ICANN79 | CF – Joint Session: ALAC and SSAC
Sunday, March 3, 2024 – 1:15 to 2:30 SJU

MICHELLE DESMYTER: Hello, and welcome to the joint session for the ALAC and SSAC. My name is Michelle DeSmyter, and I am a participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior. During the session, questions or comments will only be read aloud if submitted within the Q&A pod. If you would like to speak during this session, please raise your hand in Zoom. When called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use a physical microphone to speak.

Interpretation for this session will include English, French, and Spanish. Please state your name for the record and the language you will speak if speaking a language other than English. Please speak at a reasonable pace. And with that I will hand the floor over to Jonathan Zuck, ALAC chair, and Ram Mohan, SSAC chair.

JONATHAN ZUCK: Hello, hello. Jonathan Zuck here, chair of the ALAC. And welcome to the SSAC. You are by far our favorite SO and AC, and no, we don't say that to anybody else. So, we're always looking forward to these discussions with you, learning about the stuff that you're working on, and more recently, thinking of ways that we can be a part of more visibility for some of the stuff that you're working on. So, we've got that

on our agenda as well. So, welcome, and looking forward to a productive meeting.

RAM MOHAN: Thanks, Jonathan. I'm Ram Mohan, new chair, if you will, of the SSAC. I just wanted to take a moment to just go around the room and ask SSAC members to just identify themselves and what their roles are in the SSAC and the day jobs, if you don't mind. I'm Ram. I'm the chair, and day job, I work at a domain name registry called Identity Digital. Jaap, shall we start with you?

JAAP AKKERHUIS: Okay. So, I'm Jaap Akkerhuis, and I'm one of the eldest persons in SSAC. And my day job is still in NLnet Labs.

MERIKE KAEO: My name is Merike Kaeo. I'm a member of SSAC. And as a day job, I do all things security sometimes for corporations and right now just as an independent contractor.

GREG AARON: Greg Aaron, SSAC member.

SHAH ZAHIDUR RAHMAN: This is Shah, an ALAC member.

JEFF BEDSER: Jeff Bedser. I'm a part of the leadership team with the with SSAC, and I'm with CleanDNS, a DNS abuse mitigation company.

PETER THOMASSEN: I'm Peter Thomassen from the SSAC. I cochair the work party on DS record provisioning automation together with Steve Crocker. I'm also a member of the membership committee, and I work for SSE, which is a German IT security consultancy, and deSEC, which is a DNS hosting platform.

WARREN KUMARI: I'm Warren Kamari. Day job is Google, and I'm part of SSAC.

BARRY LEIBA: I'm Barry Lieba. I'm part of the new SSAC leadership committee, and my day job involves mostly working on Internet standards.

STEVE CROCKER: Steve Crocker, SSAC member.

JONATHAN ZUCK: I'm Jonathan Zuck. I'm chair of the ALAC, and my day job is a filmmaker.

MATTHIAS HUDOBNIK: Hi. My name is Matthias Hudobnik. I'm the new ALAC Liaison to the SSAC, SSAC member, of course. And my day job, I'm working at Europol.

TARA WHALEN: Hello. I'm Tara Whalen. I'm the SSAC vice-chair, and I work for Cloudflare.

ROD RASMUSSEN: I've been told to come up and say my name. Rod Rasmussen, former SSAC chair, now popcorn eater.

GAUTAM AKIWATE: Hi. My name is Gautam Akiwate. I'm a member of SSAC. I'm a researcher at Stanford, and I'm a co-chair of the registered name management work party.

JONATHAN ZUCK: Russ, do you want to introduce yourself?

RUSS MUNDY: Sure. Thanks. I'm Russ Mundy, SSAC member, also SSAC liaison to the RSSAC. Day job is TIS Labs research organization.

RAM MOHAN: And Benedict, do you want to say hello?

BENEDICT ADDIS: Hi there. Yes. Benedict here. I'm just next door, and I'm an SSAC member and run the registrar of last resort.

RAM MOHAN: I think that covers everyone. Anyone missing? Andre? Andrei. Hi.

ANDREI KOLESNIKOV: Hi. I'm Andrei Kolesnikov, SSAC member, ex liaison for ALAC. I'm doing a lot of music these days, but also, I'm a IOT guy.

RAM MOHAN: Right. So, Jonathan, one of the things that happened and just a little bit of backstory for everybody here in the room. In January, ICANN organized a meeting of the SO/AC chairs and vice-chairs, and it was my first introduction to that forum. So, I was there, and in one of the lunch breaks, Jonathan and some of his vice-chair the vice-chairs of the ALAC were there, and we were chatting about things that-- two things. One, that there is such a great deal of synergy between the work that the SSAC does and the reports that we produce, especially those that have relevance towards end users, registrants, things like that and the ALAC itself, At-Large, and its reach and its members.

And one of the questions that Jonathan had posed was, are there things that we can do together that accelerate or amplify the work that the SSAC has done. So, we went and we thought about it some more. And we came up with this idea that perhaps here are things that we can do that leverage work that the SSAC has already done. From the SSAC's perspective, it's a bit of a light touch. But that we could potentially provide as material as things that to the ALAC, to the At-Large, that

allow you to then take that and present it and amplify it the places where you want to go.

I mean, from its original conception to where we are right now, what's clear is that the requirements for what this is and how we want to do this is going to go through an iterative process. We're going to work together to figure out what can actually work well. But I think there is no doubt that when it comes to providing material, providing the foundation for content that allows you to go and perform your purpose and inform your audiences we're very happy to collaborate and to be part of that solution. So, that kind of is a preamble for all of that.

