ICANN77 | PF – Getting to Know ICANN: OCTO and IANA
Tuesday, June 13 2023 – 10:45 to 12:15 DCA

SIRANUSH VARDANYAN: Hello, everyone, and welcome to the session, Get to Know ICANN. Today we will have our team from the Office of the Chief Technology Officer, which is called OCTO, and Internet Assigned Numbers Authority, IANA, here with us. My name is Siranush Vardanyan. I will be your remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN expected standards of behavior.

During this session, questions and comments will only be read aloud if typed in the chat space or Q&A feature. I will read questions and comments aloud during the time set by the chair or moderator of this session.

Interpretation for this session will include six UN languages, plus Portuguese. Click on the interpretation icon in Zoom and select the language you will be listened to during the session.

If you wish to speak, please raise your hand in Zoom, and once the session facilitator calls upon your name, kindly unmute your microphone and take the floor. Before speaking, ensure you have selected the language you will speak from the interpretation menu. Please state your name for the record and the language you will speak if speaking a language other than English. When speaking, be sure to mute all other devices and notifications. Please speak clearly and at a reasonable pace to allow for accurate interpretation.

To ensure the transparency of participation in ICANN's multi-stakeholder model, we ask that you sign in to Zoom sessions using your full name. name—for example, a first name and a last name or surname. You will be removed from the session if you do not sign in using your full name.

With that, it's my great pleasure to introduce our first presenter, the chief officer, John Crain. And I used to forget his position, but for me, he's a guru on the technical side of what's going on in ICANN. So, I will give the floor to you, John, and you will coordinate the presentation for your team.

JOHN CRAIN: Thank you very much. I am John Crain ICANN's Chief Technology Officer, and responsible for a group called Identifier Research Operations and Security Function, or IROS. Welcome to the world of acronyms. IROS is made up of two parts, one called OCTO and one called IANA, and we'll get to that later. I am not the technical guru. I just have a lot of smart people working with me that I can ask questions of, and if you see me answering a highly technical question, I'm probably getting the answer in Slack from one of them; so just, you know, giving away all the trade secrets. So, as I said, we are two other acronyms.; The Office of the Chief Technology Officer, which is all about data, research, education, information, talking about technology, teaching about technology, learning about technology. On the other side is the IANA functions, which is all about registering identifiers, names, numbers, port numbers, various protocol parameters. And we'll get into those in more detail.

Okay, so these are the … I'm not quite sure where I'm handing over, so I'm just going to do this one.

UNIDENTIFIED MALE: Yes, we decided that.

JOHN CRAIN: You just decided that? Excellent. As I said, I'm not in charge here. I've got my people.

Operations. This is how we divide our work inside the OCTO side of the house. So we have an operations function. It's like a little business all in and of itself. They make sure we do things like project management, fill out paperwork, do contracts, all the stuff that you need just to keep everything running. They're sort of the heartbeat of the group because they do all the paperwork we don't want to. We don't want our researchers spending time working on purchasing contracts for servers and things like that. That's what operations do.

We have an engagement group. We call them technical engagement because they are engaging with the community around technical topics. So, they do a lot of training and a lot of presenting. And I suspect one of my colleagues here will talk to you about that.

We have two research departments, one which we call SSR (Security, Stability, Resiliency); most important is that we just have another acronym. They're really looking at things like DNS abuse, how identifiers are abused, misused, and doing research in this area. And then the other areas are what we just call research, which probably

predates the [other one] a little bit. And they're really looking at the technologies themselves and the protocols and writing papers around that kind of thing. So it's a lot of smart guys doing good research, we hope, and publishing that for the community so that when we're having policy discussions, which is what happens at an ICANN meeting (this is a policy forum, and here's a bunch of techies sitting in front of you), it's important that policy is informed. So we try and do the research that helps inform the policy makers and others about what's happening on the Internet and what's happening in the technology space.

Those two research departments, the engagement department that goes out and actually talks a lot about some of this stuff and does training, and our operations are what build that group. And we'll talk a little bit about product, et cetera, in following slides.

So, David, are you going to take this one? I will pass you this controller. Thanks, John.

DAVID HUBERMAN: Hi. Good morning, everybody. My name is David Huberman. I am the Regional Technical Engagement Manager here in North America. And welcome all of you again here to Washington, DC. So DNS was created by Paul Mockapetris with RFC 882 in November of 1988, which makes it 35 years old. Okay, that's not that old, or I hope that's not that old. But in technical terms, that's really old. And the reason DNS works as well today as it did in 1988 is because we cooperate, right? We don't just, it's just not a couple of people sitting at an IETF meeting writing about DNS. It's you and it's me and it's your colleagues in your homes and at universities and at enterprises is. And we have been building this DNS

protocol and we've been building this DNS ecosystem for 35 years. We've done it together. And so here in OCTO, we spend a lot of our time working with technical communities and non-technical communities around the world to foster more cooperation, to foster more collaboration and coordination, because it's these types of activities which allow the Internet to continue to evolve and grow on a technical aspect so that it continues working for the needs of tomorrow and not just the needs of yesterday.

So we have a department inside of OCTO called Technical Engagement, and this is what we do. We are a group of six who are distributed around the world. I'm here in the North America region. Sitting here in the front is my colleague, Nicolas Antoniello. Nicolas is based in Uruguay and handles our South American, Latin American, and Caribbean regions. We have folks in Brisbane. We have one in Africa. And our boss, Adiel, the Vice President of Technical Engagement, who you've seen walking around this week, is in Montreal. And, very importantly, we have our colleague Carlos Alvarez. Carlos Alvarez does the same kind of technical engagement with the law enforcement community and trust communities all around the globe.

And it's this type of activity that allows not only for people who are outside the ICANN ecosystem, outside this ICANN bubble, to interface with us and provide inputs into the bubble for better data, but it also allows us to provide inputs into the routing community, into external DNS communities, into the numbers communities, into all these different communities that are maybe outside the bubble. And it allows us to coordinate, it allows us to liaise, it allows us to partner, it allows us to build things together.

Very importantly, one of the things we're working on these days is a program called KINDNS. KINDNS is a global effort to better secure the DNS ecosystem through the voluntary adoption of best practices that everybody can agree are a good thing. It turns out that people who like to do bad things, bad actors, like to leverage the inconsistent consistent implementations of DNS that are found all around the world. And where there are inconsistencies, they can find vulnerabilities. And where they find vulnerabilities, they can launch attacks. We have worked with the community and come up with a set of best practices that if everybody were to adopt, we would enjoy much better consistency in implementation of DNS and we'd have fewer vulnerabilities and we'd give attackers and people who want to do bad things, less opportunities to use our infrastructure against us. And so we call that KINDNS And if we, so we're working with communities around the world to motivate them to go to the KINDNS website, the project website, take some self-assessments, and adopt these best practices.

We do a lot of technical training, whether you are a ccTLD operator who wants to learn better practices for running your registry, whether you want to turn on DNSSEC and you need help with that. We do all kinds of technical capacity development around the world (country code TLDs, gTLDs) for non-technical communities to help better inform them, especially policymaking organizations, organization policymakers, government policymakers. It's good when they get good technical grounding in what it is that they're deciding on.

