

---

KATHY SCHNITT: Welcome to today's NCAP Discussion Group call on April 15<sup>th</sup> at 1900 UTC. In the interest of time, there will be no rollcall. Attendance will be taken by the Zoom list. Kim and I will update the Wiki with the names of the participants as quickly as possible.

We do have apologies from Ros Rasmussen, Julie Hammer, and Matt Larson. All calls are recorded and transcribed, and recordings and transcripts will be published on the public Wiki. As a reminder, to avoid any background noise while others are speaking, please mute your phones. And with that, I'll turn the call over to Matt. Matt, please begin.

MATT THOMAS: Thank you, Kathy, and good afternoon, everyone. I hope you've had a nice week. I hope Jim will be able to join us a little bit later. I did talk to him about an hour ago on the phone so I know he does plan on attending. He must just be wrapped up with something. Hi, Patrick. It's nice to be both co-chairs here where we're all together. On with the agenda.

So, does anyone have any updates to their SOIs they'd like to make? I see no hands so we'll consider that a no. There are no new members this week, which is fine, but if you are knowing anyone who is interesting in the name collisions subject, please feel free to reach out and suggest that they join the group. Always looking for more participants.

The first, I think, material item that we're going to cover was the view of the document from last week. If we could bring that up, and maybe also paste the link into the chatroom for everyone?

---

*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.*

JAMES GALVIN: I'll just jump in and say hi, Matt. Everything's working great until the moment you try to join the meeting, and then my Internet decided it wasn't going to work. I'm sorry. I apologize for being late. It has taken me six minutes to get here. I apologize for the interruption. You were already going, so have at it.

MATT THOMAS: No problem. I'm glad you made it. Yeah. So, this is the NCAP gap analysis brief document that we've had out for suggestions or edits for a week or two, now. I don't see any edits in the document. That being said, if anyone wants the opportunity to make any suggestions, please feel free to go in, put suggestion mode on, and make those.

I don't know if we've set a definitive time when we're closing this but I believe we were going to try and submit it, I think, this week, maybe after the admin call tomorrow, in a PDG form out to the mailing list to memorialize this as a work product out of the discussion group. So, maybe just kind of a last call on that item, here. Jim, did you have any other thoughts or concerns on this?

JAMES GALVIN: No, you're absolutely right. Actually, we were going to try to "last call" it last week and we kind of neglected to do that on the mailing list. So, yeah, just ditto from last week. We'll actually try to get it done, this week.

---

MATT THOMAS:

Sounds good. Switching back, then, I think, out of this topic into the regular agenda, we wanted to take this moment to talk about the general plan and cadence of this discussion group going forward for the next nine, 12, 15 weeks as it is. We wanted to do a little bit of establishing a little bit more of a repeatable pattern or a cadence to having these meetings.

And so, we talked a little bit about how we want to go about that. We kind of came up with two overarching themes that we want to capture going forward when the discussion group is working.

The first one is that we want to capture relevant bibliography references out of work product one, out of Karen's report. So, as we go forward and we discuss the board's nine questions, we should really be looking in that document, and seeing appropriate content, and making notes of those references, and leveraging that document to the best of its extent. It's a very well-written and thorough document, so we should make good use of it.

And the second overarching theme is, what do we want to ask of the data going forward in studies two and three? Specifically, what questions can we ask of the data to specifically address or answer some of the nine board questions?

And so, to that point, we were thinking, over the course of the next nine, 12, 15 weeks or whatever, that we're simply going to go question by question and iterate down. And so, for one working session we would just dive into question one for a week or two, and then move on.

Of course, there is some flexibility. If someone really wants to jump on a question or thinks something's a higher priority or needs more attention,

---

we obviously have the flexibility to do that. But at least this way, we'll have a repeatable, standardized discussion group format and people can prepare in advance for the upcoming topics and do a little bit of research and reference-gathering in advance. Jim, did you have any other thoughts on that or want to comment a little bit more?

JAMES GALVIN:

No, you're hitting it. That's what we talked about. You're doing great.

MATT THOMAS:

So, to that end, we have a document up in the shared drive, and it's called "study two analysis notes." I'll paste this into the chatroom right here. Hopefully, all of you can see that. I can bring it up on the screen, as well.

