ICANN84 | AGM – Get to Know ICANN Community: ccNSO and RSSAC
Sunday, October 26, 2025 – 13:15 to 14:30 IST

SIRANUSH VARDANYAN: Thank you. Welcome to the continuation of Get to Know ICANN Community series of sessions, and in today's session, we will learn about two new communities, but before I start the introduction of those communities, let me give the floor to my colleague, Fernanda, to read the script. Fernanda.

FERNANDA IUNES: Hi, everyone, welcome back from lunch. We are in the Get to Know the ICANN Community - ccNSO and RSSAC session. My name is Fernanda Iunes and I'm the participation manager for the session. Please note that it's being recorded and is governed by the ICANN Community Participant Code of Conduct, ICANN Expected Standards of Behavior and ICANN Community Anti-Harassment Policy. The tongue twister.

Please observe the following guidelines to participate in the session. I'll also post them in the chat for your reference. During the session questions will only be read aloud if submitted within the Q&A pod. Interpretation for the session will include English, French, and Spanish. If you'd like to speak during the 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

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file but should not be treated as an authoritative record.

microphone to speak. Only questions posted in the Q&A pod will be read aloud during the session as time permits and when directed by the chair of the session. Please state your name for the record, the language you will speak if speaking a language other than English, and speak clearly and at a moderate pace. And with that, Siranush, back to you.

SIRANUSH VARDANYAN Thank you, Fernanda. Okay, let's start learning another acronym, ccNSO, which is Country Code Names Supporting Organization. And it's my immense pleasure to give the floor to the Chair of ccNSO, Alejandra Reynoso, who is coming from the fellowship program, actually, like me. We were fellows during our old times. And vice chair, Biyi, they both are here to introduce the community, ccNSO community, and also to respond to some questions which we will have for them. And without further ado, Alejandra, the floor is yours.

ALEJANDRA REYNOSO Thank you very much, Siranush. Hello, hello. Can you hear me? Okay. Shall I speak a little louder? Okay.

SIRANUSH VARDANYAN Can you be a bit closer to the mic?

ALEJANDRA REYNOSO

Okay, here I am. I just wanted to say, once a fellow, always a fellow. I was sitting where you are sitting now. A few years back, I'm afraid to say how many because then I feel too old. But I can just tell you that 10 years ago, I was here in Dublin for my ICANN Leadership Program training when I was joining the ccNSO Council. And 10 years from then, here I am as chair.

So, you are here to learn everything about what ICANN is doing, but also see yourselves in the future where you could actually be at. So, not necessarily need to go as far as the chairing community. Of course, you are always welcome to do so. But find your place. And if anything, I can help with any questions. I'm here for that. And I would like to hand the floor to Biyi Oladipo, who will give the presentation for the ccNSO. Thank you for inviting us.

BIYI OLADIPO

Okay. Thank you, Alejandra. And thank you, Siranush, for the invitation to share what the ccNSO is with the fellow community. Good afternoon. Yeah, it's good afternoon, yeah? So, good afternoon, everyone. And for those who are joining online, good day wherever you are. My name is Biyi Oladipo. I'm vice chair, one of the vice chairs of the ccNSO. And I'm going to spend a few minutes just telling you what the ccNSO is about.

If you look at the screen that you have, it just talks about the spread of the ccNSO, which is the ccTLD community of ICANN. I always like to start my presentation by asking how many of us in this room know our country code managers? Let me see your hands. If you

know who your country code manager is. Okay, okay. So, that's very good.

Then I'm going to ask a follow-up question. And the follow-up question is, how many of you have actually approached or contacted your ccTLD manager and have a relationship with them? Okay. So, we have fewer hands. Okay. That shows how we engage with the ccTLD community in our various environments. And for me, it's usually very interesting when I ask the first question and I see more hands. And then when I ask the second one, fewer hands. And that's because most of us don't really have this affiliation with our ccTLD managers. And we need to understand that the ccTLD is the identity of your region, your country, your territory on the internet. And you need to have a good relationship with them.

Okay. So, this is the spread. We're part of the ICANN multi-stakeholder structure. And we represent the different countries that we all come from on the internet. And this just shows you the spread. Next slide, please. Okay.

Now, for you to understand the power of the ccTLDs that we have, six out of the top 10 TLDs in the world, that's nearly 40% of its registrations, are ccTLDs. Of all the names in the world, the top 10, six of them are ccTLDs. And that shows you the power of the ccTLD community as far as the whole ecosystem is concerned. Next slide.

In terms of uniqueness, we represent different regions. We have different governance models in terms of every ccTLD runs itself is independent of the other one. However, the ccNSO provides us an

umbrella for all of us to come together and have conversations that will benefit us and learn from one another. And if you can see this, the registry models are different, the governance models are different, of course, as the number of domains that we've registered.

Now, one thing that is very interesting to always say is that, yes, we make policies in the ccNSO, but we don't make country or ccTLD specific policies. Those policies that govern the registration of names and how the ccTLD is run, all of that are local or domestic policies as far as the individual ccTLDs are concerned. So we don't interfere with what happens in specific ccTLDs. So your governance models and everything that you do is your own, but general policies that guide how things are done, things like IDN, how do you bring IDNs into the ccNSO and all of that are our general policies. And those are things that we make for ourselves. Next slide.

This is a quote from one of our members who said, our built-in geographical diversity lets the ccNSO bring a unique lens to ICANN-wide discussions. And like I said earlier, we are across the five ICANN regions. And because of that, we bring in our uniqueness into the whole conversation as far as the ICANN ecosystem is concerned. Next slide.

So where do we fit within the ICANN structures? I'm sure this session, the Get to Know ICANN session tells you who is who in terms of the-- I love our community. We're very acronym-centric.

So different acronyms come, like I was having a session in Nigeria last week, and I said that our community is an acronym-centric community and some of the things that we say mean some other things in other places.

So for example, if we want to develop policies, we say PDP. But in Nigeria, if you say PDP, you're actually referring to a political party. So, but I love it because it gets to make me feel like, yes, independent, yes, but it gives me this-- I don't know what it's called in other parts of the world in Nigeria. It's bougie, you know? Yeah. I know some things around acronyms. So you have SO/ACs and the ccNSO is one of the SOs, that's the supporting organizations in the ICANN ecosystem.

ccNSO occupies the seat 11 and the seat 12 on the ICANN board. So that just gives us, I think, that's called the Empowered Community. Yeah.