We also, as well, do lots and lots of engagement to talk about ICANN's technical initiatives at external events, at universities, with governments. And as well, we take activities that other organizations

do and we can do technical assessments of them. Especially with proposed regulations as governments around the world get more and more interested and more and more involved with regulating DNS space, regulating routing space, regulating number space, we can take in these proposed regulations and do technical assessments with the larger OCTO team to see what kind of impact they would have on our technical standards and in this technical community.

So that's a little bit about what Technical Engagement does. And if you're around this week and you see us in the hallways and you want to talk, if can help, please stop us.

Next up. Oh, here we are. This is pictures of us engaging technically. That's my colleague Nicolas on the left and that's me on the right.

So, John, back to you.

JOHN CRAIN: I'm going to do this on behalf of Samaneh Tajalizadehkhoob, who runs one of the research areas. So SSR (Security, Stability, Resilience) research is the area that focuses on issues more around security and abuse. So if you look at ICANN's bylaws …And of course all the Next Genners definitely have read the bylaws before they came here, right? And you read them every morning to indoctrinate yourself with the bylaws. And what does it say in there? It's all about security and stability. It's one of the fundamental things that ICANN is mandated and required to do: consider the security and the stability of the DNS and the identifier systems in pretty much everything it does. So this group is all about understanding what's happening in that world. So it's state-

of-the-art research and also educating. The purpose of our research is not always just to answer a question. A lot of the time it is about developing methodologies and spreading that methodology, maybe sharing that technique with the rest of the world so that others in industry can do likewise. place. So there's a lot of education. There is tool creation. There are tools. We had something called DAAR, which is a domain abuse activity reporting tool, which is basically a tool measuring reputation of the DNS. That was a tool both in the sense that there is a place you can go get data if you're a registry or a registrar in the future (that's coming) and the methodology that went around it, and many reports. Reports are also tools that people can use to understand what's happening in the ecosystem.

And as I said earlier and we keep alluding to here, policy needs good technical and scientific grounding. So this is all part of, again, informing the discussion. And we facilitate a lot of discussions around this. So, there are other people in the ecosystem doing similar things, and we collaborate quite strongly with many of those groups, be that M3AAWG, which is an anti-malware group, or the Anti-Phishing Working Group, FIRST, which is an incident response group. And in technical engagement Carlos Alvarez, who we already mentioned, leads a lot of that engagement and then brings in our researchers when they want to get deep into the technicalities of the research we're doing.

And the other thing that we all do within OCTO (but I mean we generally do this inside ICANN anyway) is we are subject matter experts. We're providing that expertise not only to the organization and the Board, but also to the community. So this is the SSR research area. I don't know if you put any slides in on briefing. So some of the key research areas I've

talked to include the DAAR project. That's DNS abuse. We have a project that goes a little bit further, and that looks for security threats or abuse around topics.

So one of the topics that you may have heard of that occurred in the world and caused a change in behaviors was something called COVID. I don't know if anybody's ever heard of it. I thought it was fascinating. I stayed home for almost three years for it, so it was good. Not really. So one of the things we did was we took terminology related to the COVID and the SARS vaccines and the disease, and we looked for how that was being used in strings within the DNS. And when we found things that indicated that it might be related to the pandemic, we would then go and do deeper searches to discover whether or not that was being abused. And then we would gather evidence. And then where we found a related name that was both abused and well evidenced, we would pass that evidence in a report onto the registrars so that they would take the name down. And typically, if you send a well-evidenced report of something bad on the Internet, even if it's not DNS abuse, but something that's really bad on the Internet, to a registry or registrar, they will take action. That is what we found. It's been pretty interesting.

We're also measuring some of the replacement infrastructure. RDAP is the replacement for WHOIS. We have a couple of papers out that have been measuring the response times of that system. We do that basically to see where the status is and if can we use the system.

And we have a couple of new things coming up, like DNS abuse prediction. This is the Minority Report if you've seen the movie, for the DNS. It's trying to figure out if something is going to be abused. There's

a lot of use of machine learning, etc. in that. That's ongoing work, so hopefully we'll see some papers coming out of that. And a lot of that once again will be about reporting the methodology so that others can learn from what we've done.

And we have another paper that just came out on parked pages, the pages that you often see where somebody registers a name, but they don't put up any significant content. They'll just put, like, a little advertising page or, you know, "If you're interested in this name, buy it from me." And we've been looking at how those are used and not used in relation to abuse, so we can identify the good ones versus the bad ones. And that paper's just coming out soon.

So that's the SSR research, but it's not the only research that happens within the Office of the Chief Technology Officer. And I think I'm passing it to Matt.

MATT LARSON: Good morning, everyone. I'm Matt Larson. I'm VP of Research. And this is Paul Hoffman to my left, your right. Yes, you're on my left, you're right. Left, right. So Paul is a distinguished technologist. And so I wanted to talk today about the research team. We have, as John just said, two research teams. We have the one that he just spoke about run by my colleague Samaneh, which focuses very specifically on security, stability, and resilience. And the research team that Paul and I are in has sort of a larger mandate. Anything to do with the Internet system of unique identifiers is fair game. And I would say that I like the wording here of the first bullet on the slide. One of our most important goals is indeed to provide trusted and verifiable information to the Internet

community regarding the system of unique identifiers. So we're fortunate to have various data sources, both that we gather ourselves … One for example would be from the root server that ICANN manages, the IMRS, yet another abbreviation; the ICANN Managed Root Server. So that's a wonderful source of data about the DNS that we use all the time in our research. We have other long-running processes that make DNS queries and store the results so that we have longitudinal data sets tracking, for example, the changes in TLDs related to how they're configured and DNSSEC parameters and things like that.

We also work with various other communities related to the identifier systems. So one, for example, would be the DNS-OARC, the DNS Operations Analysis and Research Center, and we collaborate closely with them. They hold workshops and we give presentations, and it's just a place for the community to hang out, and that includes us.

Another thing that we do is participate in standards development. So the IETF, the Internet Engineering Task Force, is where many of the Internet standards are developed, including those for DNS and other identifiers. And so Paul, in particular, on our team, is a very active participant and has written many, many RFCs.

Do you know the count, Paul?

PAUL HOFFMAN: 75-ish.

MATT LARSON: Okay. There's so many, he doesn't even know the exact count.

So then we have various other things that members of the team specialize in, based on their specialties and interests. You may have heard in the hall something called NCAP, the Named Collision Assessment Project, And I must have done something wrong at some point because I'm the person who's monitoring that (more than just monitoring, I guess) on behalf of OCTO. And then I also represent ICANN.org on RSSAC, the Root Server System Advisory Committee.

But I think maybe what's more interesting is to talk about specific research activities that the team has done. One is ITHI, the Identifier Technology Health Indicators, and the idea here is that we measure certain aspects of the identifier system over time. So we have longitudinal information. I go to my doctor once a year, and I have blood work done, and, you know, he compares what are my statistics, (what's my cholesterol now versus last year versus the year before) and because we have that data about my health, can see trends and make decisions (stop eating eggs or whatever). So that's been a long-running project that we have data from various sources, and we're always looking for more data sources. So if you run a network or know someone who runs a network who would be interested in contributing data, please get in touch. It's a very lightweight process to collect and contribute the data.

