
RUSS HOUSLEY: Welcome to the SSAC Evolution the DNS Resolution Work Party Meeting for May 11th, 2023. Barry, over to you.

BARRY LEIBA: Okay. Well, unfortunately, we don't have too many people who were not-- We have Geoff and Yap who were on the call when we last met, and Merike and Tara to add their opinions. Merike and Tara, did either of you have a chance to listen to the recording from the last time?

TARA WHALEN: I did.

MERIKE KAE0: Yeah. I did not.

BARRY LEIBA: Okay. So, I'll start I guess by putting Tara on the spot, sorry about that, and asking if you have any comments on some of the things that Geoff brought up last time.

TARA WHALEN: Sure. There was quite a lot there. So, I think it was helpful to have identified the need to talk about, I guess, the functionality or the reason why some of these alternatives were proposed. So, I think a lot of that was in the third section. So, I would agree that what seemed to be

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

missing there was a discussion of the fact that these systems, these alternate systems didn't just sort of spring forth in someone's head for no reason, that it was at least believed that they were filling a gap.

Some of them may have been filling a gap in the past that may no longer didn't exist and that's why they went away. Some of these other ones may there may be things that persists, there may be things that could be accommodated in different ways, but I think making those explicit is pretty helpful. Some of these may be technical, some of them may not be, and I think both of those were discussed a little bit or at least touched on in the last call, looking at some of the, I guess, business things as well as some of the technical components. And I think those were both interesting. So, I was sort of intrigued to see where the text was going because it was an interesting discussion that you all had in your small group.

BARRY LEIBA: Yes. And, Geoff, thanks so much for putting some text in the other day in Section 3. It's given us something concrete to look at.

GEOFF HUSTON: Sure.

BARRY LEIBA: Well, I guess, I don't know. Do we think it's productive to go through that text now and--? I'm really wishing we had the usual participants on. I don't know what we can do about trying to fix that. Yeah. I'm

kind of at a loss there, but do we think it's useful to go through the tech stuff road with this group?

RUSS HOUSLEY: Well, having just done it and been late to the meeting because of it. I found it very informative and people at least should read it, but I don't know if it's worth going paragraph by paragraph right now.

BARRY LEIBA: Yeah. My comment certainly is that I think it's taking that section in the right direction. So, I'm pleased use that as a better starting point than what we had.

MERIKE KAEO: Yeah. I mean, you know what I would like to see is some of the main points, right, if Geoff can maybe rather than going word for word, right, do we have everything covered? I mean, I just went through security, so I'm now in a space at the airport where I'm just sitting by the gate for the next three hours. But I took a quick look at the text, and then when I was at NIH, we had Bitnet, right, and I remember being in naming discussions at the time, and we were also starting to use TCP IP. And so that's why I had a just quick comment about whether or not organizational aspects, trying to figure out how to organize names better. I mean, not just the technical limitations, but whether that was also a component about creating the hierarchy. And I hadn't read through the text enough to see whether or not any of that was in the text at all.

GEOFF HUSTON:

I was trying to concentrate on evolution rather than history. And that's why I kind of made the note at the start that while I think ease of governance was a by product, the basic sort of motivation, and this is as much as I said, was that the hierarchy allowed you to scale and computers individually couldn't handle it. So, divide and conquer in a structured way. And that structured hierarchy was simply because we didn't have things big enough to cope with the scale. And that then caused a resolution process that was actually in two parts. It was kind of invisible, but it was in two parts. There was discover the right name server, and then secondly, once you've done so, ask the right name server the right question.

And the trick was, at the time, to make both of those phases use exactly the same query response protocol. So, discovery, and then the query become all the same query in terms of the protocol. And as I pointed out then, you get delegated governance. And so entire sub-tree is hived off to other parties to look after, and so you haven't got a single authority trying to look after everything. And the point that I was then making is that the motivations to change that, because that is a very effective design, the motivations to change that come from certainly, a technical capability.

You're not forced into a hierarchy with these names. And blockchains and distributed hash tables are both examples of the fact that really large systems can be crammed into computers these days, and it doesn't cause them to melt. And so, this idea that, well, we need a hierarchy because that's the only way it can operate is kind of false.

