ICANN84 | AGM – ccNSO: Joint Session with ccNSO and RSSAC
Wednesday, October 29, 2025 – 16:30 to 17:30 IST

CLAUDIA RUIZ

Hello, and welcome to the ccNSO Joint Session with RSSAC. My name is Claudia Ruiz, and I along with my colleague, Joke Braeken, are the participation managers for this session. Please note that this session is being recorded and is governed by the ICANN Community Participant Code of Conduct, the ICANN Expected Standards of Behavior, and the ICANN Community Anti-Harassment Policy.

Please observe the following guidelines to participate in this session. They will also be posted in the chat for your reference. During this session, questions or comments submitted in chat will be read aloud if put in the proper form as noted in the chat.

Interpretation for this session will include English, Spanish and French. If you would like to speak during this session, please raise your hand in Zoom. When called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use a physical microphone to speak. Please state your name for the record and the language you will speak if speaking a language other than English, and please speak at a reasonable pace to allow for accurate interpretation. Thank you. With that, I will now hand the floor over to Alejandra Reynoso. Thank you.

ALEJANDRA REYNOSO

Thank you, Claudia. Welcome, everyone, to the joint session between the ccNSO and the RSSAC, that is the Root Server System Advisory Committee. We will hear first from Jeff, the chair of the RSSAC. He will give us an introduction to the root server system. Afterwards, we will have a moment for questions and answers. And then we will hear from Peter Koch regarding the current status of the Root Server System Governance Working Group. These are two separate and distinct groups but they are related to the root servers. And in the end, we will have questions and the wrap-up. To start, I'll give the floor to Jeff who's going to share his slides.

JEFF OSBORN

Thank you, Alejandra. Thank you for giving us the opportunity. I think this is great to be able to meet as two groups. I couldn't find the last time that we all met. Now, let me see if I can make the slides work. Continue to share. And it looks good. That's amazing. That's the fastest I think I've ever had that work.

My name is Jeff Osborn. I am the chair of the RSSAC which is the Root Server System Advisory Committee. I'm pleased to be joining the Country Code Name Supporting Organization, ccNSO. Trying to get through an entire ICANN presentation without acronyms would take weeks so I will try to use the appropriate words in English when it makes sense. Feel free to call me if I'm using acronyms that are unfamiliar.

In overview, this is a presentation we prepared to address governments that are seeking to legislate or regulate some of the

actions of the Internet and for us, we worried about the root server system in particular. This is addressed at somebody who is extremely knowledgeable in many things but perhaps not the technology of how the DNS works. If I'm addressing you as people and it sounds like, "Wait a minute. Who's he talking to?" it's just that we wrote it as if we were addressing policymakers.

In the second part, we'll address how DNS is explained in a way that makes it misleading to people who operate things like governments and regulatory industries and how we had to kind of turn a technical story on its head to make it one that was more meaningful.

Then the third part, I'm going to try to explain a simple understanding of how the DNS root server system works. I'm sure all of you are aware you wouldn't last long at ICANN without at least a cursory understanding of DNS. Hopefully, if you'll bear with me, you'll gain a little more understanding, that would make me very happy.

Part one, the RSSAC. The RSSAC has a really, really narrow scope compared to a lot of organizations at ICANN. We are, and I quote, "To advise the ICANN Community and Board on matters relating to the operation, administration, security, and integrity of the root server system." That's it. We don't have a big broad remit to make policies about a lot of other things. We just ensure that we are available to advise the Community and the Board on how this one part of the system works. Obviously, this being ICANN, there's more

there. If you really want to dig into it, the next couple of pages are references for where you could go to do further work. As I said, the first thing policymakers do if you've reached them is have somebody else go chase down the details. I'm applying that rule here. Here are some details.

The composition of the RSSAC is fairly simple. There are 12 organizations that operate root servers around the world, a total of about 1900 as of today. Each root server operator appoints one primary member and a secondary member to the RSSAC. There's one vote per operator, so a total of 12 votes and 24 operator-appointed members. In addition, we receive non-voting liaisons who are still members from the IANA, the Root Zone Maintainer, the Internet Architecture Board, and the SSAC. We send liaisons to the ICANN Board, the Customer Standing Committee, the Root Zone Evolution Review Committee, the ICANN NomCom and various ad hoc work groups when it's deemed meaningful or necessary.

In addition, there's a larger group called the RSSAC Caucus. It's a group of subject member experts, something over 100 right now. I need to look that up in detail because I keep getting disagreements on how big it is. Membership to this is broadly available but because it is a working group rather than a learning forum, we don't take people who think DNS sounds interesting and they want to learn about it. We take people who have implemented DNS infrastructures on large scales like large corporations, universities,

or entire nations. It's a pretty amazing group and they are responsible for putting out much of what we publish.