ALEJANDRA REYNOSO The Empowered Community are five organizations. The ASO, that's the numbers, the ccNSO, us, GNSO, generics, GAC, and ALAC.

BIYI OLADIPO Okay. So we occupy the seat 11, which is occupied by Patricio Poblete from the Chilean Registry, and then also board seat 12 by Byron Holland, who is from the .CA, that's the Canadian Registry. Next slide.

So this is what we look like in terms of membership. We have 180 ccNSO members to date, and out of that, we have three IDN ccTLD managers who are part of us. The membership of the ccNSO is open and available to all ccTLD managers, and I underline all ccTLD managers, but not all ccTLD managers are part of the ccNSO as of today, and that's because for you to be a member, you need to apply, and the application will be processed, and then you're admitted into the ccNSO, but it's open.

You don't have to pay, but just make sure that you fulfill the requirement, which is the joining requirement, which is for you to apply. These are the number of people that we have from Africa. We have 40 members, Asia-Australia-Pacific, we have 58, Europe we have 47, Latin America 29, and North America 6. Next slide.

Okay, now in terms of what we do as the ccNSO, we undertake policy and policy-related work, and remember I mentioned earlier that we don't make ccTLD-specific or country-specific policies. We make policies that are general for all of us, and then we cooperate and learn both technically and commercially, things that have to do with domain name abuse, internet governance, best practices, and all of that.

We come together, we learn from one another, and when we were starting off at the .NG, when we're starting off, we actually gleaned a lot from the community on the things that they had done and had benefited them, and that helped us to start off and start on a good footing. I don't have to make the mistake that other people made.

I couldn't learn from them and then be able to do my own stuff differently. Okay, and then we engage with other stakeholders within the ICANN ecosystem.

So, we have regular meetings, well, meetings annually. At different meetings, we engage with the GAC, we engage with the GNSO, ALAC, and all of that, and we have liaisons to all the other SOs and ACs. Next slide.

So, this is what the council looks like as that today, I think, is only one person that is missing here. I think it's only Nick that is no longer part of the council here. So, who are we? We have 15 elected council members from the five different regions, and then we have three NomCom appointed councilors in the ccNSO, and what do we do? We administer and coordinate the affairs of the ccNSO, and then we lead the development of policy recommendations.

And also, okay, so we have Alejandra as the council chair, and we have Jordan Carter from .AU and myself from .NG as the vice chairs. This is what we look like. So, we have ladies who are very beautiful and men who are very handsome.

All right, so some challenges and some of the things that we actually take interest in. In terms of resilience, we ask ourselves the questions, how can ccTLDs remain stable in times of disruptions that we have? We have certain organs that help us look at those things. So, like the TLD ops helps us with resilience. Domain name abuse, we have the DNS Abuse standing committee that helps us

with that and in asking the questions and answering how do we protect our users and how do we trust the domains?

Internet governance, we have the IGLC, which is the Internet Governance Liaison Committee that helps us with that. Operational excellence. So, we also ask ourselves things around how do we promote security and stability. Universal Acceptance, we have the Universal Acceptance Standing Committee that helps us look at those things. And then we have the SOPC that helps us look at how ICANN can remain transparent and how to look at things around the budgeting and the operational efficiency issues within ICANN. Next slide.

Okay, so this is just what I've explained earlier. This is where things happen. The technical committee actually helps us with the tech day and I would like to invite every one of us in this room who is interested to join the tech day sessions tomorrow. We're going to have about four sessions, holding tomorrow. We're welcome to join us. Those who are interested in the issues around operational stability, technical issues within the DNS operating system, you are welcome to join us. Okay, next slide.

Yeah, Sean Copeland actually is the chair of our Guidelines Reviews Committee of the ccTLD, is one of the council members from the North American region. And he says if you aim to be a ccNSO council member one day, join the ccNSO Guidelines Reviews Committee. You will learn the ins and outs of the ccNSO. So this is

a direct, is actually a direct invitation and direct marketing for the GRC.

Okay, all right. So you can start small, think global. In terms of your participation within the ICANN ecosystem, you need to think around it, attend the meetings. You're welcome. The ccNSO sessions start on Tuesday and we have Tuesday and Wednesday as the ccNSO members' days. You're welcome to join any of the sessions and it would help you to at least understand better how ccTLDs, one place within the ICANN ecosystem and then two, the specific things that we do.

And then we have, you can subscribe to the newsletters, you can connect with your local ccTLD managers and they would help you to understand the processes better. And like Alejandra has told you earlier, your participation within the ICANN ecosystem, you are the only one who can determine the limits that you will go. Nobody can do that for you. You are the only one who can determine the limits. I used to come just to understand or just look at what's going on. For me, it was another thing, another place to go to and all of that.

As soon as I started getting involved, joining communities, I began to see how useful I could be within what was happening. And a few years down the line, I'm vice chair and I don't know where I will go next. But it all depends on you. You are the only one who can determine exactly how far you can go. She started a few years ago as a fellow like those of you in this room. I never had the privilege

of being a fellow. I really wish I could when I see the camaraderie and the way the fellows all intermingle.

SIRANUSH VARDANYAN You are our honorary member.

BIYI OLADIPO So now I'm an honorary member. I think we should have an induction ceremony here. So that's it. You can follow our work, our ongoing work on IDNs, governance, technical operations and then how do we run and guide ourselves within the ccNSO. Okay. So do you manage a ccTLD? If you do, please join the ccNSO and become a member. All right. This is from Alejandra. Maybe I will allow her to read it herself since she's here.

ALEJANDRA REYNOSO It says your dedication makes the ccNSO stronger. It's an honor to work with you whether you join in person, remotely, actively or as an observer. Actually, it's very important the participation of everyone in the ccNSO. We know it's voluntary. We know you do it because you want to and because you're looking for a better world in the internet ecosystem and I'm grateful for that. So, thank you.