The underlying piece that we were speaking about in January is that across the board, I mean, this is everybody wants to talk about security, everybody wants to talk about cyber security and things like that. But inside here, there is a more specific, much narrower remit, and a much narrower focus. There are some things that we care about and that we actually know about. Right? There are many things that we care about, but that we don't know about nearly as much. And the way SSAC generally works is to try and engage with you or try and provide counsel to the communities that we serve on the things that we know about rather than speaking generally motherhood and apple pie statements. Right?

So, our aim and focus is, therefore, potentially, to take some of the information, the materials, the recommendations, things like that that we've got, and hand them to you or collaborate, work alongside with you, and see how it may help you in dissemination. Inside of the SSAC, this topic, we talked about it with the SSAC membership yesterday.

This topic generated good back and forth and responses from them. You can go to the next slide, after this. Yeah. So, there are requirements. There are more things for us to work, both inside of the SSAC as well as alongside with you. But at a high level, as a general-purpose statement, we're together with you in helping solve the issue of how do you provide material, how do you provide information that can help keep the DNS secure and stable. Are there other SSAC members who would like to speak to this topic as well? I'll turn this to you then.

JONATHAN ZUCK: All right. Thanks, Ram. Yeah. We had a really productive conversation in DC in January. What I shared was that we are trying to really create a free flowing sort of engagement engine, if you will, in the At-Large community because we have a deeper bench, if you will, than most people in the ICANN community with these incredible regional organizations that in turn have both individual and institutional members that are referred to as ALSs, who in turn often have their own members. Right? And so, there's this incredible virtuous tree there that has a lot of potential, and we want to look at how to best use that potential, how to decide what we should use it for.

So, we talked a little bit. I saw on a previous slide that maybe it's for feedback on a particular issue to get broader feedback on end user positioning. Maybe it's for message amplification if we're trying to get a message out to a broader group of people. And then finally, it could be just on straight education. So, DNS abuse or cybersecurity is one of those areas where we both, the

ALAC and the SSAC, are trying to nudge ICANN Org, the contracted parties, and others to do what they can to prevent malicious registrations upfront. But what can we be doing to help promote digital hygiene at the point of the end user? So, that's one possible area where we might try to do some end user education. We've decided to do some end user education.

And so, maybe the idea is to work together with you as we develop the materials for this course to say, oh, well, you should include this too, or if we describe it this way, then end users will be able to understand it to put that together. And then the idea would be to try to execute something similar to UA Day, right, where it's sort of an event in a box that goes all the way out to our ALSs that can then put on an event.

And so, when we talk about a primary or secondary audience, our group, if you will, are largely IT savvy or smart end users, but their audience could vary. It could be just trying to bring in end users into their seminars, or it could be to mobilize to send a message to another type of organization. Right? And that's why we talked about something like DNSSEC. Maybe it's not an end user message other than to tell-- It's sort of like how the pharmaceuticals say, ask your doctor about clozapine or whatever. Right? And you've got people asking their doctor about drugs they've never heard of. Right? And is there is there a way to say hey, end users, ask your bosses about whether or not they're DNSSEC ready or something? Just so that it becomes something that they're hearing from more places.

So, those are, in very broad strokes, what we were trying to bring to you as an idea. And so, ideally, if we do this right, we'll have a process in

place to evaluate requests for campaigns and then a process for then executing them as well, and we're working on all of that. But I think we will do in the first something related to safer cyber or phishing or something. It all began with phishing Friday, and then Heidi reminded me that that might confuse the Catholics. So, it's good to have people thinking these things through. I was so excited, and then she popped my balloon. But something like that that's catchy and that we really package it in such a way that it's very easy to deliver and very easy to promulgate out to a broader end user community.

So, that's what we're working on, and we'd like to work with you when you identify messaging that you think should either get to end users or be delivered by them, either of those things, then we'd like to hear from you and ingest that idea into the system we're building to decide whether or not we would launch a campaign on that topic. That's the idea.

As follow up, we're working on materials. We'll come back to you with them on the fishing in particular and then maybe get your input or maybe explain this differently or maybe you're missing this point or something like that and then we'll go try and roll that out as an experiment. But the other action item is just to always be aware of the possibility of something like this so that you might bring it to us.

RAM MOHAN: Thanks, Jonathan. That's super helpful. We're happy to help, and we have material. And in some cases, you may not have material, but we have access to experts who can provide guidance on these things, it sounds better this way than this other way. Right? So, even that kind

of guidance can be provided. So happy to help and I look forward to collaborating and doing something together on this.

JONATHAN ZUCK: All right. And maybe you just even know of a resource. I'm sure that there's million available videos of Steve Crocker explaining to kindergarteners how DNSSEC works or something on a whiteboard. And that video, if we could get our hands on it, might be part of an outreach effort or something like that.

STEVE CROCKER: The video you really want is a kindergartner that's explaining to me how DNS works.

JONATHAN ZUCK: Yeah. There you go. Exactly. That's how you know you've learned something is when you're able to teach it. Right?

RAM MOHAN: Okay. Great. Thank you. I think that covers this topic. Let's go to the next one, Urgent Requests for Disclosure of Registration Data. Steve, you're on.

STEVE CROCKER: Thank you. We have some slides perhaps. Good. So, the concept of urgent requests showed up in the expedited—I have trouble saying that word in the context of the policy development process—Phase 1

requirements, in that there was a little detail that was left for implementation phase several years later of how many days the registrar would have to respond to an urgent request.

