KATHY SCHNITT: Welcome to the SSAC Evolution of the DNS Resolution Work Party teleconference on Thursday the 7th of April 2022. Back over to you, Barry.

BARRY LEIBA: Thanks, Kathy. Okay, so where we left it last time is that Andrew was going to knock out a initial outline of what the document might look like so we can start filling in topics. And start discussing those topics. So Andrew has done that. I guess Kathleen is displaying. So let me switch to that.

Okay. So stop at the table of contents. And let's just start with that. Just kind of a overall outline of how we might put this forward. We looked at two different ways of arranging it. And I think what's here is good for getting us discussing things and getting topics written down. We might choose to do some reorganization of the document later and change the outline.

But for now the usual sort of executive summary introduction stuff and then a section that says here's how things briefly, briefly—because this is not intended for complete novices—but here's how DNS has worked. Then a section on here's why people thought we needed to change some of the aspects of that. And then a section on what those changes have been, are being, will be. What the technologies are. And then a section on what the implications of those are.

When we were discussing this Russ Housley suggested that maybe we want to have it be here's a topic and here's why we did it. Here's the technology. Here's the implications, separately for each rather than having it laid out like this. And we might shift in that direction or to some hybrid. But I think what we have here will get us started on discussion.

Does anybody think this is completely wrong-headed or have other comments on this overall outline? Okay, so something to start with. And I propose that what we should be doing now is forget about the first three sections. Start with the executive summary, introduction, and traditional DNS resolutions bits for now. And spend our time looking at Sections 3 and 4. What were the reasons to evolve this? And what technologies are we going to discuss in this document? So Kathy, please scroll down to Section 3.

And Andrew has started with notes from the discussion we had last week and done perhaps some of his own things. And I guess, Andrew, you can speak to that.

ANDREW MCCONACHIE: Yeah, so these are just kind of collections of things that I've heard. And I kind of roughly categorized them. Very rough, as you can tell. Just to kind of start the conversation off. I don't know. Should I walk through them all?

BARRY LEIBA: Yes.

ANDREW MCCONACHIE: Should I walk through what I have under Section 3?

BARRY LEIBA: Sure, let's do that.

ANDREW MCCONACHIE: Sure. So under Pressures, I've heard a lot of talk about incentives. So different kind of incentives for vendors and operators to change how DNS resolution works. So kind of a question of why is there all this effort being expended. And then in terms of privacy I heard a few things. Users want privacy but I think it's not necessarily—thank you, Geoff, for saving me. Go ahead.

GEOFF HUSTON: I actually don't think users want privacy. I think that's an invention of the IETF. But that was just done as bait.

There is an issue about when you have this distinction between operators and vendors. You've also got to think about applications. And it's almost a third leg of this problem. And part of the motivation to actually encrypt the DNS channel in resolution is to actually make sure that middleware doesn't intercept you. That you're not trapped by legacy infrastructure. Because for innovators who actually wish to alter the model of resolution from operating system platform to CPE resolver, stub forwarder, to recursive blah, blah, blah, you can't do this and use Port 53 UDP often.

You actually have to break out of that entire model. And privacy allows you to innovate inside a channel that tunnels through existing infrastructure, right? It's sort of the VPN of the DNS for innovators and applications. Tunneling through legacy or existing infrastructure.

ANDREW MCCONACHIE: Am I running the queue? Warren, go ahead.

WARREN KUMARI: I think also couple facts that we're missing. First off, I think the title needs changed to Pressures to Evolve the Traditional DNS Resolution Model. Because otherwise there's a whole bunch of other things we need to add into this.

ANDREW MCCONACHIE: Yeah, sure.

WARREN KUMARI: And privacy is a small portion of what's driving a bunch of the encrypted DNS. But a huge and, I think, much larger amount is not the privacy part. It's censorship resistance. Like it's only a little bit that people are concerned that the ISP is going to spy on your queries and much more so that the government is going to block your ability to access censored material that people think they shouldn't.

Another big one which we left out is split DNS, right? The names within ICANN.org resolve differently to names outside ICANN.org or MyCompany.com. and we haven't mentioned that.

BARRY LEIBA: So I understand that that will come up later in implications. But is there really a pressure on that? VPNs certainly do this now. And I don't think that's been any factor in a move to change things.

WARREN KUMARI: I think I don't understand your question. I mean it's something that happened.