The document contains the nine different questions the board has put forward, and it has also captured the notes that we had going in that Excel Spreadsheet, underneath them. Hopefully, we can use this document as a little bit more of a collaborative space to enumerate and expand on each one of those questions.

I went through and added some substantial notes, or comments, or thoughts under questions three and question seven. I'd encourage you all to take a look at it, at some point, and maybe we can address or discuss at some point in the future.

But I think, going forward, we'll probably just start down the list, look at question number two right away, and then start working on that. I don't know if we'll get to it today or if we're going to plan on that for next week's call, but we'll just simply work our way through these questions

---

and talk about, exactly what could we do with the data in studies two and three, and what questions can we ask of that data to help answer these questions?

Does that make sense to everyone? Anyone have any questions or concerns at this point? Seeing no hands, so I guess that sounds like a good plan.

So, maybe if we do pull up number two right away. We were looking at what the role of negative answers currently returned from queries to the root for the [strengths, play], and experience of the end-user, including the operation of existing end-systems.

I think we have a couple of notes down there about looking at NSEC caching and QName minimization, as well as other technologies that have been employed out into the ecosystem that might disrupt some of this behavior, such as hyper-local roots or things of that nature. Does anyone have any thoughts of what kind of questions we might want to ask out of the data to specifically address this question? Yes, Jim. I see your hand. Go ahead.

JAMES GALVIN:

Yeah. Maybe I'll just jump in a little bit. I appreciate that folks might not have felt like they got enough advanced notice about exactly how we want to accomplish and what we want to do, here. Matt and I had synced somewhat late on how to approach these things and move forward.

There is a part of me here that doesn't really want to go too far down this list right here, today. I would suggest that maybe we try to see if we can't

---

get a conversation started, here, about one of these questions. And maybe we end our meeting early today if nothing is forthcoming from the group.

The model here is this really is, now, our discussion opportunity. This really is, now, our analysis part. We are moving into study two of the project plan, overall in this project. So, this is our opportunity to think about how we're going to approach data that we hope to get access to and how we're going to approach answering these questions.

We want to encourage everyone to say whatever is on your mind. This is a brainstorming session. The rest of this next nine, 12, 15 weeks as we start to kick this off is all just open-ended brainstorming discussion. Feel free to toss anything out. Go into this document. Just add text, whatever stream of consciousness you've got for right now. We're trying to figure out what the best way is to move forward, here, and to kick this off.

So, I'll just expand a bit on what Matt had said earlier and make this very concrete. You can look back at the JAS and the Interisle reports from 2012. They, obviously, had a set of questions that they asked of data. We can think, here, about gathering up those questions and distributing them inside each of these nine questions left to be answered to the board.

And so, thus, as an indicator to ourselves, "Gee, we want to ask these questions of whatever data we can get access to to help inform how we're going to respond to this particular topic area of the board."

The idea at the moment is to collect what we're going to do here, and then we'll try to, once we start having some thoughts down here, some

---

bullet points, some provocative potential answers or, at least, more questions to ask, then we'll try to structure actual task work, here.

I mean, hopefully, we will actually get a contractor who will go out and do some of the research for us that we're going to set up for ourselves, here, so that they'll bring data back for us and actually sit, and then look at, and really think about and figure out what the data means in the context of answering these questions.

So, honestly, going forward, this is the time when we expect, and we need, and, frankly, it's a requirement – if you're a panelist in this group, you're supposed to be contributing.

So, now is the time to think about actual work and for people to do some preparation to really go back, look through the bibliography, figure out what to bring forward. I'd much rather that we come prepared and we just start ... Other people should be talking besides just Matt, myself, and Patrick in trying to pull this stuff forward.

Let's try to do that. I don't really want to sit here and try to tease out of people thoughts. The emphasis here is on coming prepared to start talking about what these topics mean, and what we can do about them, and what we can say about them, and draw from the bibliography. So, go back and figure out what from bibliography works, here, and pull that forward.

Jeff is writing a comment, here, Jeff Neuman, in the chat. Yeah. So, Jeff, why don't you ...? As long as we have some questions and some things to talk about with this one, let's focus on this and let's bring some discussion out.

---

And so, Jeff, you wrote something in the chatroom, there. Why don't you pick that up and talk here? See if we can get a little discussion going on item two, here, and we'll just continue as long as we have contributions and something to talk about. So, thanks. Matt, back over to you.