So, the sequence of events is that the Public Safety Working Group operating within the GAC requested in the phase 1 the inclusion of treatment of urgent requests. The spec, as I said, the resulting spec left a detail unspecified and pushed it to the Implementation Review Team. I was not part of the initial run up, but I got involved in the Implementation Review Team. And there was a heavy debate about whether it should be one day, two days, business days, calendar days, etcetera. And I looked at this closely, and I said, "Wait a minute. What are we really talking about here? " So, there was a closer examination that revealed deeper problems. The net of it all was that the whole idea was scrapped or at least scrubbed for the moment. And in addition, I brought all this back to SSAC, and we wrote a report laying out issues. So, that's SAC122.

So, here's the language of what an urgent request is. This is verbatim only, but I've changed the formatting to make it just a little easier to highlight things, but the words are, I believe, 100% equal. So urgent requests for a logical disclosure apply to imminent threat of life, serious bodily injury, critical infrastructure, or child exploitation. These are very, very heavy, big issues. And imminent threat suggests that urgency should be measured in like right now, not next week or even tomorrow, but how about drop everything and attend to this right minute?

So, response times, if you look at this from a systems point of view, you really want measured in minutes ideally and worst case, hours. And

then there's a definition of what critical infrastructure means of physical or cyber systems that are vital in their incapacity or in that their incapacity or destruction would have a debilitating impact on economic security or public safety. To put this in perhaps a dramatic language, these are wake-up-the-president level issues.

So, what would an implementation of a system that accommodates, that provides for identifying, submitting, and responding to urgent requests look like? So, these are my deconstruction of this. So, I would say, first thing is, it's not going to be simple and easy to do this. It's going to take resources and a lot of structure. So, is this a real problem, and has it been documented, and what do we know about it? How often does this happen, and what are the consequences and so forth? So, some sort of study or white papers or a collection of statistics. So, I would just label that as justification of the requirement.

And then to implement it, you should have a well-defined submission process that actually works, and as I say, time frame for that such a thing measured in minutes. And then you don't want it abused, and you want it accountable, so you got to have some control over how the submission is done, who does it, and to be able to tune the process in case it gets overused or even underused. And then how long is the response time for the registrar's response and how does the work should be measured in minutes to hours. So, that's what I put on the expected side of this.

And then on the actual side, so what actually was proposed? Well, there wasn't any documentation, zero, of the requirement. There was no separate submission process, no SLA on submission time, no control

over submission, and the response time was a best effort with a compliance specification, which was what was actually being argued, that was measured in some number of days. So, the registrars were quite concerned that if they didn't respond, would they be in trouble from a contractual point of view? I can tell you that that's the last thing that anybody would be worried about. If there really was an urgent request and it was meaningful is arguing about, well, we wouldn't have to respond for two days. Meanwhile, very bad things have happened in the meantime. Short answer is stop, hold. This is just not fit for purpose. We should not be having such a discussion if we're going to have a definition of urgent request and an implementation of it. We've got to start over.

So, that leads to what seemed to me straightforward list of things to do if we want to proceed on this document, I mean, design, and implement a submission process that is consistent with the need. Set up some sort of accountability for submitters that controls abuse and so forth. And then on the response time, again, design and implement an automated response process. You're not going to get proper responses if you have to bring in people on a 7 by 24 basis. The costs are too high and the ability to provide the human support is not going to be there. So, that's the short story that I wanted to share. We documented the same things in SAC122. I'm happy to respond to questions or comments or whatever.

RAM MOHAN: Okay. Sebastien, let's go to you.

SEBASTIEN BACHOLLET: Thank you, Steve, for this presentation. And what next?

STEVE CROCKER: Here's the answer that you're going to love. Not our problem. Actually, I would say what next is two or three things. One is going back to the beginning, it would be, I think, very helpful to have documentation of statistics or exposition of what the issues are. One of the ideas that has come up in informal discussion is, well, how would you get the data about urgent requests? Urgent requests do get made in other circumstances and other settings. So, for example, law enforcement goes to Google and says this is a big problem with this website and something.

So, one thought is, well, why not ask the registrars how many urgent requests they actually have gotten by any form, never mind the ICANN process, and see if there's any data there. And the response is, oh, what a good idea. So, that's a thought. Another thought that is kind of in a whole different level is, how did we get to the state where we have this much energy put into arguing about to defining and requiring urgent requests and arguing about number of dates and so forth when somebody should have said, other than me, should have said, well, wait a minute. What are we really talking about? How did it get out of the GAC? How did it get through the staff that this is a non-working idea? And so, there's an opportunity, I use that word advisedly, for some reconsideration of what the processes were and who was watching over this and said, oh, yeah. Let's just pass this on. And there's more that one could ask, but those are a couple of ideas.

RAM MOHAN: Thank you. Hadia?

HADIA ELMINIAWI: Thank you. I had a question similar to Sebastien's and based on Steve's answer. So, what's next? More data is required. That's one thing. But then--

STEVE CROCKER: When it came up to the attention of the Board, there was some discussion that I was not privy to that took place among the relevant parties. And the net of it was that the notion of urgent request simply got removed. And the notion of expedited requests, which have a less specific definition, got inserted. All of this is a sort of way of papering over where things are and sort of removing the immediate embarrassment of having an inconsistent approach to urgent requests. So, it's back to the beginning in some sense.

Multiple possibilities. One is, okay. Forget about it. Another is, okay. Let's take it seriously and start from scratch and get some data about urgent requests and get a definition that is meaningful and has to be addressed. Those are two quite different kinds of ways of going forward, or you can view them as related to each other in that you kick the cane down the road until you decide that you actually have to put the time into it.

RAM MOHAN: Amrita?