We do all manner of research with all levels of the DNS ecosystem, whether it be related to stub resolvers that are on every device, like your phone or your laptop, or recursive resolvers that are typically at ISPs, or authoritative servers, like TLD servers and second level domain servers. So we do various things.

One interesting bit of research we did recently, during the pandemic, we had some good data from ISPs in France, and were able to study how Internet traffic (DNS traffic in particular) changed as a result of the pandemic, and it was very interesting that you could clearly see the shift to working at home in the data.

So one area of research is not related to domain names is addresses, another other identifier that ICANN is certainly helping to coordinate. And there's something called RPKI, the Research Public Key Infrastructure. And this allows route network information in the routing system, specific networks. You can cryptographically sign them and know where they're coming from. And so the short version is it's a way to add security to the routing system. And we've just started to engage and do studies on that. that. We've written a paper about RPKI.

As an example of some standards work we're doing, another colleague and I are working on a new protocol, so it'll be a change to the way DNS works, called DNSSEC error reporting. And this will allow … I don't know if you've heard of DNSSEC yet, but it's the way that DNS information can be cryptographically signed and verified. Well, what happens if that verification fails? If I'm signing the data, I would want to know if out there on the Internet somewhere it's failing. Maybe I forgot to re-sign it. Maybe something happened. Maybe somebody's trying to attack my domain. So wouldn't it be nice if there were a way to know when that verification failed? And that's what this new protocol will do.

We published various data sets related to TLD information. And then also the ICANN IMRS, the root server, is a wonderful source of data. And one of the most recent things we've done, hot off the presses, is a paper

that Paul wrote about tracking round-trip times from different nodes that make up IMRS.

So with that, I'll turn it over to Paul to talk about the OCTO document series.

PAUL HOFFMAN: Hi, I'm Paul Hoffman, and as you've heard before, and one of the things that I do in OCTO is I'm sort of the publisher, editor of the document series for OCTO, for all of the parts that John talked about. So these documents are written generally by people in OCTO, sometimes outside. We started the series about four years ago. So, we're at about 35 documents. Some of these documents get revised over time. So if you look into the document series, the URL is at the bottom. You'll see some of them are in version two or version three. And sometimes when you go to a publisher, you expect this publisher will only have science fiction, or this publisher will only be doing this. As you've heard, the kind of research and the kind of work that we do in OCTO is quite varied, so the documents in the OCTO document series is quite varied. Here are some titles of the recent ones. From looking at the titles, I think what you can see is that we cover lots of parts of the DNS this space in different ways. Some of these documents are very research-oriented. "Look, we've got this pile of data. We did this analysis, and here's what you should know." Some of them are really meant to be not setting policy, because we don't get to set policy. Remember, that's what you all do. We on staff don't set policy, but we help you do that. So some of them are really aimed at if you're working in policy on a particular area,

and there is a technical aspect to it, helping you know what's behind it. What are the technologies?

So for example, let me just pick one out here—well, actually, one that I wrote on quantum computing. So one of the policies that people are concerned about is how do we deal in the future for DNSSEC about new adversaries, which in this case … And I won't go into it. I love mentioning quantum computing. Everyone goes, "Ooh, quantum." So I got tasked with learning what the ooh-quantum stuff was and how it affected the DNSSEC, as well as other parts of the DNS, because, in fact, the DNS uses TLS sometimes, and TLS is affected by changes that will be coming from quantum computing. So it's a short document. It gives you, like many of the OCTO documents, an introduction to what we're talking about. It then dives into the technical realities, which might be more than what you're interested in, especially if you are a policy person. But then it wraps up with either suggestions on where you might look more, ideas for what's important for you, and things like that. So it really is meant to both help people who want to understand the technology more, but also the people who want to understand the technology to go into other things.

So, again, at the bottom there's a link. Upcoming documents are also sort of a wide variety and such like that. We're always interested to hear from the community about whether we hit the right spot on some of the documents. Sometimes we get comments saying, "This was great, but I wish you had talked about this more." And we might come out with a new version on that. So, that's that.

And I think then that finishes us. Yeah.

JOHN CRAIN: Okay, so that's the OCTO side of, or the Office of the CTO side, of what we do.

I think we're going to take a few questions about that, and then we'll get back to the IANA side after that. Is that correct?

SIRANUSH VARDANYAN: Yes. Before we move to the second part, if there is any questions, please raise your hand here or … I don't see anyone in the Zoom, but I see the questions here.

Yes. Please, Hafiz, go ahead.

HAFIZ FAROOQ: Thank you very much. First of all, thank you very much for giving us a wrap up of the complete OCTO. Just a question about the domain abuse, especially the DAAR project. Is it only for the review of the bad domains or the bad TLDs also? Because I will give you an example of .zip and .mov recently released by Google, which is a cyber security concerned for many companies around the world because of abuse. So is it under the DAAR project or is it something different?

JOHN CRAIN: No, no that's a very different problem space. DAAR is very specifically focused on the five security threats that we specify, as you know—things like phishing, malware, botnet, command and control. And is really a look at the reputation of those names. The .zip issue is a

completely different issue, and that's around how user agents, like your browser or whatever, deal with file extensions. And that's a whole long discussion, and that's one that's recently started here in the ICANN world. So I think the SSAC's looking at it, and other people are looking at it.

I do need to point out that we've had file extension strings as TLDs forever. There are thousands of file extensions that are ex-executable. So .zip's not the only one. Another one you could think of is of course .com. I did a study on this, I think somewhere between 15 and 20 years ago. I'm old. I can't remember exactly. And at that point (this is before the new TLDs), there was 80 or so in the zone at that point.

So I think there's going to be a lot of discussion on what is the exact threat vector here, and what are the solutions. But it's a fairly new discussion that's just come up again after about 20 years.

HAFIZ FAROOQ: Okay, thank you very much.

SIRANUSH VARDANYAN: Namra, please go ahead. Please speak at a reasonable pace.

NAMRA NASEER: Yeah, okay. Thank you. My question to you is … Any research that happens is as good as the audience that is reaching, right? So, I really want to understand the dissemination process. I know it's published, it's online, available for the public, you disseminate this knowledge in meetings like these, but are there any concerted efforts to kind of make

this research, and sometimes recommendations in those research, to the right, relevant stakeholders? This could be government stakeholders, other people who actually implement policies, and other organizations also. This could be the private/public sector.

And secondly, I really want to know, what's the process and who decides which issues need to be addressed and which issues need to be researched on? So I really want to understand that process. Thank you.

JOHN CRAIN: Okay, so there was about 10 questions hidden in there, which is really clever. So the first question is really … I think there's two main questions. The first question is really like how do we engage around the research that we do and how do we promote it? And the second one is how do we choose what to research? So I'm going to start with David because he's the technical engagement. We have other engagement areas within ICANN that he works closely with, so maybe he can talk a little bit to that.