MATT THOMAS: Jeff, please go ahead.

JEFFREY NEUMAN: Yeah. Thanks, Matt. Thanks, Jim. Can we also, as we go through each of these questions that the board asked, try to figure out the "why." Why are they asking this question? And then, that may also help with the data.

So, in my mind—and I could be totally way off—this question is being asked because what we're trying to show is that entities that may be relying on a certain string and getting a negative response now, all of a sudden, get a non-negative response if that string has been delegated.

So, maybe I'm wrong, but it might, for our own benefit, to may sure we're all on the same page, to try to figure out the why it's being asked. And then, it might become easier.

So, if what I said is right then I understand why the Site Finder note is there and some of the other stuff. But it would help to get us all on the same page on the role or the why of all of these questions.



---

MATT THOMAS:

I think that's an excellent point you bring up there, Jeff. I think understanding the actual details of why the question is important in the relevant context of a name collision is very important for us to be able to effectively answer or address all of these questions.

I'll just personally state my interpretation of this question is that it is focused on inherent indicators, that then are receiving a negative response, then use that signal to alter application behavior. So, if something was requiring on a negative or depending on a negative response, but now it is suddenly receiving something different, how does that alter the underlying behavior of that system or application?

I think that an example of this you could use is the [Joe] canary domain. So, currently, Firefox is using a canary domain in which, if it receives a positive response, it enables DoH. But for instance, if you want to turn off DoH then a system administrator would have to delegate that domain in such a way that it would cause a non-existent domain query.

And so, that altering of the signal or response then influences the applicational behavior. I think that is where they were trying to get at with this question. Just my take. Any other thoughts from anyone? Yeah, Jim.

JAMES GALVIN:

Yeah. Thanks, Matt. I'm trying to type some words in, here. I think that other folks should feel free to type some words. I think that there are, certainly, a number of real DNS experts who are out there in part of this group, and are there, and folks should add some thoughts, here.

---

Why are we asking this question? I'm trying to capture some notes and make some words of my interpretation of what I'm hearing here, but folks should feel free to, please, type in real-time, here. Help us out. Just put some words down so that we can capture the point, here. That's it. Thanks.

MATT THOMAS:

Thanks for bringing up that question, Jeff. I do appreciate that. To Jim's earlier point, I know there wasn't really any notice for this new plan of how we were going to try and go forward and approach these questions, and basically have a little bit of homework for all of us every single week so we can come prepared. Maybe we won't dwell on this question too much longer and we can come back to it. Yes, Jeff. Sorry. Go ahead.

JEFFREY NEUMAN:

So, one thing that would be interesting is—I don't know if this is the right way to say it, and hopefully everyone understands because, again, I'm much more of a layperson than you all—how has behavior evolved or differed since it was prior to 2012?

So, prior to 2012, lots of people may have designed systems that were based on getting a negative response from strings not being delegated. But since the new round introduced 2,000 or so of these TLDs, has behavior evolved such that it is less likely for people to be relying on that negative response? Does that make sense?

---

MATT THOMAS: That does make sense. So, I think you're talking more about ... I guess, in my mind, I'm thinking about the bifurcation of the DNS protocol from the stub resolver up into the application. And again, DoH is an excellent example of this where you have, potentially, not just relying on standard, low-level [gate host] by name or other functions that are built-in, but you're actually taking that and applying more in an application context. I think that's a great thing for us to have a better understanding of how the use of the DNS since the last round has evolved. Did I capture that correctly?

JAMES GALVIN: Yes. Again, you said it in a much more sophisticated way than I did, but yes, you captured that.

MATT THOMAS: Good.

BARRY LEIBA: Oh, hi. We've got to be careful with thinking in that direction, I think. It's clear that behavior of new software has changed, and anybody implementing something like DoH or DoT now is not going to be relying on NXDOMAINs for things that it doesn't know about, like we used to.

But there is a lot of old software out there and there is a lot of stuff that was expecting things like .mail and .home to behave a certain way. And a lot of that software hasn't changed in years and won't change for years, so we need to be careful and maybe split things out into, "This is what

---

new software is doing or is likely to do, and this is what we have to worry about for old software that hasn't been changed for a long time.