We're are a forum for discussion, but as I said, we have this really narrow remit. We're basically operating a root server system that has operated without a failure, arguably, in about 40 years. So, it is a strong underpinning of the Internet, but our focus is narrow. We don't do a lot to change it. There isn't a need to change a whole lot of the operational part, because it just works really well. The nature of our policy recommendations generally is to restate things that fall within a long-established normative framework, and that is how we operate the root server system. We suggest continued improvements that grow from the existing framework. They're not as frequent as you would expect.

We are not an enforcement or supervisory body. We are not a representative body for the root server operators, and we are not a representative body for the root server system. Again, if you want to dig into it, we have a lot of detail. We've published a lot of documents over the years. RSSAC023 is a history document. The system, as in place now, has been pretty much that way for 25 years. We grew fast, and then due to a technological leap or two, we have been able to just maintain a global Internet without having nearly the growth that used to be required. We'll talk about anycast a little later. It's been kind of a miraculous multiplier of how this works. There are service descriptions on everything from how we measure, local perspectives, governance, and specific advice to the ICANN Board. Again, if you wanted to dig in deeper, this is pointers

to reference information. I'm certainly not going to go over it today. That is RSSAC as we fit within the ICANN infrastructure.

Second, if you've taken a course on how the DNS works, there is something that as a policymaker—see? I'm misleading again because that's what it was written for. You'd hear a couple of these words and they mean different things to engineers than they do to policymakers. We kind of had to turn the thing upside down. Most policymakers, I think, it's a given, don't understand the root server system, but in order to appropriately legislate or regulate, they have to. We're trying to explain what it is theoretically and what it means operationally to the existence of the global Internet.

Now, here's the normal way to explain DNS. You start at the top of the pyramid because they describe it as a hierarchy. This assumes that the Internet is a cold dead thing that has exactly one resolver that you just turned on and it knows nothing. That's not a very likely scenario, but that's literally how this is taught. The first thing that happens is somebody asks the root server, "How do I find TLDs?" Then they ask the TLDs, "How do I find a domain name?" And finally, you get an answer from a domain name. This assumes, like I said, a cold dead start. It explains a hierarchy in a way that's meaningful but it totally understates the importance of caching, which is really how most of this works.

This, for instance, is out of the otherwise really excellent ICANN document, and it shows, you can't quite read it, but the way over on the left, our resolver is trying to figure out a domain name, and

it goes to the root servers first, and then second, it goes to the TLD server, and then third, it goes to an authoritative server for the domain name. And that's how everybody expects this to work. The problem is it creates the false impression that the root server is a gatekeeper. The root server is the on-ramp. Nothing happens until you get to the gate server.

And from those original pictures, you would think this is the case. "Oh my God, all of my delay must be based on the root server because nothing happens until I talk to the root server." And a problem here is that the political definition of hierarchy leads you to think that thing at the top of the hierarchy controls us all. Well, to an engineer, the hierarchy just explains a cascading series of events which are so manipulated by using memory that it's kind of not meaningful. Let me just step ahead and say the harm is that there is a nearly zero observable effect on an end user if any component of the root server system goes out.

Like I dumped a Coke on my laptop last week. It took me three hours to get a new laptop out and have it running. This isn't even that. This is the half a millisecond it took because some component failed and you immediately find the next one. It's indiscernible to the end user what's going on when there is a failure in this system. There are literally over 1900 server instances using anycast, which means if one isn't there, you will literally go to the next closest one on the next and the next just in the protocol and you would be very

hard driven to tell that anything had failed anywhere along the way.

The failure mode also for this, as you recognize because you guys run ccTLDs is, one, you don't change things a whole lot. The numbers of ccTLD has pointed to are pretty consistent. The glide path of failure, if people stopped updating your TLD address is really long. There are nations that could make it years without getting a root server update and nobody would notice. It's really important to recognize rather than being the first thing that shows up in your failure path, this is really way down the route. An analogy I've used for people that seems to work is once you know that +44 will get you to England, you don't have to call somebody to ask what the country code is every time you call England. You just remember it's 44. This is a system that is designed for resiliency and literally has been operating for over 40 years without an end user identifiable failure.

Here's what we did to flip the system. Let's show how it really actually works. The Internet is alive. It is warm. The caches exist. People have memory of what they did before. The first time you ask a resolver, "Hey, where is www.mybigcompany.ie?" Well, when rather than saying, "I don't know anything, let's go find out what IE means," the resolver says, "I just answered that question 20 minutes ago. Here you go." And the local resolver, of which there are millions and you know them, 8.8.8.8, Quad9, 1.1 1.1, your ISPs, local providers, everybody's got a resolver out there. Anybody could have a resolver out there. There are millions of them. They're

the thing closest to the user and that's where you get the whole address over 90% of the time.

That's really different from the idea of first you've got to get a root server, then you've got to get a TLD authoritative, and then you get to your local guy. The local recursive resolver is most likely to give you all of the answer you need. Then going down through that, we end up getting statistics where about 0.02% of address queries have to make it to the root. Not every single one, one out of somewhere between 5 and 10,000. Is there anybody here who knew that coming into this room other than the RSSAC guys? I think this is kind of interesting to provide this this way because it puts a different face on just what is expected. Hopefully, we learned something else in the rest of these.