That's not the reason why we use a hierarchy anymore. You're not forced to a hierarchy. And the motivation to evolve that comes from the fact, firstly, that we can, right? So that was kind of the opening few paragraphs there that sort of look at that basic idea that the hierarchy was sort of imposed by the limitations at the time, and those limitations are changing, and those changes enable some incentive to evolve.

MERIKE KAE0:

Thank you, Geoff. Now that I know kind of the thought, train of thought of how you evolved the wording, that will help me also as I actually read it in more detail. So, I really appreciate that.

GEOFF HUSTON:

And from that point, I do it as a throwaway, that it's not intended to be slight. The first of these pressures is actually one of that speed. The hierarchies, the deeper they get, the slower everything gets. And, quite frankly, that slowness is not adequately accommodated by caching. And so, there's this increasing pressure to do it faster. And that actually results in "I need or I want to resolve fewer names in order to serve content". And so, that first kind of evolutionary pressure, which is quite strong across the entire industry is, no one tolerates delay any longer, they just don't. And there is a strong pressure to try and get names resolved more quickly.

And I haven't really talked about how that has affected the DNS because I didn't think that was actually relevant. It's more the motivation "I want it faster" that I thought was worth highlighting. Tangibly, what happens is, when you're in a hierarchy, you can't jump, you can't get

there more quickly other than caching. But what you can do is stop ancillary questions. One question solves everything, and that's the motivation, at least one answer solves everything. And that's the motivation for the SVCB records, the https records. In other words, you adorn the answer with literally everything the content client wants to know so that they get everything they need effectively in one DNS resolution operation, rather than multiples. Right?

But underneath all of that, underneath that, increasing the richness of the answers in resolution, is the simple desire I want things to go faster. And that was really what I said, that changes to the DNS that promote faster resolution of names and a reduction in the number of names that need to be resolved, and its queries rather than names in order to serve content. You there, Barry?

BARRY LEIBA:

I'm here. Yeah. I don't know. I'm kind of lost at how to steer today's conversation given how few of us there are. I was pleased that the last time there were few of us on, we were able to have that broader conversation. And I was hoping that today we could drill down on it and we really don't have the critical mass to do it.

TARA WHALEN:

I don't know. Geoff, did you have any point when you were putting this together that you definitely wanted to kick around or pieces you were kind of, I don't know, the knots or things that you thought were particularly thorny that would be useful for this group? I don't want to take over Barry's agenda. I'm just trying to make use of the time well.

BARRY LEIBA:

No. Absolutely. If you have the ideas, go for it.

GEOFF HUSTON:

There's only one aspect that I didn't tackle, and I find it a difficult one to tackle, and I'm not sure under what sort of taxonomy it hits. All of this kind of assumes names are relatively permanent, and that names are a bit like static identifiers. There is another class of evolution going on where you regard the DNS query name as almost as a piece of microcode, and you regard the server as an engine that executes that code. And so that what you find going down that path is actually dynamic names that have no permanence that are intended to take advantage of the DNS's remarkably effective query response mechanism, and overlay that with customized queries and customized responses.

And you kind of wonder if that concept, and you see it a few times in dynamically generated names and answers, you wonder if that's a trend or merely an interesting by-product that has no future. And I didn't touch upon that issue of the permanence and the static nature of the identifiers that are used inside the DNS system at the moment and this small bubble of thought that says it doesn't necessarily have to go that way, and that you can make this a fully dynamic system using dynamically generated names and the queries are customized against the name. That was the one thing I didn't talk about.

I'm not sure it's a motivation to evolve. It could all be in Section 5, an alternative naming system that it's almost the atavism of unstoppable domains. It's the ephemeral domains that have no permanent

existence. That's the only thing when I wrote this that I thought, oh, I haven't talked about that. Then I thought, well, maybe it's not really an evolutionary pressure.

MERIKE KAE0: So, one comment I will have is that the ephemeral domains were so often used by criminal organizations. Yeah. I mean, they're awesome for them.