DAVID HUBERMAN: Hi. Thanks, John. No, great question. So one of the things I really, really enjoy about working at ICANN is we're 446 employees who all work as one team. So when Samaneh's group, when Matt and Paul and their group, produce some of this wonderful research and draw these important conclusions, one of the ways we're able to disseminate this is by disseminating it first very quickly inside of ICANN and then saying, "Go forth and tell everybody you can about it to whom it's relevant." So on Mandy Carver's team, our government engagement team, when we

produce research, and when regulators, policy makers, and communications ministers they need to know about it, they go out, they can bring … If Paul writes a document, they can invite Paul, and Paul can come in and talk to an important audience about what he's written. In the technical community, we take some of these papers that we know are interested to network operators, to DNS operators, and we will talk about the papers. We will talk about the conclusions. We will talk about, what does this mean for you? And every time a new paper is written, we have slide decks and talking points, and we can disseminate them very quickly around the world, and it's really nice.

And then, which issues are subjects for research? Matt?

MATT LARSON: They come from all different ways, and I'm sitting here trying to think of all the ways. One is just stuff that we think of on our own that could be having your morning coffee—

UNIDENTIFIED MALE: Shiny object.

MATT LARSON: Shiny object, yeah. That's one way. Another is just reacting to the community, reacting to what's going on. Paul's example of quantum cryptography is a good example. I mean, that's just not as much now as it was maybe a couple of years ago. Now AI is the hot topic, right? And before that, blockchain, and before that, quantum. But so it was something that was taking a lot of space in the technical press. And in

particular, we thought, you know, people are going to want to know about this, how it intersects with the identifier system and DNSSEC, so we wrote a paper on it.

Another technology that was, I would say, really just within a smaller community but was talked about was something called new IP. And so we started to get questions about that from the community, and at some point, when the technical engagement folks get more and more questions, we realize, okay, I guess we need to do something about this. We need to be able to research it ourselves, come up with our own assessment, our own conclusions, and then publish about it.

Paul, can you think of anything else?

PAUL HOFFMAN: It's hard to actually think in specifics because of what you just said, namely that we hear it a lot from TE, but also sometimes we hear it in the hallway, and we won't say, oh, look, somebody asked us, one person asked us, therefore where we're going to run off and do it. But if I say somebody said it, and Matt goes, oh, yeah, somebody else said it, then that's an indicator to us. But we also do it when we do our outreach. And as David said, a lot of us who are not on TE get called into doing things. The blockchain naming is a good example. Alain Durand, who's also on our team, has given talks. You think you had it bad having to do some of this stuff. He has to go and give talks on blockchain naming and such because he's the author of the document. Then we hear back from those people also. “Oh, well, what are you doing about such and such?” Often we can say, “Here's the document we've already done.” But sometimes it comes back to, “Okay this is the fourth time

I've heard this. What should we do?" So it's informal, but it definitely drives some of it. And again, they drive each other. We do do a lot of outreach. I'm in California, so any European meetings I have to get up at crack of dawn for, but we do do those as outreach, and then we, you know, we hear from people that they liked the papers. So thanks.

JOHN CRAIN: Yeah, and the final mechanism that's sort of relevant here is we have advisory committees and working groups and all kinds of things here within ICANN, and they can pose a question, formally or informally. So if the board, when they're meeting everybody, hears a lot of things, they will come to the CEO and say, have you guys thought about this? And it's the same process. "Here's the document." "Oh, no, we've not thought about this. Do you want us to spend resources on it?" So some of them come in in a little bit more formal manner from this very environment where I suspect, on the user agents and DNS, if that becomes a common topic that people are interested in, then we may go and do some study and write something up or we might not need to do the study. We might just be able to write it. So it does come from all different manners, different mechanisms.

SIRANUSH VARDANYAN: Thank you. Before we come here, we have a question from remote participant Imran. If you can unmute yourself and ask your question. Thank you.

IMRAN HOSSEN: I have a questions about the two current activities. One is the DNS [TICR], and another is [park space]. So how can we know about these two activities and how can we be active in these two activities? How can we contribute or how we can get involved in these two activities?

[PAUL HOFFMAN]: So I think the question is, is once we've published something and someone goes, “Oh, this is interesting. How can we get involved?” we try (we aren't always successful, but we try) in our documents to say where this work is going on and such like that. Sometimes it's not going on anywhere, but we do try to do that. But if we have failed, and TE hears about it in one of theirs, we will sometimes revise a document to put in more links and such like that. But I will admit, we are not consistent in our ability to say, “Now that you've learned about this, either in a document or in a training, here's where you can go.” And you know, I will admit, there's only so much that gets done by “Look at the link on the last page.” You know, sometimes that link is general information and not a place where you can be active, and sometimes it is. We would love to get better at that because one of the things that I have found running the document series is the V2 documents are often much more valuable than the V1s, and those come from somebody saying, “I got active and this is what I discovered.”

SIRANUSH VARDANYAN: Thank you. We'll go to Samik now, and meanwhile, I will ask my colleague Deborah to have the new slides up. Samik, please go ahead. At the end, we'll have time to go back to questions, so please keep your questions for the end as well.

SAMIK KHAREL: Hi. Samik Kharel, ICANN 77 Fellow. I would just like to know how an end user can be involved with the technical community because this technical community is mostly a very closed community, so there's a wall. So for the starters, how can we be involved in your research projects and how can you help us and will guide us there?

DAVID HUBERMAN: Hi, great question. So the technical community is actually open for everybody to participate, including as an end user. The barrier to participation used to be cost. It used to be, were you able to show up in this city, in this random country in the world? And typically you would need a company or a university or somebody to pay for that. Today, we have so many hybrid meetings that we can participate in many technical fora from the comforts of our home in front of our devices, right? The network operators community, the DNS operators, DNS registries and registrars, the IETF, where we develop protocols—these are open. They're very often mailing [list] space, so they can just be done asynchronously over e-mail. When they have meetings, they're very often hybrid.

And what I like about technical community collaboration here in the Internet technical space for really the last quarter century is nobody cares who you work for, where you come from. They care about what your ideas are. And so if you, as someone who might be new to the industry or early in your career, have good ideas, those are just as valuable as someone like John who's been doing this for 40 years. His

ideas aren't more important than yours just because he's old and experienced.

JOHN CRAIN: My ideas are not very good.

DAVID HUBERMAN: Your ideas are really just as important. So it's actually critical that, as an end user and as a newcomer, you get involved with the things that you know about because you have a fresh set of eyes and fresh experiences, and those are critical.

SIRANUSH VARDANYAN: Thank you, David.

[PAUL HOFFMAN]: Can I toss one quick bit on that? A lot of these organizations also have How To Get Started documents. In the IETF, it's called the Tao of the IETF. But also, if you are stepping in for the first time and you're like, "Oh God, I'm going to make a mistake," just say so. Say, "I'm a newbie. David Huberman said I had a good idea," or whatever, and then step back and you will get help. And IETF has a whole mentoring team for first-timers. ICANN … You all are here on the ICANN stuff. Assume that you will get help. And if you don't, then come around again. But since you started this on the technical people, surprisingly the technical people are much more helpful than stepping into the public policy space as a newbie.

SIRANUSH VARDANYAN: Thank you very much.

So thank you, team. But we will be back, so please, if you have questions, I hope that we'll still have time for more questions. But you know the faces now, so you can talk to these people in the corridors during the networking sessions. So please feel free.

