[KATHY SCHNITT]: Welcome to the SSAC Evolution of the DNS Resolution Work Party teleconference on Thursday, the 20th of October, 2022.

Barry, back over to you.

BARRY LEIBA: Okay. Well, it'll actually be back over to Andrew. He wrote a whole bunch of text, and that's really cool. And so we get to have a look at it and see if we think that it says what we want it to say and it's going in the right direction. Andrew?

ANDREW MCCONACHIE: Sure. So we were going to start with 6.1 and then go to 6 and then go to 6.2 because we've got new text in all those places.

BARRY LEIBA: Because we like those movies with non-linear storylines.

ANDREW MCONACHIE: Yeah, we're jumping around a bit.

BARRY LEIBA: Memento. Did you see Memento?

ANDREW MCCONACHIE: I saw it, but then I forgot it.

BARRY LEIBA: Yeah, well, so did he. Anyway …

ANDREW MCCONACHIE: So we're starting with 6.1. I think most of the comments last week were from [Geoff]. So we're going to be guessing whether or not I actually wrote what [Geoff] wanted me to write. I think this paragraph here—I'll put it at the top—I added some text to based on … Right. I hadn't mentioned before the omnibox in browsers.

BARRY LEIBA: Right.

ANDREW MCCONACHIE: So I added a couple sentences to the end there to discuss that.

BARRY LEIBA: Mm-hmm. And I certainly agree with what's in that.

ANDREW MCCONACHIE: [Geoff] also said … He was saying last week that the DNS was never designed as a search functionality. And we had that discussion about the signposts in cyberspace and the real problem with using domain names as some kind of …

BARRY LEIBA: Search key?

ANDREW MCCONACHIE: Yeah, search key. Like, search engines and domain names solve different problems.

BARRY LEIBA: Yes.

ANDREW MCCONACHIE: And the big thing is that you have to write domain names precisely. And that's what really makes them un-user-friendly for that kind of search thing. Those are the two comments on this that I wrote this text for.

BARRY LEIBA: Yeah.

ANDREW MCCONACHIE: Should I move on?

Okay. I'll leave them as suggestions so [Geoff] can take a look.

RUSS HOUSLEY: Yeah. I think Warren also talked about the omnibar, but he's also not here.

ANDREW MCCONACHIE: Okay.

So moving down to this paragraph, I hadn't explicitly said in the text that a QR code has a domain name underneath it or a URI underneath it. So this text covers that. It's just something that resolves to a URI.

BARRY LEIBA: Right. The main point of that is, from the user's point of view, they're just pointing their camera at something. There's no domain name involved from their point of view.

ANDREW MCCONACHIE: Yeah, they don't see it.

BARRY LEIBA: Right.

ANDREW MCCONACHIE: Okay. And is it URI or URL, you IETF people?

BARRY LEIBA: Alright. The IETF calls them URIs. W3C calls them URLs.

ANDREW MCCONACHIE: Oh.

BARRY LEIBA: As long as we're consistent, I don't care what we call them here.

ANDREW MCCONACHIE: Okay. Well, I've been using URI. I will continue using URI.

BARRY LEIBA: The idea is that there are locators and there are names, and URIs are intended to be a generalization that includes both.

ANDREW MCCONACHIE: Okay. Well, I'll keep using URIs.

And that's it for 6.1. And then we get into the text in 6.2.

So let's get back up to—

BARRY LEIBA: Well, let's hold on on 6.1 and just see, Tara and Jim, if you have any comments from what you've read now or before on 6.1 that you don't think we've covered correctly.

ANDREW MCCONACHIE: Let me just go up to the top.

BARRY LEIBA: Jim claps his hands, so that's good.

TARA WHALEN: I think it's fine, yes.

BARRY LEIBA: Okay.

TARA WHALEN: Thanks.

BARRY LEIBA: Great.

ANDREW MCCONACHIE: Perfect.

BARRY LEIBA: Okay, carry on.