RUSS HOUSLEY: You're correct, but they're also used in workflow systems like Salesforce. I mean, so just because it's good in one case, it doesn't mean it's bad for the other.

MERIKE KAE0: No, I totally agree. I think there's benefit to articulating something about ephemeral namespace. Just because some criminals use something for bad doesn't mean we should stop it if it can be used for good, right? It's just you have to figure out how to deter the bad.

RUSS HOUSLEY: Yeah. Or at least make sure it doesn't take over the whole thing, right? I just thought it was worth pointing out that it is used for these workflow places.

MERIKE KAE0:

I'm going to have to ponder on that. I'm not so good on thinking of my feet, so apologies for that. But as I'm going to read through this, I think I like-- And thank you again, Geoff, for enumerating, right, the entire thought process here as to how you wrote the section. But I think the path is right because we're trying to also set a context for why different name spaces may exist or what problems other people are trying to solve that you maybe necessarily can't do and the DNS as it currently stands, or maybe people don't understand it or just want to create a different system for whatever reason. But in pointing out what exists in the DNS and how it gets utilized and how it's evolved, right, can set more context as to maybe what is missing or isn't missing where other people have created their own namespace.

GEOFF HUSTON:

Yeah. This entire work group is a tricky one. It's to my mind, trying to point out the reasons why they're all these other name sort of systems around, and they're not entirely evil. And I suppose to even push that a bit further, if you regard the DNS as not being eternal, and you should regard it as not eternal, it's going to get changed drastically at some point, then whatever it morphs into is likely to have its seeds in these alternate name systems in terms of change points, likely. And so, rather than regarding all of them as fragmentation pressures and distractions, there is some benefit in actually understanding why they exist, the niches they intended to fill, and whether there are gentle ways to fill those niches without necessarily throwing everything out.

And I did make a passing comment around the privacy area, that we have managed to slide in DNS over encrypted transport session level dot

go, etc., without ripping the entire DNS apart. And the status of the installed base kind of demands that, at least, at the moment, we accommodate those precious where we can in a gentle fashion that doesn't rip everything open. I think that's got the most we can do here. Just like I said, it's kind of difficult.

I ranted last week, and I'll say it again, the aim of this paper and this document and this work group is not to identify all the animals in the zoo. And so, listing all of the alternate name systems and the characteristics is less important. Because that won't stand the test of time. What is more important is to actually understand the rationale behind these approaches, what are you trying to solve that the DNS doesn't solve and highlight that as the reason why these things have seen the light of day. And it's not to say much more than that, that the DNS is not necessarily all things to all people. It's a particular set of trade offs, and there are alternative trade offs that exist and there are reasons why those alternative trade offs should be considered in some form or fashion. And I'm not sure there's much else to say, personally. It's kind of the rest is left to the forces of whatever operates in the marketplace in the future.

MERIKE KAE0:

I agree with you, Geoff. And what I'm wondering is, I mean, is this paper going down the road of just being informative as to what exists, rationale, and not make any recommendations, or has there been discussion about that we want to make some recommendations, we just don't know what? I mean, I would actually veer away from that, but I don't know if that's been discussed.

ANDREW MCCONACHIE: It has. The goal is not to develop recommendations. If we wind up with recommendations when we're done, that will be great. But it's entirely possible that we'll wind up with a study rather than a set of recommendations.

MERIKE KAE0: Thank you. And I actually think that makes sense to me.

BARRY LEIBA: Okay.

TARA WHALEN: I wanted to have a little more time with Section 3 to see where it might line up with Section 5, so I guess we have the section with the systems, and I think some of the points are touched on in both but perhaps lightly. For example, talking about tour and you would sort of expect that if there was a motivation, for example, about evolution that talks about why, what are the particular goals of Tor and whether you'd then come to Section 5 with Tor. You would sort of maybe expect that's where it'd be discussed. I think we probably wouldn't even be talking about Tor at all, in fact, if it wasn't for dot onion. It's a very particular component of it that I think brings Tor to the forefront. And the same with multicast, it's really dot local is the aspect that gives relevance to this particular paper.