Now we get to part three. Alejandra is complaining. We had breakfast together and I had six cups of coffee since then and she thought it's a bad idea so she's trying to help me mitigate the effects. Thank you, Alejandra.

I'm going to step through something that we worked on a few of these models for how do we explain something that's really fairly complicated to smart people who just don't sit at terminals on command lines all day. We're going to attempt to do that. And when I get in here and explain what I'm doing and I'm going to ask you, really pay attention for about three minutes, and that's where we'll get through the whole thing. And I find that people really feel it's valuable to learn what we're trying to do here.

The goal of this, we're trying to increase your understanding of the system, the root server system and the root server operators by explaining the role and purpose of DNS, the roles of the resolver and the authoritative servers, and how, when, and why a resolver consults the root server system, and then the common misunderstandings. Again, this is things that leadership needs to know. I think it relates well to all of you in this room as well.

Here's the intro to DNS. We started this two years ago when we were in a room full of engineers not dissimilar to this. And I thought about it and I said, "Show me by a show of hands who in here your mother knows what you do for a living?" And other than Eberhard, nobody's mother knows. My mother is a bright woman, she's got a master's degree, she has never known what I've done since I got out of college. The idea here was to explain to somebody who is clever, thoughtful, and knowledgeable but not technically enabled. This one page, I'm kind of proud of because it sort of explains the reason DNS exists.

DNS uses human names to find computer addresses. That's pretty much it. Humans know domain names like www.amazon.de. Computers need IP addresses. We all know what those look like. I'm using IPv4 in here, 18.239.62.181. And DNS is simply the way that the words find their way into addresses. The DNS is how we know this word is this number. And the neat thing is the numbers can change. So, the operators have a lot of control but they don't have to change the name.

This was actually the address for amazon.com when I made this slide about a year ago. It's probably something different. Nobody cares. Nobody missed Christmas shopping. Nobody missed buying a raincoat. The numbers can change while the names remain the same. And almost all connected devices out there use the DNS to find each other. It was computers and servers when I was a kid, and now it's smartphones and my washing machine and my dryer. And I think my washer and dryer are having a communication in IP and I have no idea what they talk about, but they found each other through DNS so it's kind of horrible. DNS questions, ask domain names and the answers are IP addresses. The benefits, obviously, are it's human identifiable. It's easy to remember a string of words. Numbers are hard. Service portability, like I said, means the numbers can change and the words are the same.

But it's amazing when you look into it just how big this is. It is a huge distributed network that's easy to use. Many of you are aware each of your nation has a fairly large authoritative server of a bunch of domain names. What isn't always obvious is some of those domain names, each have dozens of millions of next level up things. Because some organizations—like we use an HR company called bamboohr.com. Well, I'm to the left of that. My company is called ISC, so isc.bamboo.hr, that's up there in that third level where .com, their sign's got to tell me how to get to the one address and then Bamboo is authoritative, server is going to tell me how to get to them, and then Bamboo has to tell me how to get to where I am. And then further, I can put my names in front of that. The fact that

these all resolved in milliseconds around the world and it all works is kind of amazing.

I'm going to skip a little of this because I don't need to just drive it all home. Let's get to the guts. Here's the fun part of this. This is the part where if you—I tell the beginning group in the newcomers that if you understand this well after three minutes, you can get a job as a DNS engineer and it pays really well and you get green cards and H1B visas and all. The world becomes your oyster. But you guys look like a less easy to convince group. You're more jaded, and that's probably good.

Anyway, in the case one, which is the vastly most frequent case, we're going to start up here on the top with a question. We're starting with, "What is the IP address for www.example.com?" And at the end, we're only going to stop when we have an answer to that. Okay, simple concept. My computer needs to know where's www.example.com, and I'll be done when I have that answer.

In the vast majority of cases, do I know the answer? Mr. Resolver—the gray box is my local resolver I'm asking, and as we see over there on the side, it's got a cache memory. So when I say, "Do you know where it is?" The answer is, "Yes, here you go." You didn't have to go anywhere else. This is literally the answer over 90% of the time. That's all you have to do. You go to your local resolver, which could be in your office. For me, it's on my desk. It's somewhere in your ISP or at your company. It's a pretty nearby

thing. This is not a global system because it's usually remembered. Now, in the second scenario…

What did I do? I just touched something wrong. Okay. Sorry about that.

Next, let's try it where we're missing a piece of information. We start now with what is that IP address? Do I know the answer? And it turns out we don't have it. We don't have www.example.com stored in memory. Well, then you peel off to find out what's the next recursive step, hence, the term recursive resolver. "Do I know where example.com is?" "Yes, I do." "Okay, tell me where www is?" I ask example.com. It gives me the answer. I put it in cache memory. Now do I know the answer? Yes, I do. It just took another level of looking for that information. It wasn't in local storage. All right.