BIYI OLADIPO All right. So as an observer, you're welcome to come participate with us and see exactly what we're doing. This is what our week looks like at ICANN84. So Monday we have tech day. Tuesday are

members' meetings. I would especially like you to, if you can join all, that would be beautiful, but if you can, please join the ccNSO news session. That's very topical and people always come from different parts of the world to come and show us exactly what they're doing and things that they have done that is working. Okay. And ongoing works as it's going. It's always very interesting.

On Wednesday morning, we have a newcomers' informal meetup session. So for those who are coming to ccNSO meetings for the first time, Wednesday morning between nine and 10, you can come join us. And we have informal sessions with a few people from the council and the members of the onboarding and mentoring committee of ccNSO to actually meet you and get to see how we provide guidance. All right. So thank you very much for having us. I wish you good ICANN84.

SIRANUSH VARDANYAN Thank you, Biyi, and thank you, Alejandra. Thank you very much. And we can take one more, one or two questions for our presenters. Sameer. Okay, go ahead.

SAMEER GAHLOT My name is Sameer Gahlot. I'm a current fellow. I have one question. I can't see you guys because the mic is this side. Sure. So the question is, though I'm part of a ccNSO back in India with .in and the work which ccNSO is doing when it comes to the heat map, which is very like recognized by the global team also. I just want to

understand what is the current status of the heat map because I can see it was last developed in 2023 and not yet updated. So if you could like shed some light on that.

ALEJANDRA REYNOSO

Thank you so much for the question. It's just as a bit of a background for everyone. The heat map that's been mentioned, it's a tool that it's used in our Internet Governance Liaison committee and the tool serves the purpose of understanding what are the hot topics on each region regarding internet governance. And it's been updated frequently. I know it's been updated after 2023.

These specifically yesterday, now I need to know where I am in time. The IGLC, that's the Internet Governance Liaison Committee, had its meeting and they said that in their next call, they will address the heat map to have the latest, latest version of it. But usually the topics that are mentioned there are not too far from what we see in ICANN. I remember from the last time I joined them, they were talking about DNS Abuse, Artificial Intelligence and Governance overall because of the WSIS+20, that its upcoming soon. So I would invite you to check on the website at least in a month. I think it's when their meeting where they will update it for the latest of it. Thank you.

SAMEER GAHLOT Thank you, Alejandra. And the reason for pointing out that is because it's very helpful. So it could guide us in which direction should we move on when it comes to what are the priorities for the ccNSO and whether that aligns with the GNSO or not. Thank you.

ALEJANDRA REYNOSO Thank you so much.

SIRANUSH VARDANYAN Thanks for this feedback. It's really important. Yes, Ernest, go ahead.

ENERST MAFUTA KATOKA Hello, everyone.

SIRANUSH VARDANYAN Enerst, your mic.

ENERST MAFUTA KATOKA Awesome. I guess I can be heard now. My name is Enerst Mafuta from Zambia. I'm a returning fellow. I've got a question for Mr. Biyi with regard to policy. He said that the ccNSO makes policy in general as a whole. But as you know, that the international arena is just a starting point. But it's up to each country to take what's relevant for them. So my question is that do you have any performance indicator to measure the type of policies that you develop if what countries are able to take or not?

And secondly, do you have a mechanism also to reach out to other ccNSO members who are not active? For example, in my country, I've never seen the manager here. I don't even know he or she. So I wanted to know what mechanism you have to follow up on those who are not members of the ccNSO because I believe it's very important in that regard.

BIYI OLADIPO

Okay. So I will answer the second first, and then take the first one. I'll answer the second first and then have Alejandro also say some things about that. Now, in terms of reaching out to people who are not active, there are a lot of outreaches that we have, especially through the regional organizations.

So we have AFTLD for Africa. We have APTLD for the Asia-Pacific region. We have LACTLD for Latin America and the Caribbean. And then you have CENTR for Europe. And each of these regional organizations are actually active parts of, in actual fact, the managers are observers on the ccNSO council. So they know exactly what is going on. And from time to time within, at least I can say for the African region, I have constant discussions and we have constant outreaches with AFTLD in inviting African ccTLDs to participate, not even just to join, but to participate in the ccNSO.

The Afri-ICANN is having a session in Accra in early December, and I'm actually going to be there to still talk about what the ccNSO is doing and outcomes from, we have a readout from this meeting at that session. So we have constant discussions and constant

interactions with the regional organization. I think it's actually better like that because the regional organizations are closer to the ccTLDs than the ccNSO can be.

ALEJANDRA REYNOSO Building on that, the ccNSO also has its own Continuous Improvement Program. And what we are developing very recently, it's a series of going ourselves to the ccTLDs instead of inviting them to come to us, we are reaching out to them. So this year in particular, we did two activities. We started in January with the Caribbean region and we approached them and let them know what we do or work and invited them to join us.

And early October, we did another one with the regional organization of LACTLD. That's the Latin American and the Caribbean organization to have a broader reach in that sense. And we are planning on now moving to the rest of the regions with the same program. So that's on the outreach. Shall I respond to the first question as well? Okay. So the ccNSO, it does develop policy, but it has a very, very limited remit on what policies it can develop. And these policies are addressed to IANA as to the well name service provider. And it's regarding on how to deal with ccTLDs at the root level, because they are the managers of the TLDs.

So those are the policies we develop. So we have one for retirement of ccTLDs, for example, if the ccTLDs come from the ISO 3166, that's where they are listed. And when a territory or a country cease to exist or changes in that list, then it needs to be retired. So

that is the type of policy we actually develop. We do not develop policies for ccTLDs themselves, for them to implement or to do. What we do provide in the ccNSO, it's a platform to share knowledge and to share good practices that others are doing.

So if anyone has done something that they feel like sharing to others then other ccTLDs from different regions can say, oh, that idea is helpful and they adapt it to their local needs because no ccTLD it's the same even though they do the same they do not operate the same so I hope that clears it up. Thank you.

BIYI OLADIPO Just to add a little to that is that, don't forget that the cc names are country codes and every country have what they can take and what they can't take. So, we can't make a general policy to enforce within your environment. You, will need to look at your environment and know what works and what can help to propagate and promote your ccTLD. All right.