BARRY LEIBA: Yeah, I don't look at that as a pressure to evolve it. I look at it as an effect of the evolution. Let's see what Suzanne has to say, sorry.

SUZANNE WOOLF: Yeah, I think split-horizon naming and a bunch of other stuff, these aren't driven by privacy so much. But they're driven by people being interested in security in the sense of protecting assets inside their enterprise scope, making it harder for intrusions to happen, and so on and so forth.

It's almost immaterial how much difference these specific things make as far as do they make your operations more secure in meaningful terms or not. But, for instance, people have been told that split-horizon

DNS is more secure. It protects you against information leakage. And even if you don't believe that, which not everybody does. It was driven by specific security concerns that are not related to privacy so much.

WARREN KUMARI: And related to that there's also other ones, like we haven't actually really mentioned in here adding DNSSEC. That was a big one.

SUZANNE WOOLF: Yeah. Sorry to step on you, Warren, but just to close out the thought. What I put in the chat was maybe add security as a separate topic. Not all security concerns are driven by privacy.

WARREN KUMARI: Thank you, yeah. And, yeah, I think under security there's DNSSEC. There's things like aggressive NSEC. I don't know where we put in things like having local copies of the zone. Things like local hosts, sorry, local roots-type stuff. I don't know if that goes in security.

RUSS HOUSLEY: But in doing those things I think we need to be careful not to mix confidentiality and integrity. They're two very different things, especially in this environment. I mean, DNSSEC makes no attempt to address confidentiality. Whereas DoT and DoH do.

WARREN KUMARI: Yep.

GEOFF HUSTON: I think your issue about local hosts, Warren, is well taken. Part of the traditional DNS resolution model is the information is flooded on demand, just in time. And TTLs do caching. But we never did just in case in the DNS until we started doing zone distribution with local hosts.

And it's actually quite a fundamental change when you do just in case. Because now I'm provisioning the entire zone into the intermediaries in the DNS without an initial triggering individual query. And so it's not local host per se that's actually the evolution. It is indeed the model of provisioning the intermediaries with a here's the lot just in case you get a query. Which is quite a major change to the DNS model.

ANDREW MCCONACHIE: And Jacques's had his hand up for a while. Jacques?

JACQUES LATOUR: Hi. So the operators want to shave milliseconds. Isn't that the browser vendors?

BARRY LEIBA: Yeah, we don't have efficiency or whatever you want to call it. It's not really efficiency. Speed.

JACQUES LATOUR: Of the browsers or, not sure. I think that's a first-line incentive to meet the DNS resolution. That's a browser vendors trying to evolve the name resolution to something else.

BARRY LEIBA: Yeah, Tara put the right word in, performance. That's the word I was groping for. Yes, that was a large reason for DoH in the first place was the browser vendors wanted to be able to build it into the browser to save, as Jacques says, milliseconds.

WARREN KUMARI: The browser vendors really weren't interested in doing DoH, right? DoH was largely Paul Hoffman was like, "Ooh, we can do DoT. And it's really easy. We'll just do it wrapped in HTTP." And DoH is not faster for the browsers.

BARRY LEIBA: It allows the browsers to build it in directly rather than having to call out to the OS.

WARREN KUMARI: And the browser vendors had already done that largely.

BARRY LEIBA: Okay.

WARREN KUMARI: Anyway. But, if we care, we may also want to have the, wow, I had the words when I started that sentence. And then I lost it midway through.

BARRY LEIBA: I hate it when that happens.

WARREN KUMARI: Oh, sorry, protocol differentiation, right? Like P4 and P6. There's all sorts of crazy that's added in to try and make P6 go faster than P4.

ANDREW MCCONACHIE: Okay, Suzanne, did your hand go back up?

SUZANNE WOOLF: Yeah, it did.

ANDREW MCCONACHIE: Okay, go ahead.

SUZANNE WOOLF: Noting Russ's comment a couple minutes ago about the different parameters of security and so on and also the question about local caching, and I would put any cache in the same category also. What we've got as security here I'm almost tempted to expand to risk management. Where there's a particular set of threat model is sort of implied by the things that people have been optimizing for.

WARREN KUMARI: Sorry, can you say that again in different words? Because I—

WARREN KUMARI: Risk of what?

SUZANNE WOOLF: Exactly, thank you.