AMRITA CHOUDHURY: Thank you, Ram. Amrita for the record. My question is these are real time issues of crisis. Many times, what we see is law enforcement directly gets in touch because if it's a matter of life and death, coming to ICANN, going who it is, rather there are bilateral ways in which data sharing is done or that's what we are seeing these days to reach to the particular person, the entity from whom they want the data to get it done. So, yes, it needs to be simplified, but is there a synergy between the two processes? If you're talking about real time issues, the back-and-forth time zones, etcetera, and everything.

STEVE CROCKER: So, there's a couple of interrelated questions there, which I'm happy to share. One is, okay, so when these things do come up, what happens? And the answer from law enforcement is we do everything, we explore all possible channels. we are not limited. We don't have to do it this way. We use everything, informal means, formal means, and so forth. Okay. Then another side of this is, in the discussions within the Implementation Review Team, and there's transcripts that you could dig and find the dialogue that I'm talking about, was one of the very earnest and conscientious registrar representatives says, "Well, they know how to reach us." And instantaneously, the law enforcement representative who was in the room says, "Well, no. We don't. " And the actual meaning of all this, yeah, we know how to reach some of you for some of the time, but there is no organized thing about this.

So, there was a lot of discussion about what would it take to have a formal access point. And part of that discussion says, well, you have

one for abuse. Can't you just add on that requirement? Oh, no. That's a whole different skill set and requires an attorney to get involved, and that would be expensive. Well, I don't want to put on the record what I said to that.

RAM MOHAN: Thanks, Steve. So, I have Claire, then, Alan, then I have you, then Hadia, and then Merike. Merike, do you have just a-- Oh, there's an online as well. Okay. Just a 2-finger intervention from Merike.

MERIKE KAEO: Yeah. Because one of the things that I did want to call out is that there's something called the MLAT, the multilateral agreement. And to actually get information about law enforcement between different geographic regions can take months. And so, that's been a problem for over 12 years.

RAM MOHAN: Thank you. Claire, let's go to you.

CLAIRE CRAIG: Hi. Claire Craig for the record, ALAC. I'm a little confused at this point. I'm hearing the conversation, and there is what is quoted by one party as an urgent request. It goes through a set of procedures. And the outcome is no, this is lacking too much information. We can't deal with it. So, it is not pursued. Right? Am I correct there?

STEVE CROCKER: Are you referring to the RDRS treatment of urgent requests?

CLAIRE CRAIG: Yes.

STEVE CROCKER: So, that's a subsidiary and tangential thing in that what I'm describing was part of the mainstream development of requirements, phase 1 requirements for the thing. And then as the RDRS came alive, somebody said, "Well, what about urgent requests?" So, they said, "Oh, let's tack that on." And so, they put a little check. And then when people tried to exercise that, they got the kind of responses that you're talking about.

The Board response on both of those basically was, oh, we aren't there yet, and removed urgent requests per se and changed it to expedited requests, which are-- And the big improvement is that there's no specification as to exactly what that means. And so, therefore, you can't say, well, it's not fit for purpose because we didn't say what the purpose is exactly.

RAM MOHAN: Thanks, Steve. Yeah. Sorry. Matthias had, again, a very brief intervention on just this topic, then I'll continue it on the line.

MATTHIAS AKKERHUIS: Thanks. Matthias speaking. Just maybe to clarify your question a bit. So, I mean, there are legal frameworks like, for example, the Budapest

Convention or also MLAT as you mentioned where law enforcement can more or less do their work, and there are also definitions based on this treaty or framework. And then the question is also in how far the national applicable law comes into, let's say, force. And that's why maybe, for law enforcement, they would more or less get some of the data if there's really some urgentness. And also, it depends then on the national procedural law, data protection law in how far and how fast and what the definitions are. But there are treaties like the Budapest Convention where you can check. There are quite clear definitions what is meant by it.

RAM MOHAN: Thanks, Matthias. I'm going to draw the line that whoever has already raised their hands right now. So, I'm going to go online to Olivier.

OLIVIER CREPIN-LEBLOND: Thanks very much, Ram. Olivier Crepin-Leblond speaking. I hope you can hear me because I'm in a terrible--

RAM MOHAN: We can.

OLIVIER CREPIN-LEBLOND: You can? Okay. Good. No. Just a very quick intervention on this. We've discussed this issue in the past already of [audio glitch 00:34:35] RDRS and through other means. But I think that because we live in this increasingly litigious world, well, very few organizations are going to

want to commit to these things in contract. If you can find a way to get these fast tracks put together without actually having contracts, that mean that you end up being responsible for someone's death or whatever, if something like that happens, then you're probably going to get some solution. But as long as you want to put it in contract, I really see many organizations pulling back and saying, well, you can just pick up the phone, and somebody will answer. But it's not going to be a contract as such.

RAM MOHAN: Thank you. Alan?

ALAN GREENBERG: Thank you. Steve mentioned one issue that was "not fit for purpose". There are lots of others, and I as one of the people who sat through the discussions about this particular issue for interminable hours and ending up with what the result is you saw, I'm more concerned with how we got here and how this is allowed to happen and how it happens regularly. I was encouraged a week ago when the Board held a session on what do we do when we approve something which proves to be unfit for purpose. So, the fact that it's being looked at is really encouraging. The fact that no one seems to be looking at how it happens repeatedly is the part that really concerns me. Thank you.

RAM MOHAN: Thank you. Bill?