You have heard a lot about IANA. What is it? This is another abbreviation in ICANN world? What does it mean? So, Internet Assigned Numbers Authority, how important it is, and it's my great pleasure to introduce you Vice President for IANA Services, Kim Davies, and his team, who is with us here. Mariana is here with us, as is Marilia. So, Kim, please, the floor is yours. What are you doing, and how important is IANA for the entire technical world?

KIM DAVIES: Thanks, Siranush. Hello, everyone. As you heard, my name is Kim Davies. I head up the IANA team, but I'm certainly not alone. We have several of my colleagues here. Merilia was just mentioned. I will also point out Candice and Selina who are over on the side there. My presentation is really just a high-level introduction to what IANA is, what we do, and why we're important.

So you know what IANA stands for now, the Internet Assigned Numbers Authority. IANA is essentially a global resource that holds the definitive records of the Internet's unique identifiers. Unique identifiers are very

prevalent in our usage of the Internet. There are identifiers that you see every day as you interact with your computer and navigate the Internet. And there's many identifiers you probably don't see that are behind the scenes, but nonetheless are essential to successful Internet communications. You might hear terms like protocol parameters, Internet number resources, domain names, etc. All of these are in our wheelhouse. What we do for any individual identifier type varies depending on the policy, but generally speaking, IANA's responsibility is to be that global centralized clearinghouse for all these different types of unique identifier systems.

The IANA as a function is something that predates ICANN, probably predates most of us in the room. It's really a function that evolved from the very earliest years of the Internet. As the Internet emerged as an experimental project, there was a recognition as it grew beyond a couple of devices that there needed to be some kind of central function that coordinated some of the elements to keep it working. And what we now know of IANA today really traces its roots back to the 1970s. On the picture there, you see in the white shirt a man named John Postel. He is my predecessor in this role and really the progenitor of the IANA functions. In that picture, he's there with Vint Cerf, one of the co-inventors of the Internet.

Next slide. So what are the IRNN functions? We talked about what they are in terms of conceptually. Today, they are a set of functions that are within ICANN's area of responsibility. You heard about IROS and how we are under the rubric of John Crain as CTO. There is some complexity there that I'll get to in a moment, but generally speaking ICANN was established in the late 1990s to be the home of the IANA functions. It

was one of the original reasons why ICANN as an organization was established. functions were established. Up until that point, the IANA functions were the result of different contracts issued by the U.S. Government to research organizations and so forth, but the 1990s really spurred an effort to professionalize how a lot of Internet operations were done, the Internet was really expanding in popularity, there was a lot of interest, a lot of people started to depend on the Internet in the 1990s, and that was the motivation for establishing ICANN to be a home for the IANA functions.

So today, the IANA functions are operated by a team. It's a department within ICANN, and that team is responsible for coordinating unique identifier assignments. What that involves I'll get to in a moment. So when you hear the term IANA in common dialogue, it's not a standalone company or entity that operates by itself. It is a set of functions. It is operated by ICANN. ICANN technically subcontracts a different organization called PTI. I'm the head of PTI. PTI is an affiliate of ICANN. It's closely connected to ICANN. PTI fulfills a specific governance role but in terms of the operational performance of the IANA functions, which I'm here to talk to you about today, that distinction is probably not that interesting. At the end of the day, the IANA team sits in the ICANN office, we work closely with our ICANN colleagues, so it’s, in many respects, a distinction without, without a difference when it comes to day-to-day functioning.

So I talked about IANA being this central clearinghouse for unique identifier assignments. Why do we need such a function? Why is it important for us to have some entity that is centralized? After all, you hear about it when we talk about the Internet. The Internet's

decentralized. There's no central authority. You don't need permission to connect to the Internet. You can go anywhere. It's really this global network comprised of many thousands or tens of thousands of networks all connected together. Why do you need some centralized authority? Well, essentially every device on the Internet needs to speak the same language to one another. Devices communicating to one another need to express themselves in a way that each device will mutually understand one another.

Some of these are very easy to reason with. When we talk about domain names, which is the bulk of what we talk about during an ICANN conference, when I type in a web address like ICANN.org on one device, I expect to go to the ICANN website. If I type that web address on a device in another country, in another place, I expect to get to that same website. The ability for identifiers to go to the same place and have the same meaning the Internet to operate as we expect it to operate. If one domain name took you to a completely different place depending on where you were in the world, or an IP address took you to a different device, et cetera, et cetera, the Internet wouldn't interoperate and it just wouldn't work as we would expect it to.

So our remit is very narrow, but it's essential. And the essential component is that for those parts of the Internet that need to work the same no matter where you are in the world, through protocols that are devised typically at the IETF and through assignments made by the IANA to give certain names, numbers and parameters certain meaning, that is how the Internet continues to work the way we expect.

Now the registries that we maintain (the bulk of them) are used by software vendors, service providers and the like, but are typically not used by end users. It's relatively rare that an end user of the Internet would know about us or interact with us. The vast majority of the work that we do is with software vendors who implement the codes that we manage into software.

So I know today I mentioned PTI, so I won't go on about that too much. We are a team of around 16 people, although this slide is a few weeks old. We just added two more, so now we're 18 people, and we're predominantly based in Los Angeles at ICANN headquarters.

All right, so let's dive into the functions themselves and talk in a little bit more detail about what those functions actually are. When we talk about the IANA functions, we typically reason with them, describe them, in three main categories. These categories are a little arbitrary, and I'll explain why in a moment, but what they have in common is that they all require unique coordination at a global level to work in the way we would expect.

I'm going to start with protocol parameters in this discussion, because in reality, everything we're talking about is a protocol parameter. Number resources and domain names just happen to be highly specialized form of protocol parameters that have different kinds of oversight, but fundamentally protocol parameters and unique identifiers are somewhat interchangeable as terms.

So on protocol parameters, there are literally … I don't know the count at last count, but we maintain roughly 3,000 different types of identifiers, and in each type there can be anywhere from a dozen to

100,000 registrations. So for our very popular registries, we have many different allocations. We make new allocations every day. For other kinds of allocations, it's very rare.

Just riffing off one of the questions and comments from the previous part of the session (they're talking about file types, .zip, and .move), a good example of something that the IANA team does is we actually manage something called media types. The way the modern Internet works is that actually file extensions don't matter. When you download something or you visit something on the Internet, the file extension is typically not how your computer works out what kind of file it is or what the content is. Instead, what happens is, in the transmission, behind the scenes, there is a transmission of something called a media type. And the media type instructs the device receiving the transmission what kind of data is associated with this transmission. So we maintain a registry of all these different media types, so anytime someone creates a new format, a new transmission format, they will come to us and apply for a media type allocation, we will give that to them, and then they have a unique identifier that from that point on, any time that their file transmission or data transmission happens over the Internet, they can mark with that media type and it will be consistently recorded at the other end what to do with that data transmission.

So that's just one example of, again, many thousands of registries. You can scroll through a big, long list at IANA.org/protocols. But that gives you a sense of something where you as an Internet end user probably have no idea that media types even exist. It's something that software vendors need to implement as they build Internet technologies. If you're writing a web browser, for example, you absolutely know what

media types are, but an end user wouldn't ordinarily need to know what that is.