ANDREW MCCONACHIE: Okay. So [Geoff] had made a comment in 6.1 about wanting to talk about context and how humans all … Well, here. This is [Jeff's] comment down here. And I implemented this comment in the introduction to 6—so not within end users or software engineers or anything but just introduction to this section.

BARRY LEIBA: Okay. Let's take a look.

ANDREW MCCONACHIE: And so this is that text.

RUSS HOUSLEY: I think I see what you're going on at, but I think users draw on their own experience to assume a context that in turn is used to interpret the identifier.

BARRY LEIBA: Yeah, I'll go with that.

ANDREW MCCONACHIE: So users … So are you talking about the first sentence?

BARRY LEIBA: Well, users will draw on their own experiences to assume a context to interpret end user Internet identifiers. I think that's what Russ was saying.

ANDREW MCCONACHIE: I see. Sure. And it's each group in this section. So we can say the same thing about software developers.

BARRY LEIBA: Yes. Right. So the telephone number thing is a little funny because, for instance, you can't tell whether a number is a mobile or a landline from the way the number is presented to you, right?

ANDREW MCCONACHIE: You can in the Netherlands.

BARRY LEIBA: Okay. You can't in the U.S.

ANDREW MCCONACHIE: Correct. That's true.

BARRY LEIBA: Yeah. Alright. Yeah, let's not worry about that.

ANDREW MCCONACHIE: I can change this "often" to a "may."

BARRY LEIBA: Alright. Sure. An interesting thing about that is that, when I left IBM and was sending out resumes and stuff, IBM had contracted a company to help us to do that as sort of a transition out. And when I put my phone number as +1-whatever-whatever-whatever, this professional company that was doing this said, "Why did you write it like that? I don't understand." I said, "Have you ever looked at telephone numbers

outside the U.S.?" Apparently not. But in the U.S., we don't say "+1." We just say "1."

ANDREW MCCONACHIE: Yeah. And here we sometimes use +31, but not always—in fact, rarely. I think it's becoming less common here in the Netherlands for people to ask for a +31 number.

BARRY LEIBA: Well, in the U.S., you generally don't see the 1 at all. You just see the number because everybody assumes it's in the U.S. So it's also part of assuming the context. If you have the phone number and are lacking any other information, you assume the phone number is in your country.

ANDREW MCCONACHIE: Yeah.

BARRY LEIBA: So, yes, I think the reference to telephone number does give a sense here.

ANDREW MCCONACHIE: I mean, it's not precisely what [Geoff] asked for because when we're talking about different resolution contexts, with both the telephone numbers, and with search list processing, we're kind of selecting a sub-tree of the larger global tree of identifiers. And I guess maybe [Geoff]

was talking more about, say, whether something should be using DNS or it should be using mDNS or Tor. Now that I read this again, I'm like, "Well, maybe that's not precisely what he was talking about."

BARRY LEIBA: Yeah, I don't know. So if you and I see something that says "printer.local" or whatever, we have a sense of the context there, but if my mother sees that, she'll get a sense that it's a printer, but the local part doesn't mean anything to her. And she's not going to assume the same context that we are.

ANDREW MCCONACHIE: Mm-hmm. Well, that's covered by the first sentence, I guess: their own experiences.

BARRY LEIBA: Exactly. It may be that we want to expand that a little bit and say something about it. Highlight that one of the points is that different users will make different assumptions.

ANDREW MCCONACHIE: Mm-hmm.

BARRY LEIBA: And I think Russ's suggestion to add the "assume" there is important for that. So, yeah.

ANDREW MCCONACHIE: I changed “users” to “people” because …

BARRY LEIBA: Yeah. That’s fine.

ANDREW MCCONACHIE: Yeah. We’re not just talking about users, even though every person is a user in the end.

BARRY LEIBA: You can say “different entities,” but, no, just leave it this way. It’s fine. Somebody will pick on that another week.

ANDREW MCCONACHIE: Okay. Are we good with this? And this is a reference to the SSAC document on search list processing.

BARRY LEIBA: Okay.

ANDREW MCCONACHIE: Alright.

BARRY LEIBA: Anything on this from anybody?

Alright. Let's move down to 6.2.