BILL JOURIS: Okay. One thing Steve mentioned was the case where a law enforcement said we need to get ahold of somebody, and the contractor party said, well, you know how to get ahold of us. And the response was, no. We don't. How difficult would it be to simply have a tiny database at ICANN that says, here's who people should call at each registry. And then we could tell law enforcement, come to ICANN this way, and we can get you a response in seconds of who you've got call. It's not a whole solution, but it might be a worthwhile step that could be taken much more quickly than we can solve the whole problem.

STEVE CROCKER: It's a very reasonable suggestion that got an immediate but what I would call unreasonable response from the registrars.

RAM MOHAN: Thanks. Hadia.

HADIA ELMINIAWI: Thank you. This is Hadia. So, I was thinking if actually we could get off the ground system that can provide answers to requesters in minutes or seconds. And I assume that would be some kind of automation. Then it doesn't make sense to say that we will be running this system for urgent requests. And I see that there is a need for such a system more for old requests than actually for urgent requests. And I would say that urgent requests might have a way forward through other ways and paths, though difficult. But having a system that-- It doesn't make sense that today with all this technology around us, and we are talking about

requests being answered in terms of weeks. Even if it works today, five years from now, would it work? Thank you.

RAM MOHAN: Thanks. You're echoing in many ways things that you've said in our papers, so thank you for that. This has been a good, spirited discussion. Appreciate it. We have a couple more things to share with you and again, invite your input and your thoughts. So, let's go to the next one, which Warren is going to walk us through. And the title of that is implications for end users on the proposed top-level domain for private use.

WARREN KUMARI: Thank you. And hopefully, this one will be a lot less controversial and exciting. That is not me. Next slide, maybe. In the slot. Oh, which is 122? No. That's not me. There you go. A little bit of background as to what this is and how we got here. So, this was originally a concept which we had discussed in the IETF, and it was an Internet draft called the .internal TLD, which suggested that there should be a TLD reserved for private use.

The IETF DNS participants rightly observed that this was actually related to domain policy, and so the IETF is not the right place to discuss it. And they suggested it should be brought to ICANN. And so, we did that, and it got progressed and became SAC113, SSAC advisory on private use TLDs.

So, what does that document actually say? Well, basically, it suggests that we should set aside a TLD for private use. Basically, take a domain, reserve it for private use. And why? Well, the obvious answer is people keep doing this, and we can't stop them. So, instead of having everybody just choose their own name at random and we end up with this spread all over the name space, we should have ICANN choose one, set it aside, and we'll try and convince people to use that.

So, this particular screenshot comes from the ICANN DNS magnitude data, I think it's called. That is a set of data collected from the l root servers, and it looks for the largest number of TLD lookups. And these are sorted according to some fairly complex formula, but it lists a bunch of the undelegated TLDs that people are using. And obviously, you can't read this, but that particular page has, I think 1,000 names or so, the top 1,000 undelegated TLDs that people are using for this.

So, the SSAC113 document recommends that the ICANN Board ensure a string is made available for this. And the string should be reserved and never delegated, and then it went on and provided some criteria for the selection of the string. It needs to be a valid DNS string, it needs to not already exist in the root zone, it can't be confusingly similar to other TLDs, and it should be relatively short and memorable. But the SSAC felt that it was inappropriate for the SSAC to suggest that a particular string be chosen, and so it asked the ICANN Board to ensure that this happens.

And the ICANN Board delegated this to the IANA and said, dear IANA, you please choose the string and propose this. There was initially a public comment period on sort of how the process should happen, and

I think it was largely agreed IANA's a reasonable third party to do this. And recently, the IANA made the provisional determination that .internal is a reasonable string. They say it meets all of the criteria that SSAC113 laid out. And there's currently a public comment period open, it's linked on that top link, asking for feedback on whether the string .internal meets the criteria from SSAC113 and any other observations.

One thing that I will note regarding the public comment is some set of people have suggested .local might be a reasonable string. It's used in a bunch of places, like people have seen printer. local. That string actually does not match the criteria because it is already used for other purposes. It is reserved specifically for multicast DNS.

So next slide. This is where we get to sort of meat of this. So, the implication and benefits for end users. And I think the primary one is until such a string is set aside, people have just been choosing a string at random whenever they need one. That's how we ended up with things like .home, .Corp, .Mail, .server, .internet, .TLD, .internal, etcetera. So having one clear string set aside for this purpose, I think, will be helpful, especially for companies, home users who want to do their own private DNS thing.

It also reduces the risk of collisions when we have another round of new gTLDs. And sort of the only or the primary potential challenge is users are probably not going to be aware of the string whenever it is actually chosen and set aside because this is ICANN inside baseball. And so, people might not know about it and so might not use it, but I think that's sort of the primary concern.

And then down here, we have SSAC and ALAC's role. SSAC comes up with technical advice, and we think ALAC advocating for end user interests. Possibly once the string is chosen, whatever it ends up being, it would be helpful if ALAC could help share this information to end users. So, if they want a string that's not going to be delegated, is not going to collide, they know which one is a good one to take. So, questions, comments?

JONATHAN ZUCK: Yeah. This is Jonathan. I'll just jump right in here and ask, when you say end users might not be aware of it or readily adopt it, those end users you're talking about are system administrators? I mean, this isn't a choice made by a typical end user, is it? It's more like when I'm setting up active directory or something like that and I'm assigning it. So, it feels like it's the IT department.

WARREN KUMARI: Yeah. Probably, IT departments and often small medium-sized businesses, for example. Potentially, system administrators for huge corporations might have heard of it. But there's a lot of many, many, many millions of small businesses who potentially do stuff like that. But yeah, you're right. The IT admin for an active directory system or potentially weird Internet geeks who want to do this at home, want to name their printer or their Raspberry Pi something, and they might want to do this. But they're more sort of the weird geeky folk. Or software developers. I mean, there's also large vendors like D-Link, Belkin, etcetera, are very well known for just choosing .dlink or .belkin

and then bad things happen. Oh, is somebody running the queue? Not me.