Now we go into the harder one. That last one only takes about 5% of the time you have to even go that far. Now, at around 2% of the time, I'm asking the question, "Where is www.example.com? Do I know?" "No, I don't know." "Do I know example.com?" "No, I don't know." "Do I know .com?" "Yes." Okay. We'll begin the recursive process of putting the rest of the address together. I ask, "Do you know example.com?" It says yes and I put it in memory. "Do I know the answer?" "No, I'm still missing a section." "Do I know example.com?" "Yes, I do." Ask example.com, "Where's the www?" And it tells me the address, I put it in memory. Now do I know the answer? Yes. This is exhausting. It's a good thing locally you get

over 90% of this information because this really gets complicated the farther you go.

We're now going to take the fourth case, the one out of 5000 to 10,000 case where we don't know anything. The question starts, "What is the IP address for www.example.com?" "I don't know." "Do I know example.com?" "No." "Do I know .com?" And if the answer to that is no, you probably have a resolver fresh out of the box that just got turned on. I'm the president of the company that makes the resolver software that's probably still most widely used. It's been around forever. It's called BIND, the company's ISC. And one of the things that contains is root hints that basically says, "You don't know anything but you do know how to find a root server." What it does is it wakes up and it goes, "Well, how do I find .com?" And the root server, we're finally getting a root server involved in this, goes, "Here's the IP address for .com. Stick it in your memory." "Okay, thanks." Do I know how to find example .com? No. But now I can find .com. Excellent.

So, I ask for example.com from .com, gives me a memory. Now do I know the answer? No, because we've got to go back and recurse. Because all of these are held at different organizations' places because they control what they control. "Do you know www.example.com?" I'm going to get the words wrong. I ask example.com, "Yes, it does." It puts it in memory. And do we know the answer? Yes. In the one out of 5000 or 10,000 cases where we had to go all the way to the root server, we went recursively, which is to say through all of the layers and we found out the information

from at least three places spread across the globe on databases managed by different people and we put together the information. That's kind of cool, isn't it? It's sort of remarkable that it works and it's literally been a part of the Internet for 40 years.

This is the ugliest slide in the deck, but the information is really useful. There are about 350 million domain name zones out there and they're primarily maintained by the domain name registrant. These are not the job of the TLDs. These are by the people who actually operate their domain name. The number of total TLD zones is probably not surprising to you. There are about 1450 right now, although ICANN is trying to add more. And these domains are maintained by the TLD registry, as you guys well know. There is one root zone. It has about 1450 TLDs on it, not surprisingly. And it is maintained not by the root servers, but by IANA, which then cryptographically signs it—or the root zone maintainer cryptographically signs it and hands it off to the root servers.

In review, a root server holds a copy of the root zone. The root zone holds addresses of about 1450 TLD registry authoritative servers like .com, .nl, .jobs. The TLD has an authoritative server that knows the address for the next step. You could take .com and make it amazon.com or tiktok.com. You get the idea. I have various audiences and I'm sorry if I'm infantilizing the information, I apologize. I just don't know how else to get through the slides. The domain names' authoritative server obviously knows the third thing to the left, and the resolver finds and returns the answer.

Here's the surprising part, and I made this big slide because someday I'm going to put it on a t-shirt. In the millisecond world of a resolver, queries to the root server system are rare. That is really not obvious. You could spend a lot of time looking around in the DNS and not realize that. I'll go quickly through these, but we try to make a couple of points that I think you all know. The root server system does not carry or manage content. All we provide is TLD addresses. That's it. Not even a domain name address. That's it. We give you the country code for England is 44. We give you the—I can't fake an IP address off the top of my head of any of the TLDs, but if I could, this would be a good place to do it. We literally just give you that bottom loop and that recursive search to let you go back to another TLD, to let you go back to an authoritative server.

We don't manage content. We are not a gatekeeper. We just hand out the address so you can go look stuff up. We are stable, secure, and resilient. There are just short of 2000 globally distributed server instances, and each of them contains the entire root zone. It's not a partial thing. Every one of them contains all of it. The 12 of us operate diverse hardware platforms, running diverse operating systems, running diverse DNS applications, and diverse data routing. There is no single point of technological failure to a degree that's kind of incredible. It is so hard to try to come up with disaster scenarios where the system fails that it's really, really remarkable.

My boss, who's a real Internet pioneer, I won't go into that, but I asked him, I said, "I got to come up with a failure scenario for the root server system." He thought about it and he said, "In about the

seventh month of nuclear winter...” Okay. That’s kind of a stretch. But it’s unbelievably difficult to imagine a fall-off when you have the same address feeding hundreds of servers. You could literally lose all but one of ISC’s devices and pretty much answer all of the questions. And if that wasn’t enough, it doesn’t matter because there’s another 11 out there with the same situation.

You think of banks that use a device and a fall-over. This is hundreds of fall-overs with the same address. It’s a remarkably resilient system. There’s no point of institutional failure because the 12 operators are entirely disconnected from each other. They’re in different industries, different places. There is no imaginable force-majeure event, for instance, like what happened in Mauritius, where a single-court injunction could cause problems across operators. So there’s no point of institutional failure.