SIRANUSH VARDANYAN Thank you. Thank you, Biyi. Thank you Alejandra and appreciate your time being with us to explain what ccNSO does and for those who are interested, I would encourage you to follow tech day tomorrow and participate in some of the ccNSO sessions. Thank you very much. And with that, we will move to our second community for this session to learn. This is RSSAC and we heard a

lot this abbreviation. So, this is Root Server System Advisory Committee in ICANN.

It's my pleasure to introduce RSSAC chair, Jeff Osborne, who always is happy to come and accept the invitation to talk to newcomers and to fellows to Next-Geners. It's also my pleasure to see the group behind you and Hans seeing you here. It's pleasure, head of RIPE is here with us.

HANS-PETTER HOLEN Well actually, I'm here as the next RSSAC vice chair so that's why I'm here but you're right I am the CEO and current chair of the NRO EC. Excellent.

SIRANUSH VARDANYAN I didn't know that you are coming as vice chair for RSSAC, but I knew that you are RIPE NCC, head over RIPE NCC. So, it's pleasure to see you here and Jeff, we want to know how it's everything works and what is RSSAC doing and how we can join there. Thank you.

JEFF OSBORN Thank you, Siranush. Thank you, everyone, for your time. I appreciate it. And hopefully, we can learn a little bit more. The primary function we have involves the DNS. And if you've spent any time at ICANN, DNS is important. So if I'm going over something you already knew, I apologize. But I guarantee it's impossible to know too much about DNS if you want to be involved with ICANN.

So in three parts, I'd like to explain what the RSSAC, the Root Server System Advisory Committee, does and its place in ICANN, and then a little bit about how we kind of look at DNS a little differently from some of the other folks you might look at. And then I'm going to try to introduce at a fairly basic level how the DNS root server system and DNS actually operate. So pretty ambitious for half an hour. Let's see if we can pull this off.

The Root Server System Advisory Committee, RSSAC, has a fairly simple scope. We basically advise the ICANN community and the Board on matters involving the operation, administration, security, and integrity of the internet's root server system. Okay, that's pretty much the remit. And then there's a bunch of detail, which if you really want to know, you can go in and see the bylaws.

We are part of the ICANN organization. We're an advisory committee. This really is just saying what I said in longer language. The RSSAC is made up of the 12 groups that operate root servers. There are 12 organizations. Each of them has a primary member and an alternate. So it's a total of 24 people. One vote per operator, so 12 votes. There are liaisons that come in from other organizations, and I'm sure you recognize a lot of these acronyms. I like saying them as words instead of acronyms because it just makes your head spin. The Internet Assigned Numbers Authority, of course, is the group that puts the root zone file together, which is the for all of the country codes and other TLDs.

The Root Zone Maintainer cryptographically signs that so that it's safe and secure and everybody on earth gets the same version. The Internet Architecture Board is a part of the Internet Engineering Task Force and is involved with the architecture of the internet. And SSAC is our sibling organization at ICANN that works mostly in the security arena.

We then send liaisons to the ICANN Board, the Customer Standing Committee, the Root Zone Evolution Review Committee. That's one of my favorite acronyms to pronounce, RSSAC. Someday I'll have a dog and name it RSSAC because that's just a great name.

The ICANN Nominating Committee, which of course picks a lot of people on the boards and in other positions, and then on an ad hoc basis when there are work groups that it seems necessary to be involved with, we'll send a liaison. Additionally, there is a--

SIRANUSH VARDANYAN Fellowship mentor and selection committee member for the fellowship program.

JEFF OSBORN Yes, absolutely yes we ran out of space you know this is really dense. I put these slides together with my lawyer who's actually right here and a dear old friend and if he had his way there would be three times as many words. So I apologize we skipped a couple of things it doesn't mean they aren't important.

SIRANUSH VARDANYAN But we know that.

JEFF OSBORN We know this. The RSSAC caucus is a larger group there are about a hundred experts on the subject matter. We take applications for it. I will warn you this is not a learning organization it is a working organization. This is for people who have basically built and designed DNS systems for their nation or large universities or things like that. So if you have that kind of hands-on DNS experience, we'd love to hear from you but I'll warn you off that it's a bad place to learn. People get in over their head very, very quickly.

RSSAC basically is a forum for discussion of how we operate. It's entirely structured, open, transparent. You can attend the meetings if you'd like. They're pretty geeky but you can. We get feedback on policies and plans entirely concerning the root server system and that's it. It's really our one thing. Basically what we end up doing is restating the long established normal framework of how the root server system operates.

This has been around for decades. It works really, really well and so we don't mess with it very much. Not in scope is we are not an enforcement or supervisory body. We are not a representative for the root server operators and we are not a representative body for the root server system in total. We're simply an advisory

committee to the community and the Board. We put out documents if you're one of those people who wants to learn more there are lots of documents out there. ICANN keeps track of these.

They all start with RSSAC and then a number. I'm sure you've seen this system before with other groups. SSAC has a similar one. The history of the root server system. I've actually read it. It's 47 pages. You can go to sleep four or five times before you get to the end of it. So if you don't use NyQuil consider reading the history of the RSSAC or the root server system. System descriptions about measurements, governance, advice. All of these things are available if you want them and I'm certainly not going to read it all.

All right second. Who here feels like they could explain in a couple of sentences how the DNS works? I'm glad to see Hans-Petter had his hand up. Okay we have a couple. That's good. We feel like there's a misunderstanding and I really want to explain a little better and we'll try to walk through what the DNS is and how it works and explain where the error is. So this is primarily a document that was designed for some government people that we were speaking to because they're putting regulations and legislations together.

So if I'm addressing you as policy makers, I'm sure you're up to the task. You all could be policy makers someday or now and so that's the role we're talking to. Most people who have to make policies about things do not understand the root server system. Most people in this room don't understand the root server system. I

think most people at ICANN don't understand the root server system because it's complicated.

Unfortunately to make policy you need to understand it. So we have to explain theoretically and operationally just how this works. This is the normal way to explain the DNS. Basically you go to the root server and say how do I get to a Top Level Domain? We give you the list of Top Level Domains. Then you say, okay, .com how do I find out what the domain name in .com is and you go there and you get all of these answers until you find www your domain name dot whatever it is at the bottom of this and it all starts with the root.