MERIKE KAEO: Yeah. That's what I was going to mention. I mean, there's been a lot of cases where developers who just didn't know either building new IoT devices or some phone related item where the developer has made the decision and it's more organizations also. It's not always the IT department that has control over the DNS. It's almost any random person in the organization, which is, again, a problem that needs to be resolved.

RAM MOHAN: Thank you. Satish?

SATISH BABU: Thanks, Ram. This is Satish, and I'm a member of ALAC. Coincidentally, Eduardo and I currently are composing a statement of response from ALAC, and we have already made the first presentation to the CPWG where we discussed these things. And I'm happy to report that there was strong support for this move. And all these points that we had discussed right now, even the point about whether this is really an end user issue or is it a thing for a big company. The discussions basically said that even for a home network, this is important. So, it is actually an end user issue. So, yeah, we support it completely. Thank you very much.

RAM MOHAN: Thank you. Sebastien.

SEBASTIEN BACHOLLET: Yes. A question. I think for end user, it's a little bit far, but can't it be embedded in the universal acceptance work? Because it's one technical aspect that need to be spread around as a lot of other six with universal acceptance. And I suggest that it may be a good place to go to disseminate topic. Thank you.

WARREN KUMARI: Yeah. I mean, I think that at the moment, it's primarily the public comment staff to choose the actual string. But then, yeah, I think as widely as it can be explained, sort of announced, the better and universal acceptance does feel like a place where we're trying to reach as many users as possible. So, shout it from the rooftops whenever it is actually chosen.

JONATHAN ZUCK: Yeah. I guess the reason where I was coming from is it's obviously in the interest of end users. When we were talking at the end of your slide about advocating for end users to focus on it, it feels like UA in that respect and that end users aren't the ones that are going to fix the websites that are non-UA compliant, but they could be the ones that draw attention to it. So, that's the question is who delivers the message and who is the message ultimately for? That's why I was asking that question, and I don't think the message is for end users even though the outcome will be good for them.

WARREN KUMARI: Sorry. Just a quick additional note. I think that this is largely SSAC and ALAC's role in general, not specifically on this topic. Right? This is a generic role thing, not thou shalt go forth and do this, please.

RAM MOHAN: Other comments on this? Please, Amrita. Come on over.

AMRITA CHOUDHURY: Thank you. Amrita for the record. While it's not an end user issue, but we as At-Large can help to spread it to the right communities in our regions, in our geographic regions because it's very diverse. It could be the software developer, startups, etcetera. So, I think that's the role we play. It's not the end user per se, but there are different communities whom we are connected to.

RAM MOHAN: Yeah. And just to expand on that especially those of you who are engaged with technical communities, I think there is great relevance in passing this message along to start.

JONATHAN ZUCK: Once it's created and done.

RAM MOHAN: Once it's done. Yeah. Of course. I mean, so far, it doesn't look like there is any part of the community that is standing up and saying, this is a terrible idea. You should stop it. Warren?

WARREN KUMARI: Yeah. I guess just the last note. The public comment is store open, and closes in, I think, three weeks or so. They're generally, at least in my view, seems to be good support, but there is still some discussion on what the actual string should be. Does anybody disagree with the IANA's one? I personally like .internal, but obviously, that's my personal view.

JONATHAN ZUCK: They're up for that being one of the ones that's happens to be used the most already. Right?

WARREN KUMARI: That that one's already being used by Google Cloud, and Google made an announcement, we already used this. And Microsoft Azure and Amazon AWS all already use them widely. It's the highest, if we can go back four slides, it's the highest used undelegated TLD by a fair bit. So that was why, potentially, why IANA thought that would be a good one. It's already being used for this sort of thing very widely.

JONATHAN ZUCK: I'm amazed local domain is so highly priced.

WARREN KUMARI: Yeah. That's because local domain was sort of implied in a couple of RFCs as being a useful string, but I don't know what the official actual status was. I think people got sidetracked when writing the reserve names list or something.

RAM MOHAN: Steve.

STEVE CROCKER: Very, very briefly. Two things. First of all, congratulations, Warren. Nice piece of work and very helpful. And the second is, I think this is a good example of an identification and allocation of a string, a TLD string that is not part of the standard commercial use of top-level domain names where the idea is to sell as many subdomain names as you can. And so, I think in terms of the broader purpose of ICANN operating in the public interest, this is a positive worked example.

RAM MOHAN: Thank you. That wraps up this topic. Let's go to our final topic. Barry is going to lead us through that. We're going to share with you some of our thoughts and findings on the evolution of the name space. Barry?

BARRY LEIBA: This is Barry Leiba. I talked in Hamburg about this, and there's really not much change since then. This is the report, not the slides. Yeah. We need to go to the slides. The main thing that changed since Hamburg is we actually officially published the document, but there wasn't a

significant change in the document since then. So, this started with our revisiting what we had said about doe and dot and how that changed some aspects of how the DNS is used. And we started discussing what else is changing and where some of those changes are going to lead us. And that's what we wound up writing something around 30 pages about that. It's one of our documents that's more for information and is kind of light on any actual recommendations. It's mostly for you to read and enjoy or not.

The background on it is that while names play an important role in what we do with the Internet, and we know the DNS is there to turn names into Internet addresses so we don't have to deal with the Internet addresses, but we rely on the names. One of the key points is that the names are not just a convenience, but they're actually part of the trust model of the Internet. The names are an important way in which we, as humans, identify things and decide whether to trust them or not. And we're finding that the domain name resolution is becoming more ambiguous as different systems develop to use the names.