So, I was sort of trying to pull these two sections together, but it sort of just recently got the Section 3 in. I'm trying to see where those might fit together well so they're sort of not just, I guess, floating out there as lists of systems but so much are examples of what was talked about in Section 3 and why don't these exist in regular DNS or to the degree to which they are covered where are they in conflict? Sort of why not do this in traditional DNS in the first place? Why have the alternate? Where do they might they cause conflicts to have been naming or where might they be accommodated? I expect there would be some places there to maybe draw threads between the motivations and the examples of these things in Section 5 or where they are never going to work at harmony will always live in parallel land.

ANDREW MCCONACHIE: Well, and I think that that was part of Geoff's point was that that, yes, they will likely remain living in parallel. Keep in mind that both Section 3 and Section 5 needed a bunch of changes to fit in with the discussion that Geoff and I had the other the other week. And Geoff has made a start on Section 3 in that regard, but has not-- we don't have changes to Section 5 yet. So, yeah, there's going to be a disconnect right now.

TARA WHALEN: Sure. Yes, I guess that was kind of evident. I'm not surprised that it's been flagged, but that was immediately where I was looking to try to pull the thread. So that that's, I guess, I should say also that's why I'm glad the foundation of 3 is there because that sets that up very nicely for 5.

ANDREW MCCONACHIE: Yes. So, what I took away from listening to last week conversation was that I need to make some changes to Section 5 that kind of highlight how these systems are different than the DNS and different than the status quo. Or, well, not the status quo, but well, the DNS status quo, I guess, you'd say. That's on me. Yeah. And I'll hope have that.

GEOFF HUSTON: A combination of how they're different and why they're different, what were the needs that caused that.

ANDREW MCCONACHIE: And the motivations for why they were designed to be different.

GEOFF: Yes.

ANDREW MCCONACHIE: Right. Okay.

BARRY LEIBA: All right. Are we in a position right now where the best thing to do is to have you, Andrew, go off and take a stab at that and for us to continue next week and give people a half an hour back? Or is there something productive we can do with the remaining half hour?

ANDREW MCCONACHIE: Do you want me to merge all the text, like all the changes in Section 3, so it's a bit easier to read?

BARRY LEIBA: Yeah. This lets us still go back if people say, oh my God, this was completely crazy, right? Even after we accept the changes.

ANDREW MCCONACHIE: Yeah. Yeah. There's actually a feature in Google Docs which allows you to name a version and then you can go back to a named version.

BARRY LEIBA: Okay. Well, then why don't we name this version and then go ahead and accept all of Geoff's changes so that people have something clean to look at if it's a little easier to digest? There we go. Okay. So, yes, that sounds like a good idea.

GEOFF HUSTON: Hopefully, it's clear as well. I tried to mark it in red, and I was a bit reluctant to simply strike out a whole bunch of text, but I have put two markers there because I wanted to cut down the size of this section so I lifted out some paragraphs that were relevant, and then just said the rest of this section, all up to the final paragraph, which I kind of liked, just get rid of it. Because it's not really about motivation.

ANDREW MCCONACHIE: Sure.

GEOFF: Okay. Well, that's good.

ANDREW MCCONACHIE: Well, I think we have some other leftover bits of text towards the end of this document that I've chopped and stored there for possible later use, so I'll do the same thing with this. I like to do that when they come to text.

BARRY LEIBA: Yeah. Perfect.

GEOFF HUSTON: There are two comments here. I'm not sure exactly what to do with it. The first is this previous conversation we just had about this idea of ephemeral queries and ephemeral answers. There is an evolutionary pressure to take the bits of the DNAs, the query response, which work actually remarkably well and use that for other purposes, so you're not querying them with the names, you're just querying. And I think as Russ pointed out, Salesforce does that.

The other thing that struck me too is there is a pressure, has been forever, about the global assumption of the DNS as being inappropriate localized domain names. And what we've done in the past is left the DNS alone and changed the infrastructure to give you locality. This area of the net, these names, are defined and relevant. If you physically move away from that area, that context is dropped. There have been

efforts to try and put that into the DNS itself. That's that. The answer depends on the relationship of the query up to some form of locality or scope. And there have been evolutionary pressures to improve that mechanism such that you don't need to change the infrastructure of the network, you change the DNS so that it is context aware of the place or the location between the query or the answer.