Again, the systems operated since the 1980s without a system-wide blackout. We’re a favorite target of DDoS attacks. We are really, really hard to attack. I don’t want to say we’re unimpenetrable because somebody will go try to prove themselves by ramping it up, and who needs the grief? But basically, this is 30 years of successful 24/7 by 365 operation.

You guys know this, but I’ll have to say it. We don’t decide what’s in the root zone. We don’t cut countries out. We don’t add countries. We don’t add TLDs. We don’t cut them out. We literally get handed a cryptographically signed file and serve it. We serve address information, not content. I’ve said that before. We don’t decide

whose address shows up. We're not a gatekeeper. And this is a remarkably, remarkably resilient system. Thanks for your time.

ALEJANDRA REYNOSO: Thank you very much, Jeff. That was an amazing presentation. And I think we all have it very clearly now how the DNS works. Do we have any questions for Jeff? Jodi?

JODI ANDERSON: Thank you. Jodi Anderson, Internet New Zealand, for the record. You said at the beginning that this is the presentation you do for governments that are seeking to regulate some of the Internet. What is your sort of connection at the end once you've done this presentation? What's the thing you then say, "This is why you don't need to regulate the Internet"? Are you saying that this is a great system, it's not going to fall down so you don't need to regulate? We don't do anything to do with content so you don't need to regulate? What is the through line?

JEFF OSBORN: That is a really good and insightful question. We're mostly a bunch of people who write software and give it away. We do it without the benefit of a shareholder anywhere. I run a nonprofit company and we give away. BIND is probably operated by—what did we figure?—a million and a half organizations around the world right now. As of last week, 103 of them pay us any money. There's another of those 10,000 to one things. We basically are driven to provide free and

open-source software to the world for the good of the Internet. And similarly, the root server system, we don't make money doing this. This is an expense. I run a small nonprofit and spend $2 million a year running our portion of this because the Internet needs it. And when we were first asked to do it nearly 30 years ago, we said we would, and so we still do.

The motivations are a little funny. There's a lot of money sloshing around in the Internet and we can't match all of that lobbying money or anything else. All we can do is try to ensure that people realize that—for me, at least, there's a part of the Internet where we genuinely are trying to do the right thing. We are not providing harm in any way. We are not making obscene profits in any way. We're not trying to manipulate truth in any way. That by itself is a hard sell. Jamie Hedlund and some of the other folks at ICANN, we have worked with, and I think they're much better at directing very specific information. But for us, I just kind of want people to know here's what's going on. Is that any kind of helpful?

ALEJANDRA REYNOSO Wait a second. I just will put the queue. We have Olga, we have Liman, we have Charles, and Pierre, and Wes. In that order.

OLGA CAVALLI Thank you. Thank you very much for your presentation. Very, very interesting. I have a question about the content of the root servers. It's me here. Sorry. The back of the room. Olga from the ccNSO. All

the root servers have the same content, I can imagine. How often are all they updated per day? Per hour? If you can give us some information about that.

JEFF OSBORN Sure. I think Wes was upset that I didn't let him ask the last question so I'm going to hand this to him now.

WES HARDAKER That's not quite the way I put it. I'm Wes Hardaker. I actually work for Google, but I help with the USC ISI Root, which is B-Root. To go back to the previous question for one second, it was a great question. I think that the message that we really want to take away, Jeff answered it from his perspective, thinking about the bigger perspective of all the operators. The root servers answer far less queries than most people think. And that doesn't mean that it's not important as a critical bootstrapping element, but it's not as important as people think it is. The ccTLD zones serve a lot more data, probably almost all of them, than the root system does every day. At the same time, it's still designed to be incredibly resilient. That's the message that we really want to get on to the government systems.

With respect to how often the zone updates occur, on average, it's two to three times a day. And the way the notification mechanism works technically, they all update very rapidly. In fact, we typically notice. I monitor all of my particular systems to make sure that

there's not a significant lag. And if there is, I get woken up in the middle of the night. But actually, the reality is is you don't need your data updated that quickly either. Thank you.

ALEJANDRA REYNOSO Thank you. Charles?

CHARLES NOIR Thanks. Charles Noir from .ca. The question I'm about to ask, I know it was a difficult one, and probably we can have a lot of discussion about it, it might seem fairly simple, but we often talk about the need and the benefit of a Unifiedroot. It's taken as gospel as we're headed into discussions around digital sovereignty, increasingly governments thinking about looking at how they can manage some of this on their own, or potentially in different ways, what would you argue, given your presentation, would be the benefits of a Unifiedroot within that context?

JEFF OSBORN I'm not even familiar with the term Unifiedroot, unless it's called—oh, God, I was spelling root wrong in my head. That's so funny. Anybody want to take a swing out? I can try this. Wes?

WES HARDAKER I should also mention I'm the RSSAC representative to the ICANN Board. I sit on the Board of directors on behalf of RSSAC as well. The best document to read about a unique DNS root is actually

published in the IETF by the Internet Architecture Board way back in May of 2000. It's RFC 2826, and it talks about the need for the global namespace to have a unique mechanism, a unique naming structure, so that when I send mail to somebody on the other side of the world about some new website I've discovered, we don't get to two different places, right? It's very important to have a unique naming system. I think that's what it means.

There is a lot of talk lately about alternate naming systems, and there has been work in various venues, I can't dive in the technical details here, about trying to make sure that they are compatible or that they will be interoperable. And the IETF actually has an active draft right now talking about what does it mean to have multiple naming systems that actually interoperate properly.

ALEJANDRA REYNOSO Okay. I have Liman, then Jordan. Anyone else? And Pierre. Liman, Pierre, and Jordan, sorry.

LARS-JOHAN LIMAN Thank you. From the back here, I'm Lars-Johan Liman with Netnod, also a root server operator and a member of RSSAC. I would like to add a different angle to what Jeff said, how we respond to regulators, why they don't have to regulate, because the root server operators need to be neutral. In order to preserve the unique root, which is as important as having a unique numbering system for the telephones. I don't want to sit here in Ireland and dial my

kids at home and end up in someone else's telephone because they have a different numbering system in Ireland. So, the naming system is just as important as the telephone numbering scheme.

And to preserve that, the root server operators need to be able to operate in all jurisdiction all across the globe. And if a certain limited of these jurisdictions start to regulate how the root servers operate or want to have an interest and have special treatment, that is going to break down. The unique root is utterly important and in order to preserve the operation, the technical maintenance to make this infrastructure work, the root server operators have to be independent from regulators, from political systems, from high influence, from specific business entities, and everything like that. That's why we have our operational principles, that I don't remember if you alluded to Jeff. But we have a set of operational principles that we operate accordingly and the independence from ICANN and the IANA from as many other entities as we can, and also independent from each other. The root server operators try to be independent from each other. All that is very important to keep this infrastructure running across the entire globe.

I don't serve Sweden where we are based. I don't serve Europe. I serve the world. We have servers in Bhutan. We have servers in Vanuatu. We have servers in Rwanda. We try to serve the world as best we can. And if we were to start to see strings from various entities that want to regulate, that's not going to work anymore

because we will have conflicting requirements from different parts of the world, and that just doesn't fly. Thanks.

ALEJANDRA REYNOSO: Thank you very much, Liman. Pierre and Jordan, and we will close the queue there because we need to continue with our next presenters. Pierre?

PIERRE BONIS: Thank you very much for this presentation. Surely we will use it when we discuss with our government. Pierre Bonis from .fr. I have two questions. The first maybe the easiest one is you show examples of queries of real domain names in real TLD. But what happens or what would happen or does it happen that at the root server level, you have a lot of queries of non-existing TLDs? Which means that in what you show us, I understand that it goes directly to the root because, of course, the resolvers, I don't know. Is it something that you see as a danger that someone would query millions and millions and millions of fake TLDs to test the capacity of response of the root system? That was my first question.

And my second question is a little bit more political, let's say. I've seen, I think, a report that was open to consultation from RSSAC, I guess, talking about the 13 root servers and there is no need to extend that. I just want to understand, at one point I was explained that we could not have more than 13 root servers because of technical capacity in the—I don't know in what but that was a

question of bit, no more than 13 root servers. And now there is another explanation. Why should not be 20 or 25 or 50 root servers? Especially in the context where continents ask for that. I don't say that I agree with it, but I think I would like to have your vision on it. Why is there no root server in Africa? Why is there no root server in some parts of Asia? Thank you.

ALEJANDRA REYNOSO I'll interject here really quick, just because that was a brilliant question, but it requires a longer answer. So if we can hold it a little bit so we can hear from Peter first, and then we come back to it. And if we don't have the time, then we can do—

JEFF OSBORN I and many of the other guys are free after the session. So we have a lot of time, if you want to. It's a long answer, but it's a great question.

ALEJANDRA REYNOSO Exactly. Jordan, would you mind if we defer yours? Or is it a super-quick three-second type of question?

JORDAN CARTER It's kind of quick, and unfortunately for you, it's a kind of comment, but on this question that's come a few times, like what about the root? One of the things I've tried to say to people as an example is, how would you feel if the country you're trying to phone could

break your phone dialing? Or what your local control could give you, if you had digital sovereignty over the root zone? All you could do is either not change it anyway, or mean that your own people couldn't get somewhere. Is that a reasonable analogy? Have a think about it. We'll talk about it later.

ALEJANDRA REYNOSO Let's keep that also on hold, and if we have time in the end, we will come back to those. Right now, we are moving to Peter who will talk to us about the Root Server System Governance Working Group, again, a separate group from the RSSAC. Peter, the floor is yours, and when you need the next slide, please request it.

PETER KOCH Thank you so much, Alejandra, and thank you for your patience and waiting. I hope that much of what you asked for will get a bit of an angle in this presentation. Now, for something completely different, which is the RSSAC Root Server System Advisory Committee. It's an AC, it's a Bylaws-mandated Advisory Committee. And now there is also the Root Server System Governance Working Group. It is a Board-scheduled, at that time, ad hoc, finite working group composed of members across the community. All the root server operators are represented, as are the ccTLDs. That would be myself and our friend and colleague, Luis over there, and other ACs and SOs have representatives as well.

First of all, and you know this, I'm not speaking for the committee, nor is anybody else. The state of this working group actually is that we have produced a draft report. There was a Public Comment period. The comments are in. And during this week, and actually the week before, we have been digesting the comments, and then we are still working on next steps in terms of how we deal with those comments in the report, and then the report will go to the Board. And of course, it is not an RSSAC activity. The reason that we put it here was that in the preparation the Council thought, while this group is there and this governance questions are on the table, it would be a missed opportunity if we did not add this governance angle to the technical angle that that Jeff so well presented.

Now, that said, again, we are members. We are appointed, in the case of Luis and myself, by the ccNSO. We are not representatives. All errors are mine. And there are other members of the group here who might chime in or might be available afterwards for conversation.

So what you've seen here already is two things. One is that there's obviously a certain dissonance between the operational importance that we learned, like every so many millions queries there and here and there, and the governance effort. We have RSSAC, we have a Governance Working Group, there's also RZERC, right? So it can't be so totally unimportant, and that might be a bit difficult, and let me try to untangle that in a way.

Another remark, what you've also heard is that—and that is probably resonating very well with ccTLDs, the root server operators are different from each other, so are we? And they are loosely organized, right? Even though they serve the same zone, the root zone, it's a loosely organized system. So, no formality, nobody speaks for the root server operators, right? Similar to us. Nobody of us speaks for all the ccTLDs, not even the ccNSO does. So there are lots of policy similarities that makes this, I think, interesting and noteworthy for ccTLDs, this independence and such.

One word on the GWG, I just looked it in my inbox. It was March 2022 when the GWG wasn't instantiated, but restarted. And the call for volunteers said the GWG is expecting one more year for its work. That was three and a half years ago, but we are really there now, I think. That is owed to the complexity of the questions, and also these questions, part of these have actually mentioned, Pierre was going into that direction, ccTLD colleagues will remember that there was a Delegation-Redelegation Working Group, for example, which was issues on the table that weren't completely covered in RFC 1591, which was a very important policy step for us. And this question, what would happen if a root server goes out of operation for whatever reason? Like a succession, there is no current policy for that. Same holds for what is the number? What would happen in here and there?

However, the task for this Governance Working Group is not to answer these questions. The task is a meta task. Design a

governance structure that then will go deal with these questions. We need to first have this body and compile the body to make sure that these questions can be responded to. And this is where we currently are.

So they're the Principles document and the Draft Governance Structure document, which both were put up for public comment. As I said in the introduction, we are digesting this. And let me have a short look onto being mindful of time. Can I have the next slide, please? Hopefully the graph. Yeah.

How to address these questions? The proposed result of the Governance Working Group is that there, first of all, be a Council, or maybe I should say the working group proposes a staged approach with phases one after another that should be initiated once the Board has accepted and evaluated the final report, and in the initiation phase—heavens, in the first phase—there will be a Council comprised of root server operators and what the working group calls stake participating stakeholder groups that have a direct operational dependency on the root zone. And in our deliberations, we came up with resolver operators as directly dependent, but we also were thinking of the roots. However, for TLD operators, we already have two distinct and mature organizations, namely ourselves and the Registry Stakeholder Group/GNSO on the gTLD side. Currently, there was no representation of resolver operators, and we did not only look into ICANN structures because it doesn't have to be in ICANN structures, but also outside. And while there are potential candidates, nothing

like that is organized at the moment. So we started from what is existing. And given that you heard that there are 12 root server operators to balance both the discussion and maybe any decision-making, there's also 12 representatives from the TLD side, like 6 from the Gs and 6 from the Cs. That's the suggestion. And then there should be 3 liaison seats from IANA, the Root Zone Maintainer, and the IETF/IAB for being responsible for the protocol, but also having a stake, for example, in the arpa TLD. The Council will have a secretariat, and so on and so forth. That's how to start with this. And may I have the next slide, please.

This is, by the way, just copied and pasted from the report. This is nothing new. So this is agreed upon, by the working group. And also you can reference that if you go back and read the report. So, in the establishment phase, which would be declared, started after the initialization phase, based on some criteria and milestones that are laid out in the draft report, there would be a number of—I think we call it functions. You can think of committees, but we chose the name functions. Of course, strategy, financial, the performance monitoring, security, incident reporting, and going back to what we had as Delegation-Redelegation Working Group, a designation and removal function that would deal with this. And that's where we are currently. There's a bit more phases, but I'll spare that. Mindful of time here. This is where the governance is.

Currently, again, there is a lot of voluntary reporting already by root server operators. We've heard about potential failure modes. The governance structure is expected to deliver more details and

organize more monitoring that then, for example, could be presented to regulators, policymakers as this is how the system works, and this is going to foster even more transparency and install accountability so that the other interesting governance question can be addressed. I'll stop here and maybe we open the floor for questions or remarks from other members of the Governance Working Group. Thank you.

ALEJANDRA REYNOSO Thank you very much, Peter. Any questions, comments for Peter? I see no hands up. Pierre?

PIERRE BONIS Very quickly. The report I quoted was the one you presented. That was just my point. When I asked, "Why is there 13 and not 14?" and what are the reasons behind. It was in your report. Or it wasn't, but it was expected in your report. I mean, the question I was asking, and maybe there is no answer, but there is so long answer that we have to spend two hours on it, is that on the root system there are, maybe I'm wrong, 13 root servers.

PETER KOCH Yeah. There's 13 names for root name servers in the root zone. And the reason for that used to be technical. But now you know the protocol has evolved over time, so these size restrictions can be—I'm going to be careful here with so many involved people in there—the sole size restrictions, there's debates about that. Okay.

However, the question of removal or addition, and the question whether it makes sense to extend this, like the system as it is obviously works. Is there a need for, and that is going to the principles in the document, technical need for changes to a running system, that is then to be discussed within the context of the appropriate committees and the Council as provided.

ALEJANDRA REYNOSO Eberhard?

EBERHARD LISSE Eberhard Lisse from .na. With regards to whether we want root servers in Africa, I don't think that's actually the right point. We want anycast instances in Africa. If I'm sitting on the fiber optic cable to Portugal, I reach the root server that is in Portugal quicker than the one in Johannesburg even though the distance is much less. And if you put an additional, let's say, 14th root server into Africa, then we need to get anycast instances to protect it from denial-of-service attack. So while it is an acceptable question, and it's probably not even politically correct, but as sitting on the recipient side of this, I would think it's not really a good idea to change a running system that is not broken.

ALEJANDRA REYNOSO Thank you. Jeff?

JEFF OSBORN

God bless you. The system—okay, as I understand this, I mean, I'm old enough. I've been around the commercial Internet for 40 years, but the way it's been best explained to me that made sense was the first root server operators opened up in 1984, I want to say. And I don't know about you, I didn't know DNS existed. We were still writing Etsy host files and trading them around with our friends. So, this was really early days. And sometime in either the very late '90s or 2000, the last of these root servers was added. And then soon after the implementation of anycast kind of made it unnecessary, like all of a sudden we had this crazy good tool. Eberhard is talking about how implementing anycast in Africa is more of a challenge. It turns out, the obvious thing isn't how many boxes do you have running, how much software. It's what's your routing fabric look like. And anycast is just a freaking miracle that you can say, "I want to find .82," and it can go through however many there are, topologically, finding you the most available one. And, of course, geography isn't topology. Yeah, I can do this only a four-hour answer with six of my closest friends talking more or I don't know what else to do with this in the interest of time, but it's a great, really important question.

PIERRE BONIS

Thank you very much. Just very quickly, because there are answers to my questions, but this is also a debate. And I just want to say that the question I asked was not is it technically necessary. It's because it's wanted by some people, and we have the same debate. And I will end there. We had the same debate before the transition of

IANA. If it's ain't broken, why do you want to fix it? It was political. So, is there any problem to have a 14—I don't ask if it's necessary. I ask if it's problematic. And if we don't answer this kind of questions, we are going to maybe think that the decision-maker and politicians are a bunch of stupid people, but at the end of the day, they have the power. So having being able to say, if you want to do a 14, 15, or 16 root server, we don't care, as long as you do it the way the others are done are not the same kind of answer as to say with anycast you don't need it, so don't ask the question.

ALEJANDRA REYNOSO We are running out of time, and this is a great debate. I will give the last words to Liman in the back.

LARS-JOHAN LIMAN Thank you. Yes, there is actually a problem because the root server system is technical infrastructure. What I hear you say is that there are political forces, in some cases, also financial forces, business forces that want to do this for reasons that have nothing to do with technology, nothing to do with infrastructure. And if we let one of them across the bridge, there will be a line of 250 more, and that's going to be a problem.

ALEJANDRA REYNOSO Thank you very much, Liman. Before I hand it over back to our presenters to say their final words, I want to put your eyes into your screens and join Menti. We would like your feedback on this

particular session. You know the drill. You have the QR code. You have the code on the top of your screen. So, Jeff and Peter, any final last words?

JEFF OSBORN Thank you very much. This is an ongoing discussion and I'm interested in any and all outputs. I'm easy to find. Bring it on. Thank you very much. I really appreciated the opportunity to talk to you today. Thank you very much.

PETER KOCH Yeah, same here. Thank you very much for coming. Thank you, especially for staying and for being involved. This is a policy question that will continue, because, to my ccTLD friends, there's six seats to fill, should this be instantiated, and that needs people who are interested and willing to work at this very interesting junction between tech and policy, because we've heard that both is important in this case, and we should be ready to take that. Thank you.

ALEJANDRA REYNOSO Thank you. We can stop the recording.

[END OF TRANSCRIPTION]