ANDREW MCCONACHIE: Okay, great.

BARRY LEIBA: This gets really fast when there are fewer people on the call. It just means we're going to have this discussion all over again when Warren and [Geoff] and Suzanne are on.

ANDREW MCCONACHIE: So I asked for—

JAMES GALVIN: Happy to slow it down for you.

ANDREW MCCONACHIE: Oh, go ahead, Jim.

JAMES GALVIN: No, all I said is, "Happy to slow it down for you."

ANDREW MCCONACHIE: Oh, okay. Well, I have a feeling that this section will be more contentious because this is all brand-new text. And what I heard last week a lot was that contemporary software development is really based

on using APIs, and so many things about naming are abstracted away so that developers often don't even know what kind of—I don't know—naming infrastructure they're actually accessing, right? They're just calling functions. I mean, Warren talked about JavaScript, about how JavaScript doesn't even have a way to resolve names. It just has a way to get URLs. So that's it. I guess you could implement DNS-over-HTTPs in JavaScript, but it's not like people are doing that.

So I started it off by talking about how things used to be, a bit of the history here, and then went into how things aren't really like that anymore and started talking about different APIs that developers might use that end up resolving names. But there's very little in here that's contentious. It's very expository text. So if anything, we might want to expand on this section and say something with a bit more meat on it. But I'll let people read it first.

BARRY LEIBA: Alright. Well, in the second paragraph, that's when you start to get to the URI thing. And so somewhere there we need to talk about the context that URIs are meant to provide and how that's changing with—and probably in another paragraph; probably not in that paragraph—with HTTPS being kind of the default context that now goes beyond what it was originally intended to be.

ANDREW MCCONACHIE: Right. Because Tor uses HTTPS as well.

BARRY LEIBA: Right. And …

ANDREW MCCONACHIE: As does … What does an mDNS URI look like?

BARRY LEIBA: Offhand, I have no idea.

ANDREW MCCONACHIE: Huh. Curious.

BARRY LEIBA: What does a … Well, I don't even know that … It's possible there is no URI for it.

ANDREW MCCONACHIE: Mm-hmm.

BARRY LEIBA: What does using .eth look like? Does that go through HTTPS?

ANDREW MCCONACHIE: Right.

BARRY LEIBA: I don't know. And Warren probably does.

ANDREW MCCONACHIE: I know IPFS has its own URI kind of symbol. And maybe mDNS doesn't have one.

RUSS HOUSLEY: mDNS doesn't have anything that's different. You use a different protocol with it.

ANDREW MCCONACHIE: Sure.

BARRY LEIBA: But is there a URI that triggers that? Or is there not even an mDNS URI?

RUSS HOUSLEY: I don't believe there is.

BARRY LEIBA: I think that's true.

RUSS HOUSLEY: I think it's totally a different resolution protocol. And you're assumed to know what you're doing when you use it. I mean, there's no UI for it, in a sense, if you take printer.local, right?

BARRY LEIBA: Yeah.

ANDREW MCCONACHIE: Right. So if take printer.local in an Apple Finder window, for example—not a browser but a finder—I don't need to preface it with anything, like AFP. Or I guess I could, right? I could use the Apple File Protocol (afp://) and then use an mDNS name, and it will just figure that out.

RUSS HOUSLEY: I don't know if it would figure that out or not. I've never tried that.

BARRY LEIBA: And then there's Bonjour. So I [don't] want to go into Apple protocols.

RUSS HOUSLEY: Bonjour and mDNS I think have come together.

BARRY LEIBA: Okay.

ANDREW MCCONACHIE: Yeah, I thought so. I thought Bonjour was just Apple's implementation of mDNS.

BARRY LEIBA: Oh, okay.

RUSS HOUSLEY: Yeah, it just came first.

ANDREW MCCONACHIE: Right.

BARRY LEIBA: Okay.

ANDREW MCCONACHIE: Let me know if I need to scroll down.

BARRY LEIBA: Yeah, I'm ready for scrolling down.

ANDREW MCCONACHIE: Okay.