That assumes a cold dead internet and this is the first time anybody has ever looked anything up. That hasn't been true for a very, very long time. This is an excellent version of slides from ICANN showing exactly how this works where technically it makes sense to say first resolver looking for names goes to the root server.

You can't quite read it that says root server on the far left. Then it goes to the recursive server in the middle and then it goes to the name server on the right and you would think from this that everything starts with a trip to the root server. But that's not true. This creates the false impression that root servers are somehow gatekeepers and it isn't true because the hierarchy in the technical terms is very different from hierarchy in political terms.

In political terms if there's a hierarchy there's a king or a president at the top. In hierarchy in technical terms it just means there's a structure of a thing and pieces are in different places. It doesn't

imply political power. What we're going to show is there is almost no impact from a root server having any piece of it fail. There are almost 2,000 instances of root servers out there in the world across the globe and there's no technological single point of failure.

Many, many of them use a thing called AnyCast so that if you call a particular address to find one of these and the one you were closest to has been blown up somehow or hit by a bus, the address will get you to the next one. It's like going to a town anywhere in America and say I want to go to Dunkin Donuts. If one of the Dunkin Donuts burned down it doesn't matter. You'll just go to the next Dunkin Donuts and there are 1,900 Dunkin Donuts in some small towns in America almost it seems like it. So you can never run out of them.

So an individual failure doesn't cause a failure anywhere in the system. Again because there are 12 operators and they're all different there's no institutional point of failure. The other thing to keep in mind is, we're handing out the addresses to something big like a domain name or a country code and those don't change very often.

So if you were using the same list of country codes like when you're dialing on the phone if you know the country code 44 is that England?

Okay, once you know that, you don't have to keep asking every time you dial your friend in England, hey, what's the country code? You know what the country code is, and so it's very similar. The

root server system as a whole literally hasn't been shut down in 40 years. It just works. So, here's the way it really happens.

The first thing that happens when you go out and look for something, hey, where is www.iwanttobuysomestuff.nl? The first thing that happens is the resolver goes, have I ever been asked this question before? And if the answer is yes, then it will simply tell you, I know that answer, I've asked for it before, here you go. So, in something like 93 or 94% of cases, more than 90 percent, there's no root server involved, there's no nothing involved.

The big 8.8.8.8 or 1.1.1.1 knows all of this stuff, and so in the vast majority of cases, you don't have to go off to the system, you don't have to go off to a root server. There are additional steps of look at what the domain name knows about things, look at what the TLD server looks for, and it turns out in about one out of five to ten thousand queries, somebody has to go to the root server system.

Now, think about the first thing we saw was the first thing you do is go to the root server and then and then and then and then you do these other things. That sets up the idea that well, if anything's broken, it's probably the root server. But it turns out that the only time you have to look up the root server addresses is if they've changed, which isn't very frequent, or if you literally don't know anything. And these things operate, there are about 500 trillion queries a day. That's a big number, 500 trillion queries a day.

So these resolvers get asked constantly and remember everything they've been asked, and so they just tell you what they already

know. Does that make sense? A couple of shaking heads, good, okay. All right, we got through those two parts.

Now, if you understand this next section, you will be very valuable in the job market because DNS engineering jobs pay well. I have a bunch of DNS engineers who work for me, I can tell you they make a lot of money.

So we're going to try to walk you through how this system actually works, okay. I'm going to introduce the concept of the root server system and operators. I'm going to do that by trying to explain the role and purpose of DNS. DNS is important, remember that because you're at ICANN. And then there are two pieces, the resolver and the authoritative server are the two main things that are out there answering your questions, okay.

And then we'll figure out when and why and how you would ever have to talk to the root server system and some common misunderstandings, all right. And then this is the things that you, leadership officials need to know because if you're not leaders yet, I suspect you all will be, that's what fellows do.

All right, DNS, in its simplest step, DNS uses human names to find computer addresses, okay? Human names to find computer addresses. You all know names like, I hate to admit how well I know, www.amazon.com, but the computers use addresses and you've all seen these four sets of numbers with the dots between them.

That's an IPV4 address, 18.239.62.181. Well, it turned out on the day I wrote this slide, I went and looked what is the IP address for www.amazon.com and it was 18.239.62.81. If I want to go out and buy a Hawaiian shirt from Amazon, I'm not going to remember 18.239.62.181. It's not going to happen, but it's really easy to remember words.

So, you could argue the whole point of the DNS is words are easier than numbers for humans to remember. Okay, that's it. Go get jobs. Good luck. Thanks. It was nice to talk to you. There's a little more to it than that. What's interesting is the numbers can change because computers are always changing.

I spilled coffee on my keyboard last week. I literally am on a different computer, but it doesn't matter. It's everybody's concern.

The words still get you here. The numbers can change and the names stay the same. Any device out there on the internet, computers, servers, smartphones, I have a dishwasher that has an internet connection. I have no idea why. It uses DNS to go talk to whoever it talks to. I have no idea who my dishwasher talks to, but you can see it. It shows up in the router traffic, so it's doing something. So, DNS questions ask about domain names and the answers are IP addresses. That's the takeaway.

Why use this? Well, it's human friendly. It's easier to remember, www.example.com, than a string of numbers. Service portability. Like I said, if you dump coffee on your server or if there's a fire or whatever, nobody cares. You can remember what the name was

and DNS will figure out where to take you. It's remarkable when you really get into this to recognize how huge this network is.

It is arguably the world's largest distributed database and it exists in arguably hundreds of millions of places, which is just kind of incredible. If you've ever worked with databases, back in the early days, I played with databases and was just stunned at how much you could do. And this one is distributed all across the earth and can be accessed successfully for decades without a failure in milliseconds. It's shocking that it works.

So, devices get addresses from resolvers. There are millions of resolvers out there. You probably recognize ones like Google runs one called 8.8.8.8. There's a group called 9 that shockingly uses 9.9.9.9. I think IBM is 1.1.1.1. Is that CloudFlare? CloudFlare. What's IBM? Anyway. So, those are resolver addresses and you've probably used them before. And that's where you get this information from the DNS to figure out where your computer needs to go in numbers because that's what it needs to know.