MATT THOMAS: Excellent point. Yeah. Definitely going to have legacy systems still out there with that kind of logic in there. Jeff, did you want to respond?

JEFFREY NEUMAN: Yeah, thanks. I think that that is one of the keys. We would love to see if the data can tease that out. In other words, which are the specific strings, if possible, that some of the old software relied on so that you could tease out the .corp, .home, and .mail prior to the next round so that we don't have applications for those.

I mean, that would be great, but if we find that some new software in 2016 now relies on strings other than the 2000, I don't necessarily feel as bad for them, because then they've taken advantage of something that they should have known a problem of.

So, I think using the data to figure out the other strings that would fall into that category for the old software is important, and ultimately to find this risk profile for strings being delegated in the future.

MATT THOMAS: Sounds great. You brought up a point that I wanted to just, maybe, pull that string a little bit more. It made me kind of remember I feel like there have been various application systems since the last round that have purposely elected to use a TLD even though they then planned on actually

---

going through the ICANN process. And so, I guess maybe some understanding of how some of that software would ever impact this question if that TLD did become delegated would be somewhat tangential, right? Steve Crocker, is your hand up?

STEVE CROCKER:

Yeah. Thank you. I want to echo Jeff's question. I hadn't quite got the gist of it until he spoke. If there had been appearances of new uses that depend upon NXDOMAIN from the root that weren't there before. Those are people who should have, I would agree, not done that.

I mean, you could say .corp, .home, and .mail arose at a time when there wasn't a clear understanding of the consequences, but if there is stuff that has arisen in the last couple of years that's worth flagging big time, I would say.

And so, I want to emphasize the value of the question that Jeff asked and, if there is a way to pull that data out, that would be worthwhile and might raise the question of, how come people are still doing that after the message was clear? Or maybe the message wasn't clear, so that not only do we want the data but we also want to try to understand, in classic words, where was the failure to communicate?

MATT THOMAS:

Well said, Thank you. Warren, I see your hand up, as well.

---

WARREN KUMARI:

Yeah. Two things. One, the message was definitely not clear to a lot of people because they would have had no way of hearing it. ICANN, in many ways, is sort of a bubble and we often seem to believe that ICANN represents everybody and everybody gets to hear the message. There are a lot of people who have never, ever heard this.

And secondly, to that, there are a lot of people who have a need for a domain, something that looks like a TLD, and don't have a way to get it. And so, they don't really have a choice other than to just make one up. So, in the last while, there has been .localsite, .openstack, .gnunet, .bitcoin, .coin, and a whole bunch of others which have started being used.

Some of these are intentionally being used in a way that shows up as a TLD and with knowledge that it would conflict with other things, the .gnunet people being one of them.

But for many of the others, they were people who needed a way to name something that wasn't going to be connected to the Internet and really didn't have any other sort of choice or a solution. Yeah, .eth actually is a good one. Somewhere, I had started making a list and I had 15 or 20 names, at least, that were showing up.

MATT THOMAS:

That's great, Warren. Maybe you could share that list with us some time that we could possibly look at that?

---

WARREN KUMARI: Yeah. I figured you were going to say that I need to figure out where I put it.

MATT THOMAS: Sounds good. Jeff, I see your hand u you.

JEFFREY NEUMAN: Yeah. So, I appreciate that there may be a couple of people out there that might not know about this, but making that assumption for all of them I don't think is a safe one. I remember specifically talking to the people that control the .eth. They call themselves a "domain registry" and they were talking to registrars, including my own, to try to carry their names.

And when I asked them the specific question of, "Well, why did you guys do this knowing that if ICANN has another round they can delegate these strings?" and their answer was, "Well, you know what? That'll be a decade from now and we'll worry about it then."

So, there may be a .local site thing out there that someone might not have known but I think there are more of them out there that know exactly what they're doing because they have a fundamental understanding of how DNS works. Thanks.

MATT THOMAS: Yes, Patrick. Is your hand up?