And names are becoming less visible to the users or less conspicuous perhaps is a good word, as you start scanning QR codes and click on links and interact with things in ways that don't involve you actually seeing a domain name and not being aware of where you're going and what you're doing with it. Or say a URL shortener where you're going to Bitly, and that's not really where you're going. Bitly is just helping you put it in Twitter in a few characters. So, the actual name is less obvious to you. How does this affect your view of the Internet and your trust model as a user?

There are also new needs for the entities that put together systems that use the domain names. Some of them want less centralization of how this all works. We'll get into what the various needs are. But that caused the development of alternative naming systems that we're discussing in this paper. So, that's where we are. Next slide, please. Kind of go light on this slide. We all know how traditional DNS resolution works. But the main point here is that DNS resolution has generally relied on the libraries that are built into your operating system as part of the Internet stack that says we're going to go to a local thing that knows how to resolve things in the DNS and knows how to go out there.

And what we said in the Doh Dot document was that we've started to come out of that with your browser now does that or your app now does that. And somebody said, for instance, that the Netflix app actually doesn't go out to the DNS. It knows how to contact Netflix and get the optimal thing for where you are and that sort of stuff. And it's optimizing your experience using a streaming service, making it a little bit faster, but not going and using the DNS to do that. So, we've pulled out of the traditional DNS resolution with a lot of these.

So, that's how I started talking about that. The motivations for changing this and coming out of what we've been using since the 1980s is speed enhancements is one of them, authentication enhancements, decentralized governments. I mentioned that people want to not have ICANN be involved or ICANN contracted parties be involved with how this all works out. There's also privacy concerns. And Go and Dot helped with that but there are other privacy issues involved with using the DNS that caused systems to start moving away from it. And

governments often use DNS blocking in various ways for censorship. So having alternative naming systems, alternative resolution systems allows some services to get out of that mode. That's all I want to say on this slide.

Next one, please. So, we've started looking at a selection of naming systems, and we didn't want to try to be exhaustive about every alternate naming system that's out there but give exemplars of different aspects of it. Things that this one does it for this reason, this one does it for that reason, this one has this mechanism, this one has that. So, we came up with some examples. One of them here, for instance, the Tor system, and you have an application that uses it. The Tor browser uses it, and here's what that does and how it works. And the paper talks about those kinds of things, how it accomplishes its resolution, why it's different from the others, and what the motivations for that particular system are.

So, let me look at this last bit of it. Yes. I had talked about the trust thing that when you use something like the Tor browser, it takes care of following the names that it needs to follow and doing its thing. If you take one of these systems, and this kind of merges into the name collision stuff. That if a name resolves differently in different systems, then you have a trust problem. You don't know what your browser is going to do as opposed to what this app is going to do resolving names differently. So, that can be an issue that we have to keep an eye on and try to coordinate the usage of these systems to avoid the uncertainty of what's going to happen depending on the application you're using to resolve a name.

So here are some of the examples that we came up with, and I don't want to spend a lot of time on this slide, and there's some fine print that we can't really read terribly well from this distance. But each of these five examples talks about a different way of resolving names in its service and different reasons why it developed. So Ethereum, for instance, if you have a name that ends in dot ETH, it's going to go through the Ethereum name service resolution. And if you put that into the DNS, you'll get one result. If you put that into an Ethereum name service app, an app that uses Ethereum name service, you're going to get resolution within their system. So, I guess, read the paper for the details on this rather than trying to read what's on the screen here.

Okay. Here's a larger print version of this, the kind for my age eyes. Looking at some of the different aspects, is it a centralized system or have they decentralized things? Is it secure? What are the privacy aspects? That sort of thing. And are the names that it uses human memorable? Is the mechanism that it resolves these clear to you as a human user when you look at the name? And some of them have taken different approaches to it because of their different needs.

So, the Tor system, the whole point of that is privacy and confidentiality. So, the security aspects are quite high. They don't care that much about the human memorability. That's not the primary need, but they did need it to be decentralized and secure and private. Whereas, for instance, Ethereum isn't looking at the security so much, but also wants the decentralization, unstoppable domains. The human memorability is the important thing for that. People come to them for nice names that they can use.

So here we're looking at the progression of things from before DNS when we had local names and IP addresses, and we passed around host's files and the bad old days before Paul McPheeters came along. And then in the DNS where we had consistent resolution, everybody did the same thing, everybody resolved names the same way, and we all knew what was happening. And then today where the names are the trust anchor, but the names are less obvious to us. And because of some of these alternative systems, we're not sure about the ambiguity. We're not sure about the consistency of the resolution. The same name resolves differently depending on how you're resolving it and what kind of what application you're using to do it. I mentioned before QR codes and other different ways of accessing things that obscure what you're doing and make the domain names less obvious.

So here are a couple of other proposals. We just heard about .internal, which came through ICANN. Through the IETF, we have the .alt name that's setting up as a top-level domain for alternative naming systems. And that was in the time between-- I think that happened right about the same time we published our document. The RFC got published that codifies .alt. So, both of these are voluntary things. And we know, as Warren said, the whole point here is to try to proselytize it and get people aware and hope that they will cooperate and use them rather than making things up themselves, but we can't force it. Widespread adoption of these mechanisms, though, could help the trust problem and help the ambiguity issue.