So, resolvers basically can find and read what I'm going to call internet phone books. And I'm showing my age, phone books were these things they used to have that had everybody's telephone number in them, but I realize we don't use them anymore. So, the phone books are what's called an authoritative server and the listings are what's called zone content. Okay?

So, everybody who runs a TLD has an authoritative server that tells you every domain name in that TLD. Get it? So, for instance, I'm

going to ask what is the number for Amazon? It's going to tell me the number. It happens in milliseconds and our best guess is it happens about 500 trillion times a day. I never get tired of that number because I can't believe how big it is. So, the resolver remembers addresses. It's called caching. It's just the memory. Every time it's gotten a number, it's going to remember that.

So, depending on how much information it needs, the resolver is either going to just look in its own memory, which is the usual case. If it doesn't know that, it might have to go and say, I need to ask the domain name what the thing to the left of the domain name is. If you know Amazon.com, but you don't know if it's mailed on Amazon.com or WW or whatever, you have to ask the domain name's authoritative server.

If you don't know that, you have to go all the way to the TLD. And if you don't know what the TLD is, the top level domain, you'd come to us, the root server system. And like I said, one out of five or 10,000 queries have to go that far.

My favorite part, too. This is going to walk through how this actually works. And it's a little complicated, but I beg you, for three or four minutes that this is going to take or five, try to pay attention because it's really kind of cool if you grasp it. It's like, oh my God, that's how DNS works.

In the simple, the gray box is the resolver. Okay. So that's a computer. We're going to have a series of flow charts where we're always going to start with the question, what is the IP address of

www.example.com? And depending on how much I know when I start, I'm going to take a different path. Okay. There is the first question. Do I know the answer? And then next to it is the memory. So as questions get answered, they're going to get dropped in memory. And then we're going to ask the question whether we know everything yet.

Ready? So what is the IP address for www.example.com? Do I know the answer? In the first case, yes, it's in my memory. I'm done. This is literally over 90% of the time. That's how simple it is. Where's the root server? Nowhere. Authoritative servers? Nowhere.

The resolvers do the bulk of the stuff and they know this stuff because once they remember it, they keep remembering it. So let's go to the slightly harder case where what is the IP address for example.com? Do I know the answer? No, I don't. Uh-oh. Do I know the IP address of the next level down, the authoritative server for example.com? Because I didn't know www.example.com, which is different from example.com, which is different from .com.

I ask, hey, what is the IP address for www? And I ask it of example.com's authoritative server. It knows the address. There it is. It goes to memory and we say, now I know that. Do I know the answer yet? Yes, I do because it's in my memory. Hooray. This is about 5% of the cases where you have to go that far.

So I guess you can see where this is coming, right? What if I want to know www.example.com, do I know it? No, I don't. Okay. Do I

know example.com? This is a very broken computer or a brand new one or something because it goes, hell, I only know where .com is. And so that's all right. Let's ask .com, hey, where's example.com? The authoritative server goes, here's the address. We put it in our memory and go, do we know the whole answer yet? No, because we just figured out where example.com is not www.

Okay. Now I go back and ask the authoritative server for example.com, what's the answer? It gives it, I put it in memory. Now do I know the answer? Yes, I do. Hooray. This is about 2% of the cases where you have to go this far. Now we'll get to the really crazy case. This is the one where somebody took the computer out of the box. This is a brand new resolver. It can't spell internet. It can't spell DNS. It doesn't know anything. But the first thing you put in a resolver when your company like mine makes DNS software is a list of the addresses of the root servers.

So when it wakes up, that's the one thing it knows. So what is the IP address for www.example.com? Do I know? Nope. Do I know example.com? Nope. Do I know .com? Nope. All I know is this list of root server addresses that I was given at the factory and I'm going to go figure some things out. So I send the query and say, hey, root server, what's the IP address for .com? And the root server system goes, here it is. We know right where that is. Put it in your memory. Now do you know the answer? No. Do I know example.com? No. But now I know .com.

So literally now I go to .com. I ask it, hey, where's example.com? It goes, here's example.com. Do I know the answer yet? No, because I didn't get www yet. Now can I have www please? I ask example.com. Yes. Here's the number. Put it in memory. And do I know the answer yet? Yay. I know the answer. This is what's so interesting about the DNS though. The bottom one there is one of 2000 root servers out there. The next one up is one of the ungodly number of servers that VeriSign has out there for .com. And then above that is the example company has its own database of all of them there.

So it's going to all these different places. It's taking 24 different steps. You don't have to do this all the time. This is one out of 5,000 or 10,000 things. And this is where I said if you're a database person, it really gets interesting. Holy mackerel. This stuff is everywhere. Here's the interesting thought too. There are about 350 million domain names. Okay? So it would be really easy to think, well, the root server only has 1500 TLDs to worry about. And .com only has 150 million addresses to think about. These numbers aren't that huge.

When you get to the next one up though, we use a piece of software called Bamboo for doing our HR functions and telling when people are out. The way they identify us is, my company is called ISC. So I go into isc.bamboo.com. So I ask them, how many companies do you have in there? And it's like 70 million.

So the number of these things as you go all the way up is just huge. They're distributed all over the world. Everybody who's responsible for them gets to change them as they want to on their own machine. And in milliseconds, we get these answers back. It's absolutely a miracle to me that this all works. But it does. And reliably.

So here's the ugliest slide in the whole deck, which is hard because there are no pictures in the whole thing, so it's all ugly. But there are about 350 million domain name zones out there. And those are all maintained by the domain name registrant. The registrant is the person who was contracted to have a domain name. The Amazon.com is the registrant of the letters A-M-A-Z-O-N.C-O-M.

The TLD zones, there are about 1,450 of them in total. ICANN is putting out new ones at some points. I'll have to update the slide. But that's how many there are now. And that is maintained by each of the TLD registries, whether it's a country code-based one or whether it's a company. There's only one root zone. And what it's got in it is the 1,450 addresses of all of those TLDs.