---

PATRIK FÄLTSTRÖM: Thank you very much. I think we had this discussion for the last 20 years or so, and I think we have exactly the same answers, that as long as the protocols work the way they do, as long as people, just like Warren was saying, are forced to just use some kind of name for certificates, for things that are partially connected to the Internet, to be able to build PKIs and XY9 when you do not have any connection to the global root, when you need to have your own CAs, as long as we have search paths in DNS resolution and as long as we have all of those protocol things, people will pick names that are not TLDs. That's just a fact. So, that's why I think this question is kind of weird, because as long as people are forced to use names that are not in DNS, people will do so. Thank you.

MATT THOMAS: Thank you for that. Steve Crocker, I see your hand is up.

STEVE CROCKER: Yeah. I thought that I was going to be very quiet during this call but now I'm pretty exercised. So, let me suggest two things, one of which might be a bit controversial. First of all, in a constructive and cooperative fashion, for those people who think that they need their own extension, I'll call it, top-level domain-looking thing, what is the best advice that we or anybody could give them so that they would behave better than simply making up a name and then having it trickle into the root? So, that's comment one.

Comment two is more aggressive. What is the most negative thing that could be done to discourage this kind of behavior? Today, when



---

somebody looks up a string that is not in the root, they get back NXDOMAIN.

I could imagine instituting a policy that says, "Well, when the number of requests for a non-existent domain rises above some threshold, we switch modes and start giving out different kinds of responses in a random or provocative way in order to defeat the stability that they're expecting when they build this in and then get back NXDOMAIN.

So, they could get an NXDOMAIN at one moment, and they could get no response at another moment, and they could get a redirect and an answer that would take them somewhere in another moment, all in a sort of controlled way, deliberately aimed at being as obnoxious as possible to discourage that kind of system development.

MATT THOMAS:

Thank you, Steve. Jim, I see you have your hand up. Please, go ahead.

JAMES GALVIN:

Yeah. Thanks, Matt. This question always comes to my mind when we have this discussion, taking on board the comment that people will always do what people feel like they want to do. They have a problem they need to solve and they're just going to solve it in whatever is the easiest and quickest way they can.

So, I thought it was interesting just now when Steve phrased the two ends of the spectrum. What advice can you give people so that maybe they'll behave better? Or, what's the most negative thing you can do to people who just don't do the right thing?

---

I think that, for this group, one of the questions that I would think we should ask of ourselves to help influence what we think we want to write is, is this a problem that we need to solve? And what I mean by that is, people who only read half the standards, so they know enough to create a TLD for themselves so that they can go do what they want to do, are they even going to get any message that we might give to them? Is there any way that we could create advice that's at all going to be useful to them?

I mean, the IETF could write a document and, under the presumption that they've gone far enough into the IETF to read half the standards to know that they could go create a TLD to do something, you have to ask yourself, why didn't they get the rest of them and get the rest of the message so that they knew that they shouldn't do the solution that they want?

I mean, I'm just sitting here thinking that if people are going to do the wrong thing, you can't fix people who just won't take the time to really look and see what the right thing is if they're just going to do something.

I always struggle with that particular issue when we get into this discussion about, "Oh, we have to provide a solution to people who have got to have a place to do things in case they want to do it, if they need a TLD extension." I'm just not compelled by that response.

I'm really behind the idea of ... I can't solve stupid people. You just can't. One of the benefits and features of the Internet is the fact that people can innovate new, interesting things. Well, let them go do that. If they're not going to play by the rules because they don't bother to look them up,

---

how are we going to fix that? Anyway, I'm sorry. I'm probably going on a little bit. I'm done.

MATT THOMAS: Thank you, Jim. Warren?

WARREN KUMARI: So, it's not an issue that people are being stupid, it's that in many cases people didn't have a good option. So, an example is when you very first install a Linux machine and it's not connected to the Internet it has no idea what domain name it should be. There is no good choice for it. And so, what a lot of distributions have been doing is localhost.localdomain or some options which is similar to that.

What I've been trying in the IETF is if we set aside a name for this particular sort of use, like .internal is the name we've been using there, then people who want a namespace which is disconnected from the Internet have one that they can use. Currently, there isn't a name for people to do this sort of thing.

And so, our advice has been, "You should do something else. We're not going to tell you what, but what you're doing is bad and you suck." So, something which, as I say, I have been pushing in the IETF and, I think I can mention, had been discussing in SSAC, as well, is create a name, set aside a label for this, or possibly, for internationalization, a couple of labels, and say, "If you want a disconnected name, here is the disconnected namespace. You should just go use that."

