIE: It's true. We don't talk about that there. We talk about that in the preceding section. Actually, in the text we've been looking at. We talked about EDNS.

GEOFF HUSTON: Yes.

ANDREW MCCONACHIE: Maybe I could just—

RUSS MUNDY: I'd do that but not the adornment of the query versus domain. I don't think we have that—

WARREN KUMARI: Well, hang on a second. I don't really understand how that would be expected to work. Geoff, are you expecting users to type in, imontheeastcoast.google.com?

GEOFF HUSTON: Well, they don't but that's what happens when they have client subnet added into their queries automatically on their behalf.

WARREN KUMARI: Yeah, yeah, yeah or actually, they say I'm using a resolver on the East Coast, but yeah, close enough.

GEOFF HUSTON: So do I expect that to happen, well, I thought it was happening.

WARREN KUMARI: Sure. The reason that that happens is because what users type is just google.com or more likely just google.

RUSS MUNDY: That .com is so testy, right?

ANDREW MCCONACHIE: So I think we're going to adhere but we do need some kind of reference to it.

GEOFF HUSTON: I suppose what I was looking at, Warren, was something different in so far as the assertion is. The DNS has little on the way of implicit context and objects with the DNS cannot be referenced except as precisely formatted domain names. Sort of true but we've now got to this point where we start adorning the query and using that adornment as a selection mechanism around possible answers.

RUSS MUNDY: Well, I think it's an important point that some of those adornments are not under the control of the user.

GEOFF HUSTON: Right, right, that's true.

RUSS MUNDY: That's a really important point. Whichever resolver you pick may be doing things to help you get the best answer, whatever that means to the resolver.

ANDREW MCCONACHIE: Okay.

RUSS MUNDY: It would be interesting to see how often Quad8 and Quad1 give different answers.

WARREN KUMARI: Well, how are you defining different?

RUSS MUNDY: For the same query.

WARREN KUMARI: But I mean, the correct answer for somebody asking—

RUSS MUNDY: I know exactly. You're talking about load balancing and stuff that has to be supported.

WARREN KUMARI: Yeah. I mean, how often they give different answers—if you mean ones that were not intended [inaudible] domain—

RUSS MUNDY: No, I'm assuming they're not doing that.

WARREN KUMARI: Oh, I mean, sometimes they are.

RUSS MUNDY: But I haven't really tested.

WARREN KUMARI: 1.1.1.1 has been compelled to give NXDOMAIN, I believe. I can't actually show what they're returning by the Italian courts for one of the pirate sites. Actually, that brings me to a very interesting thing. I'm going to just—seeing as we're past the top of the hour, I'll just start ranting now.

So many years ago, SSAC wrote I think SAC50 and SAC56 about DNS blocking and content blocking by using the DNS. And it's largely suggesting that if countries want to perform censorship, the DNS is probably not the way that they should do it. This seems to be becoming a thing again and I don't know if we should—so I'm trying to consider updating that.

ROD RASMUSSEN: I think it still--

WARREN KUMARI: Sure. Okay, yes. Still but I mean, it seems to be becoming even more popular and being pushed.

ROD RASMUSSEN: This is one of the inspirations for this work is that, if you have a huge chunk of your populous using a smaller set of resolvers, it's a lot easier to enforce this or at least it seems that way. It is a lot easier to enforce this and for the uninformed it might be.

WARREN KUMARI: Yeah. I think, the governments are – okay, broad generalization. Many governments are very unhappy if their users use public resolvers. Because they cannot actually—it's not as easy for them to compel external public resolvers to do what they want. For example, the UK had the—and Australia also had their set of blacklists that resolvers within the country had to comply with. And if it is something like Cloudflare, it's much harder to compel them to do various sets of blocking.

ROD RASMUSSEN: If they don't have [inaudible] it sure is.

BARRY LEIBA: We do need to wrap this up because we've gone over. I'll just leave something for us think about and maybe talk about next week. If I type in facebok.com with one O into my browser, is it wrong for it to send me to Facebook? And just think about that and we'll come back to that sort of thing next week. Any closing comments, Andrew?

ANDREW MCCONACHIE: No. Other than that I'll be on vacation next week so you'll be getting Mr. Sheng as your staff support.

WARREN KUMARI: And just as a giggle, if you type in facebok.com, your browser does take you to Facebook.com. But that was I realize just an example.

BARRY LEIBA: Yes. Okay.

ROD RASMUSSEN: Try it with two zeroes instead.

BARRY LEIBA: [inaudible] that. So anyway, see you next week.

ANDREW MCCONACHIE: See everyone next week.

RUSS MUNDY: Thanks, see you.

BARRY LEIBA: Bye.

[END OF TRANSCRIPTION]