So, that's pretty much the end of it. That's all. Read the document for information. I would be interested in hearing from you if you find it useful. Actually, I would be interested in hearing from you if you don't

find it useful as well. Good and bad feedback is both fine. So, we do actually make two recommendations in this one where we didn't make any recommendations at all in the DoH DoT document. That was purely for information. But these recommendations are pretty lightweight.

They're saying because of these ambiguity issues, we need to keep track of what's going on, watch the development of new systems like this, and try to make sure we're aware of the problems that it may cause and the opportunities we have to avoid those problems. And keeping the community abreast of this by we've had some emerging identifier technologies panels at past ICANN meetings, and we're encouraging ICANN to continue doing that because we think they've been useful. So, that's all I have to say about that. Do I have questions? Please.

SATISH BABU: Thank you. This is Satish for the record. I'm a member of ALAC. I think the presentation is really useful. So far, we have had this alternate roots emerging without any degree of control or even knowledge about what's really happening and what are the risks. So, it is good that we have stopped to take stock of what is the picture now, and we're going to track it for the future and communicate these things to the community. I think that's very, very useful, and I think we will support you full on this. Thank you.

BARRY LEIBA: Great. Thanks, Satish. Others? All right. Thanks for your time, everybody.

RAM MOHAN: I think that brings us, Jonathan, and the rest of our-- Oh, Sebastien.

SEBASTIEN BACHOLLET: Yes. Just before you close the meeting, I wanted to ask you one question, but I don't feel I need the answer right now. With ATRT3, the proposal was to not eliminate for the moment, but review on security and stability. Do the SSAC thinks that it need to be back one day, or do you think that the work you are doing it's sufficient to no need so specific reviews on security and stability? Why I'm asking this question it's that we will be in discussion about possible ATRT4 someday. And the ATRT4, we will have to discuss that because they are supposed to be the one who will decide what we do with these two reviews. We stopped during ATRT3, and it may be good to have this thinking. Thank you.

RAM MOHAN: Thanks, Sebastien. It's a good question indeed. I'll provide my off-the-cuff responses, and we'll have to consider it inside the SSAC. In general, we found that when we've been reviewed and we've had feedback from the review of the SSAC and our processes and our procedures, in the main, they've been quite constructive. And in general, we've adopted them sometimes before they've even been ratified. Right? So, in general, for SSAC itself, those reviews have been good. But as far as ATRT 3 and 4, etcetera, should those continue? I don't think we've even thought about this. And to be honest, my sense is that most SSAC

members would like to focus on the tech things that we're working on rather than the policy process things.

JONATHAN ZUCK: There's a question.

SEBASTIEN BACHOLLET: Yeah. Well, I was not talking about ATRT3 or ATRT4 to be changed, but I guess I have the impression there's slight confusion between the organizational review of the SSAC and the Security and Stability Review. Yes. And I was talking about the SSR. Regarding the SSAC review, it is supposed to be evolved to a continuous improvement program.

RAM MOHAN: Okay. Sorry. My misunderstanding.

SEBASTIEN BACHOLLET: That's okay. Good to have discussion. It's exactly for that. We are here together.

RAM MOHAN: Okay. I think that's a good question. Why don't I throw it open to SSAC members here if you have a point of view? I mean, this is not official in any way, but do you have a point of view on SSR, utility of it, and utility of future SSR sessions? Anyone? See, they're all shy people. Oh, behind. Rod. There you go.

ROD RASMUSSEN: I was choking on my popcorn. So, Sebastien, when you asked the question a bit, I perked my ears up because you I think you've phrased it as the work that the SSAC is doing sufficient that we don't need SSR reviews. Or am I misinterpreting that?

SEBASTIEN BACHOLLET: That was my way to present the question. Maybe it's a good question.

ROD RASMUSSEN: Yeah, I will answer that one wearing my ex-chair hat on. I bet Steve will say the same thing. The work that SSAC does is not what the SSR reviews are for. Right? Do not count on the SSAC to do that job. Yeah. We're forward looking. We're looking at issues. We're trying to bring things in from the outside. The SSR reviews, I think, in theory, are very useful, in practice, those of us who went through the kind of the stop and the restart and all the stuff that happens five, six years ago, that was a disaster. That was more of a process disaster than anything else.

So, l my personal view is that I think it would be valuable to have them because that is the core mission of ICANN. And that if we are to do that though, we need to carefully consider how learn from the past and make sure that they we have a good process that's well scoped and is capable of delivering some useful information.

The one thing that we did in the last five years or so that would be helpful towards that was we did an environmental review. That was about three or four years ago. I think we presented it to Kobe, if I recall

correctly. That was kind of a formed baseline from our perspective of various things to look at. That would be a place to start, I would say. But there's some feedback for you.

RAM MOHAN: Thank you. Merike?

MERIKE KAEO: Yeah. From my perspective also, you want an external security review. So, you don't want it to be just internal, and SSAC, obviously, would be internal to ICANN. But having an external security-related review where you also have external security experts, I think, is extremely valuable for ICANN as a whole.

SEBASTIEN BACHOLLET: But this is not what's happened with the Security and Stability Review. It's people from the community who participate to this. You are not talking about that but-- Okay.

RAM MOHAN: Anyone else wants to respond to Sebastien's question? All right. So, second time's a charm. Thank you, Jonathan. Thanks everyone here for coming together to meet with us and look forward to seeing you in the hallways. Have a good rest of the ICANN meeting and look forward to engaging.

JONATHAN ZUCK: I can't wait for the meeting to start.

RAM MOHAN: And look forward to having other things to do together in between ICANN meetings because that's one of the things that we want to start doing.

JONATHAN ZUCK: That's right.

RAM MOHAN: Okay. Great. Thanks.

[END OF TRANSCRIPTION]