And obviously, we won't get everyone. There are certain sets of people, like the .gnunet people, the .ethereum people, some of the other .blockchain people, .okturtles, who specifically are squatting on names as a political statement against ICANN. We're not going to be able to solve those.

But people who want a name which is not connected to the Internet need a set of names that they can use. I guess I'll answer the question Jim's asking in the chat.

Distributions could set aside some part of their namespace, and when you first install a machine you could have all machines be called "localhost.randominternalId.redhat.com," but that then leaks a lot of information to Red Hat, which is a privacy leak, Red Hat probably also doesn't want that sort of information and, in many cases, the distribution isn't really that much managed.

The distribution is a base-Linux system with some amount of stuff scattered on top. There is also the risk of there are probably 300 distributions which have been created and have then gone out of business.

There was the one whose name I've forgotten that had a wizard's hat as their logo. They no longer really exist. And so, using a public name, which then ends up being released into the namespace and registered, will leak a humongous amount of potentially private information to whoever registers that name. Hopefully, that was clear and coherent.

---

MATT THOMAS: Thank you very much, Warren. Quite a few hands up, still. Steve Crocker, please. Steve, you might need to unmute yourself.

STEVE CROCKER: Yeah, sorry about that. So, I want to address Jim's point, which is why is this our problem? I'm empathetic with that kind of question. Excuse me. If you take it one step further, the question that is directly our problem is, in the next round, if there is a string which has been used this way—I'll call it contaminated—and then somebody requests it, somebody asks that it be delegated, should that be delegated to them or should it be ruled, like .corp, .home, and .mail, too contaminated to delegate? So, there is a kind of natural protecting-the-truth, in a way, role that we have.

So, as I said before, and I strongly support all of the discussion that has just taken place, any advice that we can give the world in general about how best and better to proceed, that's great.

But the other side of that is, what role should we understand is required to protect the existing namespace? I think if we say, "Well, it's not our problem," then it just circles back to, "Okay then," when we have to delegate it and somebody requests it, "what do we do?" Our successors, maybe even our children, will be having the same discussions in ICANN275.

MATT THOMAS: Thank you, Steve. Barry?

---

BARRY LEIBA:

Also responding to Jim, more in the, “Why are they doing it?” They’ve read some of the specs, they haven’t read all of them. Why are they doing it?” The answer is simple, that it’s because it works. They can’t figure out how to do it. They can’t figure out all the specs they need to read.

Microsoft has had this huge problem since, basically, the beginning of Windows, that it publishes ways to do things, it publishes APIs and various interfaces like that, and they’re always having trouble with people figuring out undocumented interfaces that simply worked, and so they used them.

And then, they’re stuck supporting them for a long time because they don’t want to lose customers that are doing that. And one way they could get around this would be, as Steve was suggesting, make it seriously not work. Stop supporting them.

In their case, they have a customer base that makes that impractical. In our case, we have a community that we need to support. I’m not sure whether what I quipped in the chat as “weaponized DNS” is or isn’t the right way to do it, but the fact is that people will continue to figure out that something works even though it’s not the way you’re supposed to do it, and they’ll get stuck on that because it was easy.

MATT THOMAS:

Thank you for that. Patrick.

PHILIPPE FOUQUART:

Thank you very much. This is exactly why I pointed out things that people do just because they have read the specifications. Just like Warren said,

---

people have read the specifications and they do not see any safe TLD that they can take. Everything is taken. Everything is allocated. It's also the case that, as I said, as long as the IETF do encourage and do have in their specification they'll be able to use SearchPath, people will use that.

So, if it is the case that we should tell people something to do, I think the people that would like to have a new TLD should ensure and go to the IETF and change the protocol so that people that do read the protocols can still not create any problems. Thank you.

MATT THOMAS:

Thank you, Patrick. Warren, back to you.

WARREN KUMARI:

Thank you. So, just responding to something that Steve had suggested earlier, and a few other people have touched on. If you try and make it so that you intentionally break these by returning something other than NXDOMAIN, or something similar to that, I think that you're going to end up in a much, much, much worse situation.

I think that this really does end up becoming a mutually-assured-destruction-type event. There are really quite a lot of people—I'll go back to GNU but there are the Turtles folk, there are a bunch of others—who are not happy with how the current set of new gTLDs happened and feel that they were disenfranchised.