So protocol assignments generally are directly allocated by IANA—again, mostly software developers, protocol developers. There's not a lot of structure around that. And this is a key differentiator between protocol parameters at large and specifically number resources and domain names. So I'll drill down into how those work in a moment. But number resources and domain names are highly specialized. They have a hierarchical allocation model.

So if you want a domain name, if anyone in the world wants a domain name, clearly you don't come straight to us. Firstly, we're, as you saw, a team of 18 people. If we had to register every domain name in the world, that would not scale. So there is a hierarchical allocation system. IANA allocates the top level. Other organizations allocate levels below that. For a number of resources, it's a very similar model.

So let's move into what that means—

JOHN CRAIN: [inaudible]

KIM DAVIES: Alright, John has a conflict, so he's going to jump in and say two words.

JOHN CRAIN: I'm just going to say, I have to run to something else. Something just came up. So I'll leave you in the capable hands of my colleagues here. Find me in the corridors if you need to speak to me in person.

SIRANUSH VARDANYAN: Thank you, John.

KIM DAVIES: All right. So where do these protocol parameter registries come from? Who invented the concept of media types to begin with? Typically, the Internet Engineering Task Force, the IETF, is the preeminent standards body from which almost all of our work derives. People create Internet standards there. They define the technical standards on how the Internet should work. These standards are published in documents known as RFCs. And RFCs almost always have a section in them called IANA considerations. The IANA considerations section will inform IANA on what it needs to do to make this protocol operational. And some typical standards have no IANA involvement, but many do. And that section will inform IANA on how to allocate identifiers, what the policy should be, reserved values, initial registrations, and the like.

All right. So let's get on to number resources. Number resources is the second of the three groups of how we organize our work. Again, number resources is just a specialized form of protocol parameters. And when we talk about number resources, we really focus on three areas. One is the allocation of IPv4 addresses. The other is the allocation of IPv6 addresses, and the third is autonomous system or AS numbers.

What are IP addresses? For those not familiar, every device that's connected to the Internet needs a unique identifier. This is how at a very fundamental level when you transmit something over the Internet, it gets to the right destination. So what is an IPv6 address? IPv6 connection. So when you connect to the Internet, you're issued with a unique number and that number is an IP address. There's two different kinds of IP addresses. Version 4 IP addresses date back to the early 1980s. This is universally deployed. It's undoubtedly something you use almost every day. The challenge with v4 is effectively we're out of numbers. The amount of numbers available is around 4 billion, and essentially all of those numbers have now been allocated. v6 is the replacement. It adds a lot more numbers, but it is not universally deployed. It is well-deployed, it's certainly available in this conference room, it's available on the most modern networks that you connect to, but it's not available everywhere. So it's not a complete replacement yet for v4, but that is the ultimate aim.

On autonomous system numbers, the way I tend to describe them is to think of them like postal codes or zip codes for the Internet. Maintaining instructions on how to direct traffic on a per IP address basis would be very cumbersome. So what network operators do is aggregate their networks into a single sort of postal code for the Internet called an AS number. And it's AS numbers that are used to dictate how Internet traffic is directed across different networks over the Internet.

So what all three of these identifier types have in common is that IANA manages the global address space but does not directly issue them to end users or to networks for that matter. We only allocate large blocks to organizations known as Regional Internet Registries. You might hear

the term RIRs. RIRs each is individually responsible within their region for firstly setting the policy on how individual allocations are made but then implementing that policy as well. So each of those five RIRs has a staff, an application process. And depending on where you are based, you would go to the RIR within your region if you needed an allocation of IP addresses or if you needed an allocation of an AS number for your network.

I will just note here that as an end user, you don't need to go through this process. When you connect to your ISP or when you connect to a hotel network like here today, you're automatically allocated a single IP address through the connection protocol with your Wi-Fi. So an end user doesn't need to get their own IP address. But where does the hotel get its IP addresses from? Where does your ISP get its IP addresses from? From the RIRs.

All right, lastly, (and I suspect the area that most of you are most familiar with, since we're at an ICANN conference) is domain names. So again, the domain name space, the domain name identifier space, is actually just a specialized form of unique identifier. In fact, it's not the only kind of identifier used in the domain name system. There are other kinds of identifiers (resource record types, security algorithm types, header flags and many more that) that, again, are only particularly interesting to software vendors that are implementing the DNS protocol. And all of those other ones in grey there are directly allocated by IANA. But when it comes to the domain name space itself, again there]s a very hierarchical allocation model.

So I think many of you will be familiar with a diagram like this. Domain name system is hierarchically allocated. We tend to depict it as a tree. At the top is the root zone. The root zone is IANA's bread and butter when it comes to the domain name system. We administer the contents of the root zone. the root zone contains the official list of what is a top-level domain, and then pointers to the technical operations of each TLD operator. In turn, operators (and on this example illustration is .org, .us and .bel, which is the Cyrillic for Belarus) in turn delegate the next level below them, and so on and so forth. So an end user who wants a domain name registration to us. They would go to a registrar or a registry for the particular namespace they're interested in. But the TLD operators themselves do come to us. We have relationships with every single TLD operator that runs, you know, .org, .us, et cetera.

So in addition to being sort of the authoritative record of what is a top-level domain. The root zone also has other kind of associated technical and operational data associated with that: pointers to WHOIS and RDAP servers, who the administrative and technical points of contact are, and things like that. Our day-to-day business for managing the domain name system is to receive requests to update these details. Many of them are routine based on personal changes, technical upgrades, and the like. TLD administrators will come to IANA and ask for changes, and we will implement them in the root zone.

But we also receive a request to actually add and modify TLDs at a more fundamental level. And when we add a TLD in particular, you can't just come to us and say, “I want .kim.” There are policies. There are rules. And those rules are set in this community. We are here to implement those rules. So we assess applications against the eligibility availability

requirements and do our analysis, do due diligence, in the case of gTLDs, make sure there's contractual relationships that have been established and the like, and only then do we make changes in accordance with those instructions.

One thing we don't do is we don't operate the root servers. The root servers are the actual servers on the Internet that get thousands of queries every second for DNS lookups. We manage the content but the distribution of that content is done by other organizations.

Another thing that we do in the context of root server operations is manage the root key signing key. This is a cryptographic security mechanism that that is part of the root zone and essentially enables devices to confirm and assure that DNS data that they receive is accurate. Because of our unique role at the top of that tree that we saw a few slides ago, it's really essential that the key signing key be protected and be kept secure, more so than any other key in the system. So for that reason, we have a very deliberative strategy in how we manage the KSK, and in fact it's a very unusual process because we don't keep the key secret in how we protect it. We actually are very transparent. We have these public events called key signing ceremonies that people can attend. People watch how we administer the key. You can follow along on YouTube in real time if you like. It's a very transparent process, because we believe that the way for the community to trust us in managing the key is to watch us do it. So there's lots of opportunities for the community to participate in key signing ceremonies, watch how we operate the key, and satisfy themselves that we're operating it in a proper way. And the purpose here is really to ensure trust in the process. We shouldn't just say, “The

root zone's secure. Trust us." We want to say, "The root zone is secure. See for yourself." And particularly security experts can look at what we're doing and satisfy themselves that we're managing the key in an appropriate way.