BARRY LEIBA: Yeah. To Jim's question, formally it is called a special-use TLD. So maybe that's what we should say there instead of "reserved." Or "private." But it's defined in a particular RFC—I think 6762 or something like that.

ANDREW MCCONACHIE: I think, when I've talked about in the past, I've referenced the IANA special-use domain names registry.

BARRY LEIBA: Right. And that's what it's formally called: a special-use TLD.

ANDREW MCCONACHIE: So we'll use that under the special-use.

BARRY LEIBA: Yeah, we can say "TLD" but "special-use" as the adjective would work.

ANDREW MCCONACHIE: Mm-hmm. I need to put a reference here.

BARRY LEIBA: Now, if you really want to do the right thing with that … I'll just edit it as the official IANA wonk here.

ANDREW MCCONACHIE: Oh, okay.

BARRY LEIBA: It redirects to the XML file, but that's the URL to reference.

ANDREW MCCONACHIE: That's true. And people—you never know—might want [to CSV].

BARRY LEIBA: Well, and IANA might change to JSON or something like that at some point because all the cool kids are using that. Except ICANN won't change to it until the cool kids are using something else.

JAMES GALVIN: Well, my favorite pet peeve in this space to always reference RFCs through the RFD editor website, not the IETF. Formal correct pointer is the RFC editor site.

UNIDENTIFIED MALE: Well, that's certainly what the ISPNs send you to.

JAMES GALVIN: Yeah. Anyway.

RUSS HOUSLEY: I always do the RFC editor as well.

UNIDENTIFIED MALE: I always point to the data tracker because it gives a lot more information, but …

JAMES GALVIN: Oh, it does. And there's no denying that. And I think, in an IETF context, I would generally do that, too. But in any kind of formal document, if you will, or something outside of all of that, I think you got to reference it the right way.

UNIDENTIFIED MALE: You're definitely right.

ANDREW MCCONACHIE: I decided a couple documents ago that we weren't even going to include links for our RFCs because everyone reading SSAC publications knows how to get to RFCs. So in the routing paper, there's no URIs for RFCs. It's just like this: "See RFC6762." And that's it.

BARRY LEIBA: Okay. And 6762, by the way, is the one that defines.local.

ANDREW MCCONACHIE: Yeah.

JAMES GALVIN: We don't have to have this discussion, but I have to be honest. I'm not sure I like that.

ANDREW MCCONACHIE: Okay.

JAMES GALVIN: And I say that only because we often have this discussion about, what is the audience of SSAC documents? And you start making assumptions about that and you can go wrong pretty quickly if things happen to get

around and go somewhere. I tend to prefer putting links in for things when you got a link. And I would personally prefer that we take that on. But anyway, sorry. Just my little thing. We don't have to have that here.

ANDREW MCCONACHIE: I'm open to whatever as long as we're consistent in the document. I don't think we have to be consistent across documents. So if that's the way we want to do it with this document, I'm fine with that.

BARRY LEIBA: Okay. Yeah, we can figure that out later.

ANDREW MCCONACHIE: Should I scroll down?

BARRY LEIBA: Yeah.

ANDREW MCCONACHIE: Any comments on these last two paragraphs? I was starting to get into why is it that developers make the decisions they do with that last sentence. Like, what are the incentives for developers to try and guess what a user means with regards to context?

BARRY LEIBA: Yeah. I'm good with what's there.

RUSS HOUSLEY: The incentive for the developer is an interesting question. They want you to use their stuff, right? Isn't that the incentive?

ANDREW MCCONACHIE: Yeah, exactly. And that's kind of where I was going with this last sentence.

BARRY LEIBA: So put another way, Russ, they want to make your user experience the best that they can. And they want you to get what you expect.

ANDREW MCCONACHIE: Right.

RUSS HOUSLEY: Least surprise. Principle of least surprise.

BARRY LEIBA: Mm-hmm.

RUSS HOUSLEY: Which of course is different if you have a different context.