One. We don't make it. We don't write it. We don't do anything with it. IANA puts that together in a very sacred and protected way. And then it is cryptographically signed by the root zone maintainer and handed out. So nobody can hack it. Nobody can change it. You can tell if somebody tried to do something incorrect with it. And it is then offered to the world by the root server system.

By a show of hands, anybody feel like they learned anything in the last 20 minutes? Two of you. Excellent. Okay. I'll keep going. So in review, the root server holds a copy of the root zone. And that's basically the addresses of the 1,450 registry authoritative servers. And so that's for everything from .com to .nl to .jobs to .anything. And the TLD then knows the next address. So .com knows everything to the left of the .NL knows everything to the left of the .Jobs knows everything to the left of the dot.

Then once you're the domain registrant, you control everything farther to the left. Whether it's www or in the case I said, it's all of your customers if you choose to handle things that way. And the resolver finds and returns the answer to all of these and almost always out of its own memory.

I was going to put this on a T-shirt. That's why I put it on a big slide. But it wasn't as clever as I thought it was. In the millisecond world of resolvers, the queries to the root server system are really, really remarkably rare. And this is not obvious if you haven't really dug through the numbers. So I'm going to zip through a couple of these. Just to be really clear, we hand out addresses. These are basically the internet's version of a country code you would dial on a telephone.

We don't have content. We don't have email. We don't have anything like that. We have nothing to do with anything bad that's ever been scammed or sent to the wrong place. We literally just tell you how to talk to dot com, how to talk to dot NL, how to talk to dot

jobs. Really basic. No content. We're not a gatekeeper. We are not the first thing you have to do. We are not the thing that is holding you up. I guarantee even if we were slow, who would notice that one out of 5 to 10,000 queries took longer? Like, that's a lot of queries. The average person could spend not as much time as you think. Probably days or weeks doing that many. But you wouldn't know which one it was.

So traffic is almost always transmitted without the need for the root server system. This remarkably resilient thing has no single point of technological failure. It's on diverse hardware, diverse operating systems, diverse DNS platforms. My company makes DNS software that is probably the oldest one out there. It's called BIND. And I think we're still the most used out there. Millions and millions of users out there using it. But even we use other people's software, too, because we want diversity, because we don't want the thing to go down. No point of technological failure.

And with 12 entirely separate organizations, there's no single point of institutional failure. This is the point where after the difficulties that happened with the RIR in Africa, there are force majeure or court things that could go on that wouldn't affect the end user's use of the sim at all. You could literally drop any one of these operators off the edge of the earth tomorrow and nobody would notice. Which seems really cruel, but it's the case.

So the system has operated since the 80s. It has not had a blackout. It has not had a failure. We are one of the most favorite things for

DDoS attackers to try to attack, and it just doesn't work. It's an interesting detail how AnyCast can beat the heck out of DDoS. But that's a geekier discussion for another day. 30 years of success, 24 by 7 by 365.

So we don't decide what goes in the root zone. You get to decide what goes to the left of your domain name. The TLD gets to decide what goes to the left of the TLD. The IANA puts all this together, puts it into place, and it is cryptographically signed by the root zone maintainer. They hand it to us. It is locked up. It is solid. We have no ability to edit or change any of it, and we have no interest in doing it.

So we provide address information, not content. We don't decide what's in the zone. We are not gatekeepers. This is one of the most resilient systems, frankly, on the face of the earth. It's also a huge bargain. All of this is done voluntarily by these organizations. I run a non-profit. We spend millions of dollars a year operating our root servers because the internet deserves it. And that concludes our show.

SIRANUSH VARDANYAN Amazing. Thank you. Thank you very much, Jeff. And we have a couple of minutes for questions. Miracle, please.

MIRACLE ONYEBUCHIM Thank you very much, Jeff, for that very engaging conversation. I have a couple of questions. The root server system has been in operation for nearly 40 years.

SIRANUSH VARDANYAN We can't hear you. Can you-- yeah, can you put…?

MIRACLE ONYEBUCHIM Can you hear me now?

SIRANUSH VARDANYAN Yes.

MIRACLE ONYEBUCHIM Right. So I have a couple of questions. The root server system has been in operations for nearly 30, 40 years. And like you mentioned, it's never failed for one day. It's pretty, it looks like it's a stable, successful system. And I'm almost left to wonder what RSSAC has done to contribute to this. Because it looks like it's pretty much set in motion.

A lot of the issues, especially as it relates to end users, seem to be influenced by resolver software companies. And I know you mentioned BIND. And I think that they're one of the very important, and I think they're quite a couple, that it was actually influence caching behavior, query time, and all of that.

So I'm just curious, if less than 2% of queries, or 0.02% of queries actually get to the root servers. So it means that end user experience is influenced by things that don't have to do with the root server, right? So I'm curious, I know maybe the IETF probably has a say on that. But unlike root servers, where all 13 root servers, they work hand in hand, and they share information. I think that resolver companies are very independent. And I don't think there's any regulation.

Has RSSAC actually, I don't know, maybe it's not on that RSSAC. But I'm just curious to know, if the root server system is stable, maybe the concern should be on the things that actually affect end user experience, like caching behavior for resolver companies. Is that something that's perhaps under your remit? Or is that under the IETF? I'm just curious to know what that dynamic is. Because it just looks like the root server system is stable, and doesn't need any help. Okay.

JEFF OSBORN To be clear-- oh, is this working? Hello? To be clear, RSSAC is an advisory group. I'm the chair of the RSSAC. I happen to also be an operator of one of the root servers. My company does the same thing. And Rob and I are actually the two. I think I'm primary and he's secondary, or the other way around. So we operate the way we do. And every resolver on earth gets a list of the addresses of the root servers, okay? We have one address for all 300 or 400

servers that we operate. I don't even know what the number is that we personally, our company does.

So we don't tell people where to go. The protocol picks sort of randomly from a list which one to go to. And the resolver that you use is something you literally enter in your own computer. So ISPs often have their own that they'll preload or I couldn't even tell you whose resolver I am using right now. I switch them around because I find it interesting.

So it's really an individual choice. We don't really get to control who uses what. It's kind of remarkable how much it just works. It's like you said, it just kind of works. It really does just kind of work. Does that answer the question?