Here are just some other functions that we have, particularly in the domain name space, that's probably interesting. These aren't core functions but functions nonetheless. And we operate a small top level domain called .int. .in is strictly limited to intergovernmental treaty organizations. So it's just a bit of an unusual TLD, but for historical reasons is why it's part of IANA, but it still remains the case. .arpa domains you've probably never heard of. It's part of sort of the technical plumbing of the Internet, again very interesting to software implementers, not that interesting to anyone else.

We also manage a repository of label generation rule sets. You may have of them as IDN tables. This is where registries who register internationalized domain names share their practices and their rules. It's important, in the context of IDN registration, that that be done. And IANA acts as a centralized clearinghouse for that work.

All right, so that is a fairly high-level summary of the IANA functions. Hopefully, that gave you a sense of what we do. Obviously, to go into the specifics of each registry would be a long discussion, but this is intended to be relatively high level.

I did want to note that security of IANA is more than just that DNSSEC, that key protection, that I mentioned a moment ago. We really try to maintain the confidence of the community that we're administering all the identifiers we do in an that we do in an appropriate and responsible

manner, and so this includes things like having dedicated workflow systems that we implement, access controls, separation of systems. And most importantly is that last bullet point. We have regular third party audits. We bring in third party auditors every year to analyze how we're doing work and make sure we're doing it in an appropriate manner.

On other aspects of the work that we do, we're impartial and neutral in what we do. We don't set policy. That's the community's job. We're here to neutrally implement the policy that the community has set. Performance is important to us. We strive to make sure that when we do transactions, we do them in an appropriate time frame that our customers expect. Continuous improvement is important. We never want to rest on our laurels. We want to continue to advance, make sure we continue to stay ahead of expectations. We have a lot of recommendations of what we do. I mentioned audits on the last slide. Accountability is important, too. There's many different committees, oversight bodies, in the ICANN sphere and beyond that are tasked with making sure IANA is doing its job effectively and appropriately, and we support those oversight mechanisms, and we produce many, many, many reports on how we do our work as well. So there's a whole performance section of our website, and you can look at dozens and dozens of performance reports every month on how we do our work.

So to summarize, IANA maintains these registries of unique numbering systems that keep the Internet interoperating. Many of them you've never heard of. They're relatively straightforward, generally of interest mostly to software vendors, and generally not visible to the end user. However, we are responsible for a number of identifiers. We are

responsible for a number of relatively high-profile identifier systems—domain name system, IP addresses, the top ones. And for those, we maintain sort of the global route. We don't maintain the whole system. We work with partners who do much of the work. We just maintain that very top level of those identifier systems.

So with that, I'm very happy to answer any questions.

SIRANUSH VARDANYAN: Thank you, Kim. I think was a wonderful presentation/deep-dive into what IANA is doing and how important it is. We have already a line of questions and hands here. I will start with the questions which were posted in Zoom. So Sandra Davey always wondered, what happened to IPv5?

KIM DAVIES: I think that's a good question. There is an IPv5. Incidentally, one of the other registries that IANA is IP version numbers. There is a table that IANA administers and there's an IPv1, IPv2, IPv3, et cetera. So anytime anyone wants to experimentally create a new Internet protocol (a fundamental one), they would come to IANA for assignment of a number. It just so happens that it was only v4 and v6 that were the ones stuck, that got adoption. So there were experimental versions, other versions, that fizzled on the vine that were still allocated a unique number, but they never took off.

SIRANUSH VARDANYAN: Thank you. I will give the floor to our beautiful NextGenner, Zoe. Please, go ahead.

ZOE FILOMENO: Hello, everyone. My name is Zoe Filomeno. I'm a NextGenner here at ICANN. I would like to, since you said that IANA likes to look forward, see if IANA has ever thought about brain-computer interfaces and IP, since the end user is not dealing with anyone. Has there been talks about it and the future related to these protocols? Thank you.

KIM DAVIES: I'm pretty sure we have not discussed it. I think IANA is typically at the tail end of innovation and invention. So IANA wouldn't necessarily be the forum where new applications of technology necessarily start being discussed. I mentioned the IETF earlier. When it comes to network protocols (things that engineers would say go over wire), they would typically be standardized in the IETF, where different people come together and come up with a common standard to do something, and as part of that process we would start to be involved in terms of allocating assignments. Yeah, I'm not aware of any work in that space, but it wouldn't surprise me if it's ongoing out there. And at the time where work like that, innovative new work, requires standardization on the Internet itself, undoubtedly, I know we'll be involved at some later stage.

[PAUL HOFFMAN]: Kim, can I hop in here just to follow up? Wearing my IETF hat (I mean, Kim's IETF hat is as big as mine), the IETF itself also generally stays the

heck away from user interface. So the IETF standardizes new applications. You've seen HTTP. That's an application that, when it got standardized, entered into a few of the IANA registries. But I think what you're speaking about is, you know, where are the buttons? What do they say? And such like that. There isn't any standards organization that really does that. The one that people would expect to is the IETF, and then, as Kim said, it would filter down to him. But the IETF has spent the last 30 years just not doing that at all and letting people innovate on their own. Once they come up with protocols, though, that will need identifiers, those get standardized in the IETF and come to Kim.

SIRANUSH VARDANYAN: Thank you. We have hand raised Ugo Akiri in Zoom. If you are here, please unmute yourself and ask your question.

UGO AKIRI: Thank you for the opportunity. I'm a research student, and I've always wondered, how is the user protected in the making of policies? Because it seems to be it's a lost world. So that's one question.

I have another question. I mean, I've heard during the course of this meeting that ICANN is maybe cooperating with or using blockchain technology in some aspects. So my question is, how is ICANN planning to make the domain name system much more secure? Maybe using blockchain?

KIM DAVIES: I think, obviously it's a very open-ended question. I think that, generally speaking, on how is policy development protecting the end user, ICANN, as a community, as a structure, tries to have models that are representative and allow for contributions and participation by all manner of people, so that the policies that derive out of this organization do that. I think that when it comes to my narrow remit, (the IANA function specifically), I think through striving to ensure unique identifiers are globally consistent is at its heart a consumer protection measure and does help protect end users, because if those unique identifiers weren't consistently applied, weren't consistently used … Again, to use the example at the start of the presentation, if a domain name meant something entirely different depending on which network you're connected to, there are huge implications for confidence in the Internet, end user protection, confusability, etc., etc. So I think our core mission is making sure everything works the same no matter where you are, and that implicitly helps with end-user protection.

I'll hand it over to Paul on the other.

PAUL HOFFMAN: So I will riff off of what Kim just said. No, ICANN is not doing anything with blockchain naming for the exact reason that Kim just said, which is that ICANN's mission is to maintain the unique identifiers. The whole idea behind the different blockchain naming is to break the unique identifier system, that each one gets to have their own. That has advantages for them, mostly financial, not advantages for the user and certainly not advantages for security. So ICANN has no policies against

the use of the blockchain names. Anyone can have any naming system they want. Almost every country has postal codes, which if you think of it, are sort of like names for areas and such. ICANN is not in charge of all the postal codes. We are also not in charge of anyone else's naming system, particularly, going to the main thrust of your question, not the ones that are going to hurt the end-user security.

SIRANUSH VARDANYAN: Thank you—

UGO AKIRI: I think I need to clarify something.