The first primitive and I think very primitive way of doing that is explicit in client subnets in EDNS, the whole ECS thing, which is a privacy nightmare. But I suppose you have to understand that this didn't come out of the blue. This came out of I want to create localized answers rather than having every answer global. So, I've put in two comments there, Andrew, and if you want to flesh it out, that's fine. I just wanted to talk about them not in terms of, well, A does this, B does that, C does that, but just simply in terms of these are pressures because the DNS doesn't do either of those things very well right now. And folk are reacting to that by sort of tweaking and trying to evolve the DNS resolution process to accommodate that. Thank you.

ANDREW MCCONACHIE: When we talk about ephemeral names, I just want to understand what we're talking about. Are we talking about names that don't exist in zone files, but that can exist when they're queried for, basically?

GEOFF HUSTON: Right. We often do this with wild cards in order to not to tweak the server too much, but you can do it in all kinds of ways. Don't forget a

zone file is merely a convenience for the server, not for the client. The client is actually unaware of zone files.

ANDREW MCCONACHIE: Sure. But if the authoritative server is reading from a database that in the database is relatively static, then those are not ephemeral names.

GEOFF HUSTON: Yes, whereas, if the authoritative server is actually saying, you asked for foo in my dictionary, the definition of foo is blah, and you never see the dictionary, it just computes it. So, the answer is obviously blah because you asked for foo. And the point is folk are looking at simply using the query mechanism, but discarding the rest of the characteristics of the DNS, the zone files and permanence, and actually making efficient distributed databases from that, because you can. Or as criminals, you can make untraceable names that are damn hard to find where they are. Now, I don't want to talk about the use. I just wanted to talk about that pressure to sort of poke at the permanence of names to go, that's just a convenience, it's not necessarily an axiom.

ANDREW MCCONACHIE: Okay. I'll see if I can figure something out to add there. I too need to read this new text.

GEOFF HUSTON: Yes.

BARRY LEIBA: Okay. So, then my inclination is to give us 25 minutes back now. Does anyone have some final comments before we do that?

TARA WHALEN: No, but I did want to get a sense of, I don't know, Andrew, what your timing might be. And I'm probably jumping ahead of chair here, but I'm trying to look ahead to the next meeting and kind of will we have had a chance for us to, I guess, see new text or get work put in the pipeline ahead of being on the call so that we can be diligent before the next time we all meet.

ANDREW MCCONACHIE: Yeah, that's tough for me to say. I will definitely try to get this done as early as I can next week. I can probably guarantee it before the call. I don't know if I can guarantee it before Tuesday, but when I do get something, and I'll send a mail to the list.

TARA WHALEN: Sure. That was not meant to be pressure on you, Andrew. We appreciate all that you do.

ANDREW MCCONACHIE: I want to get it done soon as well. And we may want to point out as well that the week after next is ripe, [crosstalk 00:31:54]. Oh, sorry, did I steal your thunder, Barry?

BARRY LEIBA: Go ahead. You got it. Continue.

ANDREW MCCONACHIE: Sure. So, the week after next is ripe and the meeting is cancelled. So, we're meeting on the 18th, but we're not meeting on the 25th.

BARRY LEIBA: And we will also have a meeting during ICANN77. So, here's for future planning. We intend to have a broader discussion on a substantially completed document at the SSAC workshop in mid-September. Which means that by, let's say, late August, we should have a goal of having a document that is not necessarily complete but substantially enough complete that we're happy to pass it on to SSAC and have them do a review of it before that discussion. So, let's keep that in mind and try to make progress. And that's it from me. Russ, anything from you?

RUSS HOUSLEY: No. That all makes sense to me, and I'm looking forward to getting to a stable document. But I think this is an improvement to the document we discussed today.

BARRY LEIBA: Okay. So, we'll see all of you and I hope a couple of others next week. Thanks.

MERIKE KAE0: Thank you.

BARRY LEIBA: And, Merike, have a have a good flight.

MERIKE KAEO: Thank you.

[END OF TRANSCRIPTION]