WARREN KUMARI: That's what I didn't understand from what you—

SUZANNE WOOLF: Okay, so what do you care about? You care about information leakage or you don't. You care about speed and performance or you don't. You care about resiliency against distant failures or you don't. And the thing that got me thinking this was the local caching piece, which I think also goes along with any cast on the authoritative side, well, either way really. What are you protecting yourself against? You're protecting yourself against unmanageable dependencies and distant failures.

So for all of these things, there's something in particular people want to protect themselves from or change. And some of those things are we want it to work better for more of the things we're trying to do. But we want it to be less likely to fail and more reliable is kind of a different parameter.

WARREN KUMARI: Okay, I think I understand what you're saying.

SUZANNE WOOLF: Yeah. And there's an implicit threat model behind a whole bunch of these things or several implicit threat models. And that might be a different way to slice it that's of some …

ANDREW MCCONACHIE: Okay, Geoff.

GEOFF HUSTON: I'm going back to performance. And there is a trade-off between performance and performance. What I mean is, if you have very generic queries, they're eminently cacheable because the query is shared across a number of different sort of points of wanting that information, irrespective of the particular motivation. So caching is brilliant. You don't sort of re-query all the way up to the authority.

But because of its lack of specificity, you've actually got to do multiple queries sometimes to get what you need. Think DNSSEC validation. Think EDNS client subnet. Think about what we're doing with the SVC records. And what we're actually doing is saying, "Well, if I can put more data into the queue name, into the query, I lose on caching performance, big frown, but I win on multiple queries. Because there aren't multiple queries anymore, I can, if you will, preload the information I'm after into the query. And get back a single answer."

And so this whole trade-off in traditional resolution of generic queries, generic answers, queries are cheap. Do lots to navigate where you want to go versus I'm going to sharpen out my focus. I'm going to tell you what my end goal is, underscore port, underscore server's name, blah, blah, blah, blah, blah. And you're going to tell me how I rendezvous with that service, single transaction.

So that's not really cacheable because all of a sudden the query's got very specific. Frowny face. But I don't need follow-up queries. And that's really cool. And so this whole issue of what's a traditional resolution model is actually now quite interesting. Because the traditional model was what's the IP address of this name. Why do you need that, not telling you. It says give me one piece of data and let me work on it.

The new model is I'm going to tell you what I'm going to do and you're going to help me get there. So, yeah, Russ, I agree with the comment. It's metadata that can't be known ahead of time. And for this it leads to the next issue. Do we have a section on answers, authoritatives? Because inexorably the model that I have a zone file, it's a static file and I serve off it got put under intense pressure from dynamic signing with DNSSEC. And is now being put under intense pressure with SBCB records. Where dynamic answers are, I think, an increasing part of DNS.

So I don't have a massive great fat text file. I actually answer on demand because you unique customer has queries and I assemble the data and answer you. It's not statically there. And again I'm trading multiple queries against single omnibus queries. And they're both performance

oriented but it's different aspects of performance it's trading off against itself. Thank you.

RUSS HOUSLEY: Geoff, you've actually gone a step further than that, I think, and said it looks like a DNS query but it's really not what we meant in the early days. it's something else. It's just using that same protocol.

GEOFF HUSTON: Right. The query name is now a piece of let's call it microcode for want of something better, a program, that directs the answerer to perform various functions on your behalf. And offer you results based on that computation. It's no longer just static file. Everyone gets the same answer. They don't. And the trade-off—

RUSS HOUSLEY: And they don't want them to, like in a CDM environment.

GEOFF HUSTON: Correct. And we're trading off the big asset that the DNS used to have—we couldn't make this work without caching—into we feel confident we don't need that level of caching. We can drop down cacheability and increase accuracy. Decrease the query rate because now I'm actually answering the answer to the question you really wanted to ask. So you're not pummeling me. Single question, single answer, yay.

It's a big issue. And I think it's important. It's not that I'm decrying it or saying it's really good. I'm not. But it is a real shift in the landscape of what the DNS is.

UNIDENTIFIED MALE: Jacques?

JACQUES LATOUR: Yes. I put a comment in the chat to facilitate this. But basically are we specifically on the DNS resolution? Because the web three, we're talking a name space provider like ENS and Cloudflare and Opera. You can put the thing in the browser bar with what appears to be a domain name. And it goes somewhere. But it's not DNS.

WARREN KUMARI: Well, but you're assuming that the browser bar is a thing that does DNS lookups.