And doing that seems like it would just be thumbing your nose at them and you will, I think, end up with them being intentionally going out of their way to break this sort of behavior.

---

I think that there will also be people who will then say, "I've been using this name. It's no longer working for me because ICANN is returning intentionally breaking answers.

And so, what I will do is I will just override that TLD in my name server and in all of my client apps, and now you have this problem but it's spread out over a much wider range and launching any new gTLDs will become way, way, way harder.

I think you'll end up, more likely, with things like whitelists of names which will get queried to the root, and everything else will get dropped. Or I would just take this set of names, I will bake it into my app, and I will never query the DNS. And now, what you've got is organization of the namespace. So, that's a really dangerous thing to try poking at, I believe.

MATT THOMAS:

Thank you for that, Warren. I see we're getting ... Steve, we'll go to you, next, and then, as we're getting closer to the top of the hour, I do think Jim had a few other announcements and business-related things he wanted to probably announce around Kuala Lumpur and travel. So, this has been an excellent conversation and I'm super-happy it went so well. Steve. Steve, we can't hear you. You might need to unmute.

STEVE CROCKER:

Well, yes. Sorry about that. Apologies for taking so much time and making it worse by saying half of it muted. Warren makes a very good point that the direction that I was suggesting could be very problematic. But I was particularly taken with Warren's point that there is a set of people who



---

feel that they were disenfranchised by the way that ICANN handles the gTLD rounds.

Let me speak forcefully as a past chair of ICANN Board, and on the ICANN Board during the entire new gTLD rounds from 2008 or before. I think it's a very important point that there may be other important ways to think about what should be in the root or how top-level domains should be used other than simply as commercial enterprises that are subject to the contractual rules under the particular contracting process that ICANN has.

And I think it would be very helpful to get a clearer understanding of who those parties are that feel disenfranchised and to understand from an objective and empathetic way. I don't know how supportive, necessarily, but at least understanding their point of view, exactly what those issues are, and not be dismissive of them. So I, at least, would like very much to have a better understanding of that and for their voices to be given proper hearing. With that, I will subside.

MATT THOMAS:

Thank you, Steve. I think we have time for one last one. Jeff?

JEFFREY NEUMAN:

Yeah, thanks. Look, the only way to effectually change is to have them be part of the system which they don't want to be a part of. And if we incentivize groups that have intentionally become disenfranchised and have intentionally warped around their own rules, then we're just going to encourage others in the future to break the rules.

---

I mean, we should not be rewarding that behavior at all. I see Warren's note, "They tried to participate." Look, a bunch of us tried to get strings in 2000. In fact, 1997, we tried to get strings. We couldn't. We tried to participate. We tried to get it. And you know what? We came back in 2000. Didn't get it. We came back in 2005. Didn't get it. Came back in 2012. Right?

You have to work within the system you have and not just cater to those that become disenfranchised. Look. I disagree that we'll end up with mutually assured destruction but we have to try to have some sort of order to the system, otherwise what's to stop me from now becoming disenfranchised with this whole thing and setting up my own space?

MATT THOMAS:

Thank you. Warren, we'll close with you, and then, Jim, we'll switch it over to you. Again, this has been an excellent conversation. Great thoughts and commentary going back and forth. Thank you so much for all of it. Warren?

WARREN KUMARI:

Thank you. So, responding to Jeff, yes, you tried multiple times to participate and, finally, people did manage to get names through. But I should point out that you, the royal "you," in this case, is largely commercial enterprises. In order to get a TLD, you needed \$185,000 for the submission fee, plus an indeterminate other large amount, probably on the order of \$500,000 if you want a single TLD.

For many of these organizations, they are non-commercial groups. GNU, as the GNU Software Foundation, they don't have that sort of money. Tor, the privacy browser thing, doesn't have that sort of money. And also, in many cases, even if they did have that sort of money, they're not actually registering names within the TLDs. It would be any name that is set aside.

I actually tried helping some of these folk. There was no way in the Applicant Guidebook to be able to do something where you have a name that is not used within the DNS. You can have a TLD but there are a whole bunch of regulations and requirements you had to follow.