BARRY LEIBA: Right. So I suppose that might be worth saying at some point, that getting the context wrong gives the user a bad experience and may drive the user to a different software product. Something like that. I mean, maybe that's just so obvious we don't need to say it, but …

RUSS HOUSLEY: May choose if there is an alternative, given that we just, two paragraphs earlier, were talking about Microsoft Active Directory. Their IT department may have picked that.

ANDREW MCCONACHIE: What do you mean?

RUSS HOUSLEY: I mean the user may not have a choice.

ANDREW MCCONACHIE: Yeah. Right. That's true. So it may not be the end user.

RUSS HOUSLEY: Exactly.

ANDREW MCCONACHIE: It's almost as if the software chooses the wrong thing for some definition of "wrong." But I ran into this where I just said the correct website and then put "correct" in quotes.

RUSS HOUSLEY: I think you should say “expected.” I’ve been thinking about that.

ANDREW MCCONACHIE: Yeah.

RUSS HOUSLEY: “Expected” might even be quoted.

ANDREW MCCONACHIE: Well, now we use that word. “Users have an expectation that putting a malformed DNS name in an omnibar will take them to what they think is the expected website.”

RUSS HOUSLEY: Oh, I see. Yeah. How about an “unsurprising” website?

ANDREW MCCONACHIE: Yeah. We could just say … I hear what you’re saying. You want to reference of this principle of not surprised.

RUSS HOUSLEY: Least astonishment, yeah.

ANDREW MCCONACHIE: Yeah. The principle of least astonishment. I think we should just call that out because I doubt people reading this will be familiar with that principle.

How does that look? [Best action].

BARRY LEIBA: I mean, I think that … yeah. Tara says in the chat, "'Predictable' is a helpful word." That's more what I was thinking. I mean, we can talk about the principle of least astonishment as a thing, but what I'd really like to say in this sentence is that it takes an action that the user would predict, like expects, maybe. Take an action that the user expects. So the software is trying to guess what the user expects.

ANDREW MCCONACHIE: [inaudible]

TARA WHALEN: Yeah. I meant that to be a very light touch. If we're trying to say "expects" in about three of four ways, we have a couple of other things in our toolkit to not overcomplicate.

BARRY LEIBA: Yeah.

ANDREW MCCONACHIE: Yeah. But it's predictably from the perspective of the user [inaudible].

BARRY LEIBA: Yes, exactly.

ANDREW MCCONACHIE: And the perspective of the user. “Should be [inaudible].” That’s better.

So what else do we want in this section? Or do we not want more in this section and we want to put stuff somewhere else?

BARRY LEIBA: Well, I don’t know. Okay, this section is basically saying, from the software developer’s perspective, the context is not always clear. So that’s the more stuff we want to talk about: how URIs have changed over time in how precise they are about … Well, the context is not always specified by the URI anymore.

ANDREW MCCONACHIE: Mm-hmm.

BARRY LEIBA: So the goal of the software developer is to figure out the right context so that the user is not surprised, so the user gets a predictable result.

RUSS HOUSLEY: By that, you mean that the scheme of the URI used to provide context, and now everything is running over HTTPS.

BARRY LEIBA: Yes.

RUSS HOUSLEY: So it doesn't provide any context.

BARRY LEIBA: Well, it sometimes does. It's just it often does not. So, yes.

ANDREW MCCONACHIE: But did the scheme of the URI provide context as to which resolution mechanism for naming should be used?

BARRY LEIBA: Well, yeah. Originally, if it said HTTP or HTTPS, you knew you were going out to the DNS to resolve that first blob. "Authority" is what it's called. And now that's not always the case.

ANDREW MCCONACHIE: Right. Or I'm thinking of something that wouldn't use DNS, like SMB, maybe?

BARRY LEIBA: SMB is an example. E-mail is an example. Mail.

RUSS HOUSLEY: Well, mail does, right? It just used a different part of the string that follows.

BARRY LEIBA: But it's looking for MX records rather than A records, for instance.

RUSS HOUSLEY: Ah, that's true. But it's still using the DNS.

BARRY LEIBA: Yes, but it's giving you a different context of how to resolve it, even through the DNS.

RUSS HOUSLEY: That's actually a really good point.