SIRANUSH VARDANYAN: Okay, please go ahead.

UGO AKIRI: So I'm not talking about the unique identifiers in terms of protecting the user. I'm talking about policies made by maybe IEEE. How is the user protected? That's my concern.

KIM DAVIES: I suspect this question is a little out of our wheelhouse for this session. Happy to talk to you offline and perhaps point you to a resource that might be useful. But certainly with IEEE, I don't have a lot of familiarity with them.

SIRANUSH VARDANYAN: Thank you. Nojus, go ahead.

NOJUS SAAD: Thank you, Siranush. Good morning, or should I say afternoon, to an incredible panel over here. This is Nujos Saad from the ICANN 77 Fellowship, for the record. So I'm really curious to know, how does OCTO regulate and monitor DNS abuses that occur on the internationalized domain names, especially in cases of those that are run by non-contracted registrars? Do you think that it is the primary responsibility of national governments? And if so, how does ICANN influence or work with these national authorities to regulate them?

And a very quick question to the IANA team over here. In terms of a technical perspective, how does VPNs, or virtual private networks, when they're used by cyber criminals, affect the cybersecurity measures of law enforcement or ICANN in general or registrars? Thank you.

SIRANUSH VARDANYAN: David, we'll start from your team, or...

KIM DAVIES: I suspect we might want to flip the assignments, because I can talk about IDNs and ccTLDs, and you might want to talk about VPNs.

DAVID HUBERMAN: Right. So why don't you do IDNs first, even though, just to be clear, I was doing IDNs earlier. I'm actually one of the co-authors of the original IDN specs, but Kim's the one who has to drill down on them.

KIM DAVIES: So I think if I understood your point correctly, IDNs, particularly for non-contracted parties, in this case, ccTLD operators, not necessarily all of them are following best practices. And best practices here (and this speaks to something in my slides called label generation rule sets) for registry operators is to adopt policies that minimize the risk of confusability between different registrations. If you just allow any possible IDN to be registered within a registry, you will undoubtedly run into problems. And so ICANN has a lot of initiatives (one is LGRs and reference LGRs and the like) to establish best practices for how IDNs are registered to minimize that risk.

Unfortunately … I'll take that back. ccTLDs have a very different governance model for gTLDs. ccTLDs are vested within their countries. ICANN has a very limited role in how ccTLDs operated. Now, this is a positive, because every country is in charge of how they run their ccTLD. They can find locally appropriate solutions and meet the needs of local stakeholders. That's a positive. But it doesn't mean that there's no unilateral application of best practices and norms and the global level down to the country level. So it does mean that we are reliant on ccTLD managers being responsible and following best practice, and we can't enforce it. We try to use our technical engagement and other things at our disposal to promote best practices, to explain why it's important. If

they need resources to better understand how to solve the problems, we're happy to provide them, but we can't force best practice.

So I think that's the status quo right now. The good news is the vast majority of CC utilities do follow best practice. I think the number of TLDs where this is a problem is relatively small, but it's not zero.

DAVID HUBERMAN: And then to quickly answer the VPN question, I also, before I was at ICANN, was the director of the VPN consortium. So VPNs, in fact, are a method to hide one set of addresses behind another. So the VPN has an external address, which is visible. It's a normal IPv4 or IPv6 address. Again, that's not an IANA thing, because IANA has already allocated those out. And then they use private addresses on the inside. We can talk offline about that, but that is not within ICANN's purview at all, because ICANN is about the address allocation, as Kim was saying, not about what you do with your address, no matter how bad it is, much less how you hide other bad people who are doing bad things. That's not part of ICANN's mission at all. We watch it. We will help other people with it. But we have no control over that at all. And Kim doesn't even get to help on that one.

SIRANUSH VARDANYAN: Thank you. We have only three minutes to go. We have one question online and one question here. So the question is from Dr. Gopal from Chennai, Anna University, India. And the question is, is there any randomizer used in the generation of the IP addresses?

KIM DAVIES: I think the short answer is no. When we allocate IP addresses, it's actually more about routing efficiency, and I think that's true at the regional Internet registry level as well. Generally speaking, to make the Internet function more efficiently, you actually want to group IP addresses together that are operated by a common network operator. This simplifies routing tables and the like[.] And then on the allocations we make from IANA to the RIRs, generally they want to be sequential so that it's much easier for RIRs to allocate them in a sensible way. So in a sense, it's the opposite. We do not seek randomization. We seek sequential allocation.

SIRANUSH VARDANYAN: Thank you very much. And the last question. Mouloud, please go ahead.

MOULOUD KHELIF: Good morning, everyone. I'm Mouloud Khelif, ICANN 75 Fellow, for the record. Just to go back to your presentation on the domain namespace, here's maybe a policy governance question. You mentioned that you maintain direct interaction with TLD managers, but only due diligence processes for ccTLD. Can you elaborate on the nature of the relationship, the difference of the relationship, and maybe explain why it is the case historically? Thank you.

KIM DAVIES: Yes, certainly. That's a good observation, and our relationship is very different, so you're correct in making that assessment. So, as I mentioned in response to one of the previous questions, ccTLDs are generally locally decided and empowered, so our responsibility is to

effectively do due diligence to make sure that any ccTLD manager we appoint reflects the local wishes within the country. And we have a bunch of assessment criteria around that. It's a very deliberative and detailed analysis that we perform, but it's around that fundamental notion that we're not picking who runs a ccTLD manager. We're seeking to satisfy ourselves that a process has been conducted within the country in an appropriate manner, consistent with local laws, to appoint a ccTLD manager. They come to us and say, "We've identified this party to run our ccTLD." And we want to confirm that that's a result of a proper decision that's made in the country, and also to just validate some basic operational competency issues. We need to make sure that any TLD manager is sufficiently capable of running a TLD. TLDs are critical infrastructure. It's really essential that they're run in a very resilient and appropriate manner. So we do do some tests on that area as well. But the fundamental test for ccTLDs is, is this what was decided in the country?

gTLDs are a completely different case. You've probably heard about the New gTLD Program. If you want a gTLD, you apply to ICANN at large through an application round. You need to comply with eligibility criteria. ICANN does an assessment outside of IANA (IANA is not involved in that assessment) to make sure that your application complies with the application round's terms and conditions. You pay money. You get a license to run a TLD. You enter into contractual negotiations with ICANN to give you permission to run a TLD. Only once you've gotten to that point where you're in a contractual relationship with ICANN to allow you to run a gTLD would you then come to IANA. And then our job is not to vet whether you're allowed to run that gTLD, because that's

already been decided. Instead, we're there to establish an operational relationship. So we identify who is in the operations team of the gTLD that should be authorized to run this TLD, and then we establish the day-to-day points of contact that allow us to do our work to make sure the TLD actually works. So that, in a nutshell, is the distinction.

SIRANUSH VARDANYAN: Thank you very much, and I would like really to thank my colleagues from OCTO team and from IANA team to be here for these really very informative presentations and interaction with the ICANN 77 Fellows and NextGenners and the community.

And yes, I agree. A round of applause is for you. My huge appreciation to our interpreters and our tech support and my colleague Deborah for helping me in this session.

And with that, the meeting is adjourned. Thank you very much.

[END OF TRANSCRIPTION]