JACQUES LATOUR: Everybody would assume that, correct?

WARREN KUMARI: I think that almost everybody in the world would assume that the thing you put in the DNS browser bar is a query, a searched query. And if you happen to search for the string Microsoft.com it will return the I'm Feeling Lucky version of that.

JACQUES LATOUR: Or if you put something .AlternateNameSpace it goes there, same way.

BARRY LEIBA: It does whatever the browser does with that name that you pull in with that string that you put in there. I remember reading at one point that one of the most commonly searched-for things in Yahoo! was Google. Because that's just what people put.

JACQUE LATOUR: But today are we past the traditional DNS already?

BARRY LEIBA: I've been thinking—and this is one of my thoughts about its more generic TLDs—is that people don't use domain names much anymore. They use search queries.

WARREN KUMARI: The vast majority of things that people put in the browser bar is not a DNS by name.

BARRY LEIBA: People don't put Starbucks.com. They just put Starbucks. And they let the search engine take care of it.

JACQUES LATOUR: What they do is today a regular user with Opera set up to his online Cloudflare will click on what appears to be domain names, something.ens or something. Whatever the alternate name space is. And it will resolve and connect to a different resource on the Internet.

WARREN KUMARI: I suspect I'm completely not understanding what Jacques is saying, which is unfortunate because it sounds like something important. But I just don't—

JACQUES LATOUR: Perhaps it gives towards the [inaudible] decentralized DNS. [inaudible]

WARREN KUMARI: Sorry, are you saying the fact that if you enter Microsoft.com into the browser, you're going to Microsoft. But if you enter wkumari.eth.link you'll go to the Ethereum naming space.

JACQUES LATOUR: Yes, so the Ethereum naming space, however you resolve it appears—

WARREN KUMARI: Okay, so that's specifically why I said pressures to evolve, that I wanted us to add the word traditional DNS in this section. Because what you're saying is a real problem. And something in a complete whole section.

RUSS HOUSLEY: I think the point that he was making is that, to the user, it looks like a domain name. And he doesn't have any clue that it's not.

WARREN KUMARI: Yes. And that's a major problem. And I think that that's—

RUSS HOUSLEY: Of course, the people doing it would tell you it's a major feature.

WARREN KUMARI: Yes. But—

JACQUES LATOUR: That's a pressure that impacts the DNS today. Because [inaudible] is a risk presumably in the .namespace, the alternate namespace that Opera is configured for. If CIRA.ca we don't register our .ca domain with them, there's a chance that .ca's are going to resolve to something else instead of us with a very specific browser that's set up to look at the different name space.

WARREN KUMARI: Yep, and I agree that's a major problem. I think that it's an entire whole section which is pressures to evolve the resolution system. And I think that it's really important and we should talk about it. It has me concerned that I don't know if we capture it in Section 3 or if that's Section 4, which I think is a huge section as well.

RUSS HOUSLEY: But it's not really resolving—I'm sorry, that's the wrong word. What it's doing is forking the resolution to support multiple mechanisms. It's almost like they're multiplexed together in some kind of bizarre way.

WARREN KUMARI: Here we go. Click on this link.

RUSS HOUSLEY: I dare you, right? Are you phishing us?

WARREN KUMARI: So what this is is the link DNS—

GEOFF HUSTON: Oh, my god. I clicked on the link. And my machine crashed.

WARREN KUMARI: Sucks to be you. Get a better machine.

JACQUES LATOUR: [inaudible] installation.

WARREN KUMARI: So the link TLD, which is an ICANN-accredited blah, blah, blah, has a second-level name in it, ETH, which is literally an interface into the

Ethereum naming service. And so if you put in some Ethereum name.eth.link, that is a—

RUSS HOUSLEY: Firefox screams at you that it's insecure. Don't go there.

WARREN KUMARI: That's because I haven't bother trying to… Oh, yeah. So does Chrome. That's because the cert that I have there is for wkumari.com because I couldn't be bothered trying to make whatchamacallit work, thingy.

RUSS HOUSLEY: [ELS] it's called.

WARREN KUMARI: Yeah. Let's Encrypt is the word I was looking for. I'd have to do things in a weird way to make them be willing to actually issue for ETH names. But, yeah. And I mean, Jacques is right. It is a name that looks like DNS but it's not using the DNS resolution model.

JACQUES LATOUR: What it comes as, I think—

WARREN KUMARI: Users don't know that. And that's a major problem. And actually, I mean, to be honest, the wkumari.eth.link is actually an ipfs:—I don't remember the full key—eth. So that's the same thing. If you watch you

actually see a redirection through that. And, yeah, that's a major issue. If you like, we can stick it in this section. I think it's a whole separate section that requires its own set of description as well. And that we need to cover all of this in that. Whether you want it in Section 3 or a whole big separate Section 4 I don't think matters at the moment.

JACQUES LATOUR: I think people need to understand the risk, the issue, the risk, the impact around this. But my computer was set up in a way that when I click on the ipfs.[inaudible].eth, it connected and opened up a webpage. And I—

WARREN KUMARI: I don't know why you sound like you think I'm disagreeing with you.

JACQUES LATOUR: No, but I was surprised that it actually worked, so.

WARREN KUMARI: Yep. It's terrifying. And this is viewed as a feature by—I don't know if you saw my presentation on the .crypto TLDs, the blockchain ones? I've got a whole presentation on this that I did to a set of people. And to me it's a hugely concerning issue.

RUSS HOUSLEY: So I think what we need to capture here at this high level is other systems are purposefully making their names look like domains.

WARREN KUMARI: Yes. Yep.

RUSS HOUSLEY: Right? And then we've got these examples, Ethereum and so on.

WARREN KUMARI: And I mean they have to make them look the same as the DNS in order for them to work, realistically.

GEOFF HUSTON: The lecture on the risks, the harms, etc., is probably another document. This is not about the threats and opportunities, etc. This is actually I would have thought more about this is the landscape that we see without sort of getting into what's good, what's bad per se. Because once you start good and bad you could go on forever.

WARREN KUMARI: I mean, yes and no. I think that there are certain things which you can relatively definitively say are not good. Like, for example, you could say cyanide exists. But if you're going to mention that I think it's probably implicitly sort of something you need to include that and you probably shouldn't drink it.

GEOFF HUSTON: But the link you posted has a whole bunch of underlying issues. But on the face of it it's just an interesting in a piece of innovation in resolution. Yes, this is happening. There are seamless mappings or invisible mappings going on that put you into other realms and spaces without you even being aware.

Sometimes breaching disparate realms is a good thing. Christ, the 1980s was all about that, making bloody BITNET mail talk to this mail and that mail. Sometimes doing that is a security nightmare. Understanding those trade-offs and what's going on is the work of another document, I would have thought. Because it's such a big piece of work. That's all.

WARREN KUMARI: Yeah, I mean, I must admit that that was much of my concern when we started this is that I had thought the purpose of this work party was originally specifically to talk about the evolution of resolution. Being the fact that a lot of it is moving from the DNS or that there are multiple resolution contexts. And that there is no way to differentiate them.

And one of the concerns, if we're going to talk about alternate resolution systems, is the GNU naming system. One of their big features is that it uses the same name space as the DNS. So with the GNU naming system, you can have www.foo.com. And those names work. And if there is no foo.com in the real DNS, it will be resolved in the GNU naming system. And depending on how you configure it, if there is a foo.com it will be resolved in the GNU naming system. So GNU is internationally—and they view this as a feature—squatting on the entire DNS name space. And—

JACQUES LATOUR: Exactly. Well, that's clearly—

WARREN KUMARI: Well, Ethereum and crypto at least have the advantage that they end in a string which looks different, potentially. And people might learn that .crypto is terrifying. But if you look at a name like www.foo.com, there is no way of knowing which naming system it is intended to be resolved using.

GEOFF HUSTON: I think those are valid points, Warren, absolutely. I just think you're opening up an awfully big space. If what we're doing is opening up a big space and cramming it into one document, it's going to be an uncomfortable fit. That's all.

BARRY LEIBA: We could decide at some point that this becomes two documents. I think right now our goal should be to just throw out topics and get them in there for discussion.

WARREN KUMARI: Okay.

GEOFF HUSTON: Well, that's a perfectly legitimate topic in that context then, absolutely.

WARREN KUMARI: So you think the high-level thing of it is interactions with other resolution systems might be a good bullet point for it.

GEOFF HUSTON: Yeah. There was an observation Patrik Faltstrom made many years ago that, if you squint hard enough, most of these systems look like the DNS. Or the DNS looks like most of the others. The corollary is in theory you could map them all together. In practice that would be bloody stupid.