BARRY LEIBA: Yeah.

RUSS HOUSLEY: And I don't think we say that anywhere, that different records are tagged by different schemes or implied by different schemes. We need to say that somewhere.

BARRY LEIBA: Yes. In fact, that's a different context. Getting a different kind of record out of the DNS is a different context.

ANDREW MCCONACHIE: Mm-hmm.

RUSS HOUSLEY: Is a way that context used to be implemented.

BARRY LEIBA: Yes.

RUSS HOUSLEY: But now it's more complicated.

BARRY LEIBA: Yes.

RUSS HOUSLEY: That belongs back where we talk about context or where we introduce context.

BARRY LEIBA: So back in the bare section, 6.

RUSS HOUSLEY: I don't know where it fits best in the flow. Probably not there. But just make a note.

BARRY LEIBA: Yeah, make a note. Andrew will do his magic. And when we next look at it, it will be amazing.

ANDREW MCCONACHIE: So something on how context is actually not just about what resolution system is being used but also about what records are being asked for—so really the difference between asking for an MX record versus asking for an A record.

[BARRY LEIBA]: [inaudible] distinguish between [inaudible] how a domain name and e-mail address can signify. But is there really …

ANDREW MCCONACHIE: Yeah. It's the job of, I guess, the MTA to look at that e-mail address and say, "Oh, I'm going to look up an MX record. And then I'm going to look an A (or a quad-A)."

BARRY LEIBA: Well, the MSA (Message Submission Agent), yeah.

UNIDENTIFIED MALE: Yeah.

ANDREW MCCONACHIE: Okay. I mean, from the user's point of view, there's really no difference between a domain name that has an @ symbol before and a user before that. And say a domain name with just … that they might type into their web browser. But it gets handled differently by different software agents. I guess that's the point.

BARRY LEIBA: Yeah. I don't know that the average non-technical user really thinks about how the right side of an e-mail address relates to a domain name they might type in a web browser bar.

ANDREW MCCONACHIE: I think they want them to match. I mean …

BARRY LEIBA: Well, what I mean is I don't think the user really … The user doesn't think about them being the same animal.

[RUSS HOUSLEY]: You don't think so?

BARRY LEIBA: When you send something to barryleiba@computer.org, you're doing an e-mail thing. And when you go type npr.org in your web browser to

visit the NPR website, you're doing a different thing. And I think the average user doesn't really relate them.

ANDREW MCCONACHIE: But they know that the person who writes them from NPR is more likely to be a reporter. Like, the person who writes them from the npr.org e-mail address is more likely to be a reporter. Because they've been to that website and they know it's a news website.

BARRY LEIBA: Right. This is part of that study we need to have.

ANDREW MCCONACHIE: Yeah, it really is. But then you're right because it's not necessarily like they recognize it as a domain name. They just recognize it as an organization, right? Like a trademark.

BARRY LEIBA: Yeah. And I think they don't really see a difference between, say, gmail.com and google.com.

ANDREW MCCONACHIE: Mmm. They're just two different trademarks.

BARRY LEIBA: Yeah. So the people who use google.com are of course the Google staff. And it makes a difference from the point of view of things like [DKAN]

and DMARC and all of that stuff and doing spam filtering, but do end users really understand the difference? I don't think so. I could be wrong, but I don't think so.

ANDREW MCCONACHIE: Mm-hmm. No idea.

BARRY LEIBA: And just for grins, I really want to look at the IANA registry. This is [inaudible] that. There we go. And the trouble with the IANA registry is it's hard to go find what you're looking for. If you search for "URI scheme," you don't find what you're looking for. If you search for URI) scheme," you do, just because of how they have it hidden in there.

Anyway, the URI scheme. So I wanted to look at how many URI schemes we have, and I think the answer is—it's a technical term—a shitload. Yes, there are quite a few screens' worth of URI schemes. But one of them is not [drop]. And Russ will understand what that means.

RUSS HOUSLEY: I was about to say I bet half the people on this call don't get that reference.

BARRY LEIBA: Actually, that's not true. [inaudible] is one of the defined schemes—

RUSS HOUSLEY: Yes.