We had tried to have some set of discussions, I can't remember with whom, on, "Can one have a TLD without registrants?" And we were told, flat out, "Nope, it can't be done." "Okay, can we have something where we register the name so they don't get delegated?" "Nope, can't be done."

And so, that's the sort of disenfranchisement. It's not they wanted their name, and couldn't get it in a short amount of time, and went home. They wanted their name. They tried to follow the process. They said, "Is there a way that we could possibly get this done as a discount for community good?" They were told no. They said, "Okay. If we had the money, could we do these other things?" They were told no. And so, they largely said, "Well, guess I can't work within your system," and left.

MATT THOMAS:

Thank you, Warren. Jim, I will kick it over to you for closing out.

---

JAMES GALVIN:

Thanks, Matt. Yeah, let me repeat what you've said a couple of times already. Thanks, everyone, really. Just an excellent meeting. Very happy for all this discussion. I hope we continue this as we move forward in our weeks going forward, as we begin to collect notes and streams of consciousness.

I'm going to stay away from having the last word in this discussion about putting names in. I have a thought going, there, but I think I'll just save it for next time because I think it's a provocative comment to make.

Let me make the two announcements that I think are important, here, I want to remind people of. By now, folks have certainly seen—I think we said this last week, I'm going to say it again this week just to remind folks—we know that ICANN68 is going to be a virtual meeting. It was previously supposed to be Kuala Lumpur in Malaysia.

Plans are being made for what that agenda is really going to look like and how that meeting is going to be organized, so that's really outside of our control, and that's fine. What's important for us here is we, as a working group within ICANN, have an opportunity by default for an all-day meeting in front of every ICANN meeting so that we can have facetime, face-to-face meetings, and a very highly productive meeting, and move forward.

And we would have had one scheduled in Kuala Lumpur on the day before the actual ICANN meeting starts, but rather than us planning for a six-hour or eight-hour video teleconference or whatever, we're going to choose to not have our long, all-day meeting, and instead we'll just

continue our cadence here, now, with just our weekly meetings, our one-hour teleconference, and we'll just continue to move things along, for now.

One should expect that, if the fall meeting does continue and is held in Hamburg, Germany, where it is currently targeted to be, we will get our full-day meeting in front of that. And folks should plan and take note for their travel plans accordingly. So, that's one announcement.

And then, on a related note, speaking of travel plans, I do want to point out to people and remind people ... We haven't done much with this up to this point but I think, as we now get into a regular cadence of meetings and start having a lot of discussions, and if we actually have a very well-orchestrated full-day meeting in front of an ICANN meeting, part of this project is to have some limited travel support available for a small set of people. In fact, we have five slots that have been allocated to this project that could be used if we need them.

There are two things to note about that. One is that it's not intended to be somebody's primary support, which means the objective here is that if you've got other sources of funding to go to the meetings then you should use that first, especially if you have competing ICANN support options. There are a number of people who have other ways within the ICANN system to get funding for the meeting, and it's important that this be the second choice, not the first choice, for where you get your funding.

The second thing is that there are only five slots available. We will, at some point here, have to get much better about this process. Kuala

---

Lumpur kind of came upon us too quickly and, well, they just canceled it so we got the benefit, there.

We'll do a better job as it comes time for Hamburg because we'll make an announcement to the list, to people: "Okay, if you're interested, you should put your hand up and say that you want it."

The admin committee will be making a choice as to who gets funding if we get more people than we have slots to give. We will be selecting people, and the primary criteria that we will be using for selecting people is who participates in discussions.

I want to put that out there now so that people know upfront and you have plenty of time to work your way into being prepared and contributing to discussions. Hopefully, we're just going to have to deal with any issues that we have as it comes along, but that's the model.

It'll be transparent. People will know who is volunteering and asking, and then we'll let you know who was selected. Hopefully, it'll be obvious to people why folks were selected because we'll be looking for participation as the criteria.

So, I think that was it as far as that was concerned, just those two announcements I wanted to make to people. We're very much at the top of the hour. Any other business from anyone about anything? If not, thanks very much. Hopefully, folks can see how we're going to try to do things going forward and where we are. We'll continue onto question two. We'll move onto question three next week if we need something else to talk about. We'll see you next week. We're adjourned.

[KATHY SCHNITT:] Thanks, everyone.

[END OF TRANSCRIPTION]