WARREN KUMARI: Yep. I mean, you kind of have to have them look the same. As part of the—

GEOFF HUSTON: That was Patrik's point.

WARREN KUMARI: Yep.

SUZANNE WOOLF: They're evolving to meet comparable names. And they're running into comparable problems, so they're going to—

WARREN KUMARI: Yeah, but I mean it's not only that. I spent a huge amount of time in this, including with Patrik Faltstrom and John [Krenzen] and Suzanne and Ted and many others on this whole thing. And one of the common questions that people come up with is why don't these other naming systems simply use names that don't look like the DNS? Why don't they have www.foo.!!!? And you can't have that if you want these names to work in normal applications. They just wouldn't work. And so those name resolution systems can't actually realistically differentiate themselves through a reasonable switch.

Anyway, sorry. We're getting—

[BARRY LEIBA:] Another version of—

RUSS HOUSLEY: Warren, you're saying that—

[BARRY LEIBA:] Universal acceptance.

RUSS HOUSLEY: You're saying that the Internet applications have been so successful that that name form has to be embraced. Instead of replacing dots with colons and see where that gets you, right?

WARREN KUMARI: Yep.

ANDREW MCCONACHIE: I kind of stopped taking notes in the document itself. Because I'm taking notes in a Notes document. So I'll go back and update this later. So don't worry if your idea isn't on it.

BARRY LEIBA: I'm not seeing any hands other than Warren's perpetual one that we need to have Zoom automatically take the hand down when you unmute or whatever.

WARREN KUMARI: Yep.

BARRY LEIBA: It looks like we have a good list of things to bat around.

WARREN KUMARI: I suspect that the list could be substantially longer if we were to look at the list of thoughts that are open in the DNS Sub Working Group.

BARRY LEIBA: So one thing that I would like to see—

WARREN KUMARI: Actually, hang on. One of the big ones that is actually important we left out is a significant one, which is things like multicast DNS. Like my phone connects to my printer using the DNS but it's in a local context. And this has required a fair amount of futzing in order to make the DNS work. And it's not a different resolution context. It's just a weirdly local version of DNS. I don't know what we call it. DNS as a discovery mechanism maybe?

BARRY LEIBA: Sounds like a good way to refer to it at least for now.

Okay. So one of the things that I wanted to have us do as we proceed with this discussion is to get some work done in the document in between calls. Well, but specifically Andrew's going to be building out the document. But he's going to take his interpretation of what you said and his notes from what you said for all values of you who are out there. And I'm not talking to Warren who happens to be highlighted on my screen right now.

But you know what you said and you know what you want to say. And we need to have that input also. So as Andrew starts putting in notes of things and his interpretation of things, we need the people who were on the call to go and look at the document in between. Like if the call's on Thursday, have a look at it on Tuesday and flesh out some of the stuff that you had put in or that you would like to add to. So I really want to encourage everybody to get into the habit of doing that so that we can get more meat in here than Andrew's able to put in from notes from what we've said on the call.

RUSS HOUSLEY: I really like this idea and I'm hoping that people will find 15 minutes here or there to do that so that we can have a meatier discussion the next time. Because it will be more than just where we left off last time. It does, of course, require you to read the things that have been added since the previous call shortly before the call next week. So it really does require time commitment of both people with ideas and the people who want to react to those ideas.

BARRY LEIBA: Any comments on that? So you're all on board with this, yes? Yes. All right. I've got something in the chat. Merike puts a chat-based thumbs up.

All right, so I think we've had a good discussion here. We've got some stuff in the document to think about and discuss next week. I think unless there's something more we need to do right now we should call the call at now and just, say, spend some time over the next week typing stuff in here. Even if it's just notes and not full sentences, something to feed the discussion. Okay?

Does anybody have any last words before we wrap up?

ANDREW MCCONACHIE: Do I have any marching orders?

BARRY LEIBA: I think just merging whatever notes you took outside this document that you want to add to it, but otherwise no.

ANDREW MCCONACHIE: Okay.

BARRY LEIBA: Russ, anything you want to say before we call it?

RUSS HOUSLEY: No, I think this was a good session. and just what's on the screen is thought provoking.

BARRY LEIBA: Yeah. So thanks, everybody. Thanks for your input. Thanks for coming on the call. And we will see you all next week. And we will see your comments in the document before that. Toodles.

[END OF TRANSCRIPTION]