BARRY LEIBA: That's right. Because he wants it removed.

RUSS HOUSLEY: Correct.

BARRY LEIBA: Right.

RUSS HOUSLEY: It's like, dude, once you get it registered, you don't ever get it unregistered.

BARRY LEIBA: We can the rest of you all in at some point if you care, but beer has to be involved.

ANDREW MCCONACHIE: I saw the remnants of that mail thread, so I roughly know what you're talking about.

BARRY LEIBA: It's worse than you think.

ANDREW MCCONACHIE: I'm sure it is.

RUSS HOUSLEY: No. Pitchers of beer will have to be involved.

ANDREW MCCONACHIE: But you want somewhere some text … I don't know if here is the right place, but you want somewhere some text on how URI schemes used to be useful for denoting context. And they could be useful in the future, but they're currently not very useful for that.

BARRY LEIBA: Well, they … no. They are—

RUSS HOUSLEY: They provide some context but not enough.

BARRY LEIBA: Right. They are often still useful. It depends on the URI scheme. But there's been enough subversion of it that more information is often needed.

ANDREW MCCONACHIE: Yeah. Sure. I mean, I think we could talk a lot about that.

BARRY LEIBA: Right. And I don't think we want to get into a lot of detail about it.

ANDREW MCCONACHIE: Mm-hmm.

BARRY LEIBA: It's just more of a—if you'll excuse the expression—context-setting thing.

ANDREW MCCONACHIE: Right. And it really is about HTTP. That's the one we want to talk about. That is the bad boy.

BARRY LEIBA: Right. That is the main one that has been overloaded so much that it does not provide the context we need in many cases.

ANDREW MCCONACHIE: Right.

BARRY LEIBA: But, for instance, there is a DNS URI scheme.

ANDREW MCCONACHIE: Ah. Can you paste the link to that IANA page in the chat?

BARRY LEIBA: Sure. Doing it now.

ANDREW MCCONACHIE: We can all marvel at the thousands of URI schemes that never get used.

BARRY LEIBA: Ha. A lot of them are provisional. I didn't realize that there were that many provisional URI registrations. I haven't looked at this in a while. TN3270. Ha.

Anyway, we can waste a lot of time on that.

ANDREW MCCONACHIE: Okay.

BARRY LEIBA: I always forget that these things are recorded.

Okay.

RUSS HOUSLEY: I guess my raspberry will be preserved forever.

BARRY LEIBA: There you go.

ANDREW MCCONACHIE: Okay. Unless we want to talk about Section 2, I don't have anything else to talk about.

BARRY LEIBA: Yeah, I think it's a good time to take a break. And I will remind everybody to read all of this so that, when we meet next week, we really have—it is next, right?; yeah, it is next week—a running start.

RUSS HOUSLEY: Next week face to face. Or some of us will be face-to-face.

BARRY LEIBA: Yes.

ANDREW MCCONACHIE: And for next week, we kind of want to do an overview of the table of contents a little bit and get some feedback from the SSAC?

BARRY LEIBA: Yes.

RUSS HOUSLEY: Yeah, I think that's right.

BARRY LEIBA: Yeah, I think that's the starting point.

Tara, will you be there in person?

TARA WHALEN: Sadly no. I'm packing up my stuff [inaudible].

BARRY LEIBA: That's what I thought because [I think we] discussed that in KL.

TARA WHALEN: I've missed you all. If I can manage to do remote participation in fits and starts, then I would like to. But I deeply regret not being there.

BARRY LEIBA: Yeah, it'll be sad, but we'll see you again soon.

TARA WHALEN: Fingers crossed.

BARRY LEIBA: Yes.

ANDREW MCCONACHIE: Good luck with your move.

TARA WHALEN: Thank you.

BARRY LEIBA: Alright. So let's call it a call. Thanks for coming.

ANDREW MCCONACHIE: Okay. I'll see you folks next week if you're there. If not, I'll see you when I see you.

BARRY LEIBA: Okay. Bye-bye.

[END OF TRANSCRIPTION]