MIRACLE ONYEBUCHIM Yeah, I wanted to know. Thank you very much. I wanted to know if you also release any sort of policy advisory for how independent resolver companies should apply standards. I'm just curious to know since you also run one of the resolvers in BIND. So I wanted to know if there's been any of that.

JEFF OSBORN I happen to have an expert here.

HANS-PETTER HOLEN Yeah, so Hans-Petter Holen and I mentioned I was the chair of the RIPE NCC which is the secretary of the RIPE community which is all

internet providers in Europe. And typically they have met over years and developed best common operators practices. Some of that for DNS may be pretty old but if you look at the history of resolvers back in the really old day, your company would put a server in their office running a resolver because you wanted to minimize the traffic going over your very expensive internet connection.

Now over the years, you get the box from your ISP with the resolver config is included in. So you get the ISP's resolver. And then you go start to go down a political lane rather than a technical lane. And then up pops global resolvers that promise you some kind of privacy from your ISP or your government. So then you rely on large multinationals to provide you privacy that your government won't provide you or is it the other way around? I mean, this is, then you move into something that's not technically driven, but driven by other elements.

So there is a lot of different drivers on why you would use a resolver in a certain way and whether you would run one yourself. You can get a cheap computer and put in your own network and run your own resolver there and then make your own decisions on configurations. But that's for the few, right? That's not something everybody can do. So is there research and work in this? Yes, but not in RSSAC.

SIRANUSH VARDANYAN Thank you. I saw a couple of hands over here. Yes, let's start with Om.

OM PRAKASH SHARMA Hi, this is Om from Nepal with NixIn Program. I want to ask about the root zone key signing ceremony. So I see in the videos, a bunch of important people coming together where every second is logged in or who does what. So I think it will be interesting to hear from you. Thank you.

ALEJANDRA REYNOSO 89 Om, just a quick fun fact. I'm one of the safe openers on the East Coast for the KSK ceremony. I have the code for the safe.

JEFF OSBORN I was going to say that all happens before we get it. And I was going to look for an expert. And of course there's one. It really isn't part of our function. We get everything cryptographically signed. I mean, we really just provide carriage. We really just make it available to everybody on earth. That's our whole thing. By the way, like I said, nobody has to pay for root server access. It's all free at delivery.

And I can't remember the last time somebody asked us, hey, I want a root server in my town or company or village or whatever. And we didn't just do it. So it's really pretty much a freely available good. It just works. And they're just available. Did that help? I wish

I knew about the key signing. I don't. Thank you. I wish I knew about the key signing. I don't.

ALEJANDRA REYNOSO 89 I can get you in touch with some people that do know, so.

SIRANUSH VARDANYAN Thank you. Shashi, I know you have a question. Can you come close to the mic?

SHASHI RAJ Okay. I'm audible. Okay, I'm audible. Yeah. So exactly like, I'm Shashi, a NextGen program fellow. So exactly like you showed the light on, a different angle light on how DNS actually works, I want to understand a similar view on DNS cache poisoning. Can you just give an idea that how it technically works? Because in the view, I can see that if DNS resolvers get infected from APT groups, that can lead to a different cache poisoning attacks.

JEFF OSBORN That, okay. Cache, I am not an engineer. I run a company full of them and I get to be the dumbest guy in my company, which is fun, because I can always refer my questions to someone else. Cache poisoning isn't something that happens in a root server. It's something that happens in a resolver. Let me see whether one of my smarter than me people here knows enough to answer that. Hang on. You guys want to try it?

KEN RENARD

Hi, I'm Ken Renard. I am one of the root server operators as well, H Root. So the first point is that cache poisoning does not involve the root server system or a TLD or the authoritative server. So it's very much outside of our remit, but we can comment on that just as DNS nerds, that the responses that would normally come back from an authoritative server, sometimes they can be changed in root.

Now there's a nice thing called DNSSEC, DNS security extensions, which can really help that and almost make that impossible. So we strongly encourage everyone to turn on DNSSEC validation and to sign their zones as well. There's other older mostly mitigated attacks that can essentially poison the cache, but yeah, it's inside that resolver. So we sort of, as end users, sort of have a choice over what our resolver is. Sometimes it's our company or ISP or however. So ask your resolver what their policies are, how they mitigate cache poisoning, but it's a largely solved issue, but there's always going to be something that crops up. So.

SIRANUSH VARDANYAN

Thank you. I saw also hands over here. One more question. yeah, Maciek, let's go.

MACIEK PIASECKI

Hello, Maciek Piasecki, a fellow, that was going to be a bit geeky. So you said that basically every computer will know how to contact the root server. My question is how? So if I have Arduino or Apple

	II and I want to connect it to the internet because I hate money, where is that information stored or how is this update? Do I need to code it in manually or will the network card have it in its firmware?
JEFF OSBORN	If I understand the question correctly, you're saying how would you run a resolver yourself on an Arduino or a Raspberry Pi?
MACIEK PIASECKI	I want it to contact the root server so it knows how to get other addresses.
JEFF OSBORN	Oh, how would…?
MACIEK PIASECKI	How will it know how to contact the root server?
JEFF OSBORN	Root hints. There's a, I was getting three steps ahead. I've run our software, it's called BIND. Go to isc.org, it's my company. We are an open source software company so the source code is free, the software's free, fire it up. One of the base things in it is called a root hints file and it basically says there's a server at this address, there's a server at this address, there's a server at this address and it's IPv4 and IPv6 addresses.

So, theoretically that's a total of 26 addresses but I'm wrong, right? Because it's more complicated but effectively there's a hints file in the software and even I have installed BIND on a Raspberry Pi, not an Arduino, I think you'd need a shield because there's a memory issue but it's not that hard and it's really fun to play with if you find this interesting and it's free because we do it for the good of the internet.

SIRANUSH VARDANYAN Thank you very much. We are no longer accepting new questions, but because we are already up to the end of our session, allow me to thank you, Jeff, one more time and the group behind you and behind the scene, but thank you for coming and taking your time with this busy schedule to be here with us, appreciate it. And with that, our meeting is adjourned. Thank you very much.

[END OF TRANSCRIPTION]