ICANN78 | AGM – How it Works Root Server Operations
Sunday, October 22, 2023 – 1:15 to 2:30 HAM

YAZID AKAHNHO: Okay, so hi everyone. Thanks for joining this session which is also streamed online. So we have remote participants as well as in-person participants and of course we are going to dive a little bit into the root server system this afternoon. My name is Yazid, Akanho. I work for ICANN, technical engagement under the office of the CTO. Most of you probably now understand or know what is the office of the CTO now. And this afternoon I'm here with four gurus of the root server system. I'm not sure we could have better people than them to drive us through this critical infrastructure of the internet. I won't take the risk to present them. They will introduce themselves and then we'll go into the presentation. Of course this session is being recorded and feel free to share your questions. For those who are in the remote, you can either type your question in the in the chat or raise your hand. Sorry, I need to cross check that with the support team. Should they type the question into the chat? Is that correct? Yes, okay, cool. And for those who are in the room, you will have chance to also ask questions. But before that, let's get our presenters to introduce themselves. Liman, start first.

LARS-JOHAN LIMAN: Hello and welcome. My name is Lars-Johan Liman. I work for one of the root server operators called Netnod based in Stockholm, Sweden. I'm happy to be here and I will talk to some aspects of the root server system, but I won't be starting. Thanks.

KEN RENARD: I'm Ken Renard. I work with the Army Research Lab in the US. We're also one of the root server operators and also acting right now as the RSSAC vice chair.

WES HARDAKER: And I'm Wes Hardaker. I work for the University of Southern California, which is actually where the root server system got started. I didn't start it, it wasn't me, but I do work for the organization that started it originally. Thanks.

KARL REUSS: Yeah, I'm Karl Reuss with the University of Maryland. We operate a root server and I'm here to answer questions.

YAZID AKAHNHO: Cool, thanks for your respective introductions and let's get into the presentations.

KEN RENARD: All right, so we can go to the next slide and talk about our agenda here. We do this session for you guys. We really want to tailor this as much as possible to what questions you might have. Make sure that you can understand what the root server system is, also what it is not. So we can get super technical. We're engineers and we love to do this. We love to talk ones and zeros, but again, we really want to tailor this to you guys. If you have questions, raise a hand. If somebody can watch the Zoom

room as well, we appreciate that. But we want you to come out of here understanding as much as you can about the root server system and what it means to you when you go back to your other groups. So we can go to the next slide. Actually, the next one after that.

So this is just fundamental stuff here. IP addresses are really how computers talk to each other. If you're connected to the internet, you have an IP address. You can see on the bottom right there, there's examples of what an IP address looks like, but we generally don't like to touch these IP addresses. Go to the next slide.

So we put names on there. Originally, the problem was that these numbers are just hard to remember. And actually, they often change. You'd be surprised how much addresses of websites and systems on the internet change a lot. And when you're talking about some of the advanced websites, some of the hosting things like that, it's not very clear exactly what an IP address means as far as a name. It's actually quite complicated. So in comes DNS, and thanks to Wes's colleagues who started this system about 40 years ago. So the next slide.

So we see this lookup mechanism. It's a hierarchical namespace. The main use of it is to map a name, like www.example.org, into that IP address so now that the computers can talk to each other. There are a lot of other uses. So this is, mail servers get looked up this way. We've got, obviously, IPv6 addresses. Security information is looked up in the DNS as well. So it's certainly a lot more than just looking up names to IP addresses, but that's primarily what we think of when we talk about DNS. So it's globally distributed, loosely coherent, scalable, and it's split up into these zones. You see the different boxes on here where at sort

of the top of the namespace there is root, and underneath that we have our top-level domains. Those are the administrative separations of who controls the information in there. Go to the next slide.

And so the process, kind of step through this and see where the root server system fits into the grand scheme of things. So on the left-hand side there you've got your device, your phone, your web browser, your server. Pretty much every computer on the internet will take on this role at some point where it needs to look up an address. So its first step is to contact a recursive resolver. Now that's something usually run by your ISP. Maybe you have it set to go to Google or Quad1, Quad9, whatever. Those recursive resolvers are kind of the workhorse of the DNS world. Those are the ones that are going to contact those servers on the right. Those are the authoritative servers. And they're going to do this dance. They're going to look at the root first if they need to, then go to the top level domain if they need to, and finally find that answer. The big thing about this is that the recursive resolver has a cache. And that is used quite heavily to optimize DNS. So if you send a query, a DNS query from your device to that recursive resolver, it probably already knows the answer or parts of the answer.

But when it doesn't, let's say we theoretically reboot the internet and there's nothing in this cache, you're looking up this name www.example.com. At this point, the recursive resolver has nothing in its cache. The first thing it would do is ask that question to a computer in the root server system. And it's going to ask usually that full question, hey, what's the address of www.example.com? The root server system is going to say, I don't know, but hey, go ask .com. And it gives the address a reference to the .com servers. That process repeats, it

recurses through the root to the top level domain and eventually until it finds its answer. Now in this scenario, it asks three different servers questions and got three different answers in that path. It's going to cache every one of those. So now when your device needs to ask another question, it probably doesn't need to go to every step of this process. If you're asking for something else in the .com TLD, it does not need to go back to the root server system. It already knows where to find .com. It will know for the next 24, say 48 hours, it won't have to go back. Go to the next slide.

That's because of caching. So I said the recursive resolvers, they have their cache, they remember a lot of these answers. Each DNS record has a TTL or time to live value. And that's how long that recursive resolver is sort of allowed to keep that in its cache before it's considered stale and it should ask the question again. Typically for the root zone, that value is 48 hours. So you only need to go to the root zone every 48 hours to look up that TLD entry that you're looking for.

So the root servers only know about the top level domains. So we don't know of anything, generally there's nothing www in the data we serve. So it's just how to find the next step, the top level domains.

So as DNS has evolved, the root server system has been involved with that as well, specifically DNSSEC. So the root zone was signed how many years ago? 2010, so 10 to 15 years now. So DNSSEC is an important mechanism that's there in the root zone. There's some work on privacy that is gaining some traction. Some of those technologies are query name minimization, DNS over TLS, DNS over HTTPS, as well as DNS over QUIC. And Anycast, we'll talk a little bit more about this

shortly, but this is a way to scale systems on the internet that provide services like the DNS can. It allows us to replicate things a lot wider than just putting additional IP addresses out there.

All right, talk a little bit about how things are set up today. If we go to the next slide here. So if you get a few definitions out of the way, the root zone is actually the data. That's the database or the information about the TLDs and how to contact their servers. In the DNS hierarchy, there's no parent zone for this. This is where we start at the top of the namespace. And the root server system, these are the set of systems, these are the computers on the internet that collectively serve the root zone. So they're just answering the queries. A root server operator is an organization that's responsible for managing that set of root server systems on the internet. We get down to a smaller piece here as a root server Anycast instance. That's typically what we think of as a single machine somewhere on the internet that's answering questions.

Then we have the Root Server System Advisory Committee. That's the group here within the ICANN space that provides advice to the board about matters of the root server system. And we will talk a good bit about that in a little bit.

All right, so one thing that often gets confused is the root zone versus the root server system. And where people think the root server system actually controls the data, we don't. We just get a copy of it and operate and serve the queries. We're kind of like the IT department for the root zone. So you see on there the root zone is managed by ICANN. That's the IANA function. The root zone maintainer is another piece contracted by ICANN to compile and sign the data and then distribute that to the

root servers. So the root server system, again, we're just the computers that respond. There are 13 identities, each with an IPv4 and IPv6 address. Those are the numbers that you go to, but it's not 13 machines. There's actually, this slide's a little bit out of date, I think there's over 1700 instances globally. So each of those 13 identifiers, those 13 IPv4 addresses are replicated all over the world. And the root server operators, these organizations, we're responsible for running and caring for those machines and making sure they work properly. Next slide.

This is basically the same message, just pictorially here. So on the left you have TLD operators that want to make changes or if there are new TLDs, that all happens outside of the root server system. That's the IANA function. They maintain the database. They give that data to the root zone maintainer, which formats it, does some of the final signing. That's where they send it out to us as operators. And as operators, we operate all of our instances and that's where queries come from, from those recursive resolvers, as well as anybody who wants to query our data.

A little bit of the history and the growth of the system. So almost 40 years ago was Wes's group that started this and by 1985 there were actually four different root server operators. Back then it was actually four different machines on the internet. As the NSFnet grew and this was still before we even called it the internet really, we expanded to seven root servers. And with 91, the first root server outside of North America, and as the technology became available in the early 2000s, we introduced Anycast. So this is what allowed us to scale from having 13 identifiers, those 13 addresses, we could have had 1700 addresses. That's just kind of a waste. Now we're actually able to scale in a different

dimension so that with those same 13 identifiers, we're able to have more than 1000 physical locations around the world.

Interesting document there, RSSAC023 discusses some more of the history of the root server system. And here's our current list of the 13 operators or 13 identities and their operators. So there's actually 12 different operators because one operator is responsible for two. Go to the next slide. And this is a way, this is root-servers.org website where you can go and you can actually look at the physical locations of where these root server instances are. You can zoom in and look specifically of where these things are. And one important thing, we'll talk about this shortly. We'll go next to the root server system principles.

Several years ago, RSSAC embarked on a process to define a governance system for the root server, for the root server system. Previously, originally, it was volunteers that decided to do this. We thought we needed a more formal structure. During the development of the original document, we came up with these principles of what it means to be a root server, what it means, what we hold as deep values. Go to the next slide.

I'm just going to call out a few of these principles. There's references to documents where you can look further. Some of the important ones are that IANA is the source of the root DNS data. So we got our data from one place. That's IANA. We all have the same IANA database for the root zone. Doesn't matter which of the 13 operators, which of the 1700 instances, they're all going to have the same root zone. We also consider diversity of the root server operations a big strength of the system. From the security perspective, we don't all run the same

software. We don't all run the same hardware. We're different organizations in different economic infrastructures, different political infrastructures that if any one of those would be threatened, if one industry collapsed, well, not all root servers are in the same industry that would take it down. So diversity is a very important point there.

RSOs do collaborate and engage with the ICANN group here in the multi-stakeholder community as well as with the more technical groups, IETF. So we're here. We're interested in collaborating with the rest of the community to understand the needs.

Another important piece is that RSOs must be autonomous and independent. So we don't control each other. We work together. But we are free to implement things independently, and that's what gives us that diversity, which is the strength. Next slide.

Okay. Yeah, so these are some of the things we actually hear around the halls of ICANN or other groups where misunderstandings about the root server system. So root server system does not control where traffic goes. All we do is provide addresses to locate the servers of the TLDs. People think that most DNS queries are handled by the root server. That's not true. Again, that recursive server is the one that takes the brunt of your queries. And because it caches, the root server system doesn't need to be contacted all that often. The administration of the zone, that's the separation of the IANA function that controls the contents. And we're just the service. We just provide the computers on the Internet to do this. We'll say this over and over again here. There are not 13 root servers. There are over 1,500 instances, 1,500 computers around the globe to do this service. And one thing about when you send

a query to a root server, originally the idea was you send the entire query, www.example.com. A lot of recursive resolvers are implementing something called Q name minimization now. So for the root servers, they know we're only going to be able to answer the top level domain. So they only send us .com. So we see less and less of the query. We have everything we need to answer the question exactly correctly. But we don't need to see that. That's good for your privacy. Next slide.

So the root server system and you, we say that we, if you operate a recursive resolver, you should typically have three to four instances that are nearby to you. What does that mean? Nearby, when we're talking about accessing a root server system or root server, this is topological, not geographic. So you just need a good route. And really when we talk about instances being nearby, the main piece of that is availability. It doesn't matter how far, what distance you go. In fact, how long it takes to get there really doesn't matter much, especially if you're only querying it a few hundred times every two days. The latency is not sufficient. We encourage everyone to use a DNSSEC validating resolver. That's just good practice no matter what, but you will get full use of it for the signed root zone. We also encourage people to participate in the DNS communities. The RSSAC has a caucus which has a group of DNS experts that can help us advise on the root server system, help us create documents and advice to the board. That group is open. If you have some background and technical experience that you want to join that group, there's information later on. There's even cases where you can request to host one of those 1700 locations. Make it 1701. Send an email to ask-RSSAC@ICANN.org. There's a couple operators that will work

with you and possibly help you run a root server instance in your network.

Another part of what we do is measuring the root server system. And this is actually one of the measurements that we publish. I believe we publish this daily. And this is just the query volume between all of the 1700 instances. These are how many queries we get per day. So we've kind of steady stayed around about a hundred billion queries per day. Which sounds like a lot, but for a computer, for 1700 computers, it's really not that much. The root server system is well over-provisioned. There's a lot of capacity. It's there when we need it. Next is Anycast. Liman.

LARS-JOHAN LIMAN: So thank you. I'll try to explain a bit about this Anycast technology that we use. And by doing so, I will take you through a way of thinking that might help you. In the beginning, if you dialed with a telephone, you would reach a single other telephone at the other end. Now, if you have a lot of people who want to call that place, it could be a large company, you cannot have a single telephone behind that telephone number. So you have a switchboard and a couple of people who receive the calls. And you might even expand that to a call center. But it's still typically a location, an office where this call center is located. But in a modern system, for instance, here in Europe, on your telephone, if you want to dial the emergency services, you dial 112. And if you are here in Hamburg, you will get to an emergency response center located probably here in Hamburg or near Hamburg. But if you travel to Rome in Italy or Stockholm in Sweden, you can still use the same number. You

don't end up at a call center in Hamburg. You end up at a call center in Rome in Italy or a call center in Stockholm where you happen to be. And that's a bit how this Anycast system works, but on the internet.

So where the DNS queries that are sent from the recursive resolvers, they have to reach name servers. And we are here talking about the root name servers. And from a specific resolver, the traffic will follow the traffic signs that sit in the routers on the way. You send out the packet on your link towards the internet, and the internet is full of these routers that will forward the packet to a next link and another link and another link and eventually towards the target. And these routers, they take instructions from their neighbors. They talk to each other and tell each other where to send these packets. So what we do is that we announce reachability. We place all these call centers around the internet and we say, you can reach us on number 112 or rather not the IP addresses for the various root server operators. Here is one and there's another one and there's another one. And the routers on the network will pick the nearest route to the nearest call center or nearest root name server.

This has a couple of good consequences. So for instance, if you look at Unicast, which is what we call the traditional way of calling between just two separate computers, all the packets will go to the same destination. You have that single telephone that you reach. But that means that if you have an attack where a lot of computers send traffic at the same time in order to block the service, they will hit the same server and that will not cope with the load. That's a simple fact. But with Anycast, with these many call centers spread across the world, they are much more resilient to these types of attacks because they will be distributed. The telephones in Hamburg will call the Hamburg center and the

telephones in Stockholm will call the Stockholm center. They will not all go to the same place. So this helps us to mitigate these types of attacks.

And it also gives quick response because you don't want to wait for your call or your packets going across the entire internet from Stockholm to Hamburg. You want to have a quick response from someone in Stockholm. And the wider we distribute all these servers, for every client, we get closer to that client and that means that we can give better service. Even though, as Ken said, the exact number of milliseconds isn't really that important, but it may make a difference if you are in Tokyo and my server is in Stockholm because that's 400 milliseconds, that's close to half a second for a packet to go over there, get the response and back again. If I have a server in Tokyo, that's four milliseconds. So that's a lot quicker. Next slide, please.

So this is how it works in the traditional case, Unicast case. A single source will send a packet across the network, across the routers and reach the destination. And it will take some version of the shortest path. But next slide, if we do Anycast, the service will be present at multiple locations and the blue dots here indicate that. And a source will reach only the destination that is closest. And again, close isn't necessarily a geographic measure in kilometers. Close can be number of hops on the network between different routers. But there is mostly a correlation between geography and topology in the network. Next slide, please.

And here we have a picture trying to explain that someone who attacks a single IP address here will only reach the nearest one because once the attacker sends the packet onto the network, the attacker can no

longer control the path on the network. That's done by the routers and they have their own opinions. The attacker cannot tell the packet to take a certain path through the network. So the routers will send them all to the nearest root name server and all the other instances, all the other blue dots will be unaffected by that specific attack. Next slide, please. I guess this is where you kick in, Wes.

WES HARDAKER: This is, but why don't I pause here? So that was the technical load of our presentation today. Is there any questions about the technical aspects? I'm going to in a minute get into RSSAC and how it works within the ICANN community and things like that. Is there anybody that had questions on the material so far?

EDUARDO DIAZ: Why 13 servers?

WES HARDAKER: So the quick answer is history. So remember that there's not 13. As Ken explained a bunch of times, there's actually over 1700. The size of the packet was actually the limiting factor that 13 was sort of the maximum that could fit. Before the Anycast technology that Liman talked about came along, the only way we could expand the service and scale it was by adding more addresses. This is no longer true. Now with Anycast, 13 is actually more than we need, but that's where we ended when Anycast got created. Does that make sense? And somebody else? Yes, sir?

UNIDENTIFIED SPEAKER: I wanted to ask about the Anycast providers. When I operate an exchange point, I get an Anycast provider like PCH. They install a root server instance on my network. I wanted to know the relationship and the responsibility that is associated with the root server operators for such an Anycast. How is the relationship like?

KARL REUSS: We partner with PCH, and they're our sort of Anycast provider, if you will. They provide the hardware and the interface to the exchange point, and they give us the facilities to run a root name server on their hardware. So we take care of the operating system and all of the various pieces here that others have described, and we run the root server along with others that we run on our own hardware. And different root server providers have different partners that they work with, and D, University of Maryland, works with PCH.

OSCAR GIUDICE: I wanted to know if the blocking is at the root level, when the attack is blocked, is it at the root level, or not taking advantage of the IP address corresponding to a site?

WES HARDAKER: That is an excellent question. There are many forms of attack on the Internet, and some are at the routing level, like you talked about, and some are at a DNS level. All of the root server operators have to be prepared to deal with any type of attack, and we get them occasionally. The reason we have scaled to 1700 different instances is to make sure that we can handle anything that comes our way. The other day, my

instance alone, I communicated with everybody else. We got a significant attack. It didn't take us out, but it was a lot of traffic. It ended up at all of our instances for my university. None of them got overwhelmed, but the rest of the operators weren't affected. That was an attack to a particular address. The way it was executed came from all over the world, so it went to more than one instance. But the type of attack and how the solution to that attack is very dependent on the type of attack. Hello. I'm

SHERIF KHALIFA: I'm totally aware about applying Anycast over OSPF protocol, but I still don't understand. I believe this solution will be applied via eBGP. Is that correct? So I'm a little bit concerned about how the route election will be made to choose the shortest path, and I just heard you say something about Anycast provider. So can you give us some more info about this?

LARS-JOHAN LIMAN: Yes, you are quite correct. You are quite correct. It is BGP announcements, and you can, when you use BGP, the Border Gateway Protocol for routing, you announce reachability. You say, I am here. You can reach me here. And you tell that to your neighbors across the link, and the neighbor is the next router at the end of the link that you are connected to. And when routers hear that, they learn and say, okay, Liman is in that direction, or that root name server is in that direction. If you say the same thing from multiple locations, these will spread as rings on water, and eventually they will meet somewhere, and that router will hear the same information from two directions and may be

confused. No, it won't. It will think that this is two different paths going to the same location. It doesn't know about Anycast. The router in the middle has no idea that we do Anycast. So it will see this as the main route, the main path to the target, the route name server, and the backup path. And the BGP algorithm, it's a very strict algorithm how to make a decision about which path to take. And that algorithm will tell you at some point in there, it will say, okay, this is the path to go, and it will select one of the routes. And the other one will remain as a backup. So if the first one disappears, maybe we stop our server in Stockholm, and you can no longer reach it. The router nearby says, okay, that path to the target is down. I'll use my backup path instead, and they will end up in London, not knowing that it's a different location. So that's one of the big benefits of Anycast, that we can take our machines out of service without even telling you. And you won't notice, because the entire network will say, oops, that one's gone. I'll use my backup path instead. And it will just continue to work. So yes, it's BGP, but it's not problematic. It's designed to handle this type of thing. But the routers in the middle have no idea that we are doing this. Thanks.

WES HARDAKER: One important point that I don't think our slides actually talk to at all. Anycast is not a technology specific to the root server system. Many other DNS servers use, including COM and NET and ORG and all the TLDs pretty much, and other protocols as well. The time synchronization protocols and some of the other protocols. Anycast is just a routing trick. It's not specific to the root server system at all.

YAZID AKAHNHO: And maybe also before reading the online question. Anycast doesn't also prioritize a root server operator or a root instance against another one, right? It's a technical configuration.

WES HARDAKER: We should take a couple more and then we should probably go on.

YAZID AKAHNHO: So we have one more question online. But okay, let's go back to the room.

EDUARDO DIAZ: The graph that you have that IANA is the one that enters information and then spreads out. You know, that's a person doing that, I believe. And I always wonder what kind of security process in that point where you entered the information. What happens there? Because if I enter the wrong information, that will spread out everywhere. So can you explain how that little part works?

WES HARDAKER: So we don't have time to get into routing security today, but there are mechanisms to help with that. I've given lectures on that, but it's another complex subject for how to make sure that the routes that you're advertising get accepted and that others do not. Go ahead. Yeah, please.

LARS-JOHAN LIMAN: I also would like to add that we use something called transaction signatures when we import the data into the servers so that we have crypto signatures on the stuff that we import. And we will recognize that, oops, something is not right here, and we will stop importing that data and go back to the earlier version. And on top of that, we are also introducing a system with the checksum for the entire data set, the entire zone file, which is becoming operational probably in the course of the next year.

EDUARDO DIAZ: I really wanted to know about the content.

LARS-JOHAN LIMAN: Okay, generation of the content. Okay, I don't think that any one of us up here is able to speak much about that, but it was Duane in the room? IANA maybe.

WES HARDAKER: The root zone maintainer, Duane, is back there.

DUANE WESSELS: Yeah, I'm not from IANA. I can say that the data that enters the root zone manager system comes from the TLD managers. They're the ones that type it in and when they make their request to IANA. And then from IANA to the root zone maintainer, this all happens over EPP protocol, and there are numerous technical checks to make sure that the data that is provided is correct. For example, if there is a change that involves the name of a name server, there are checks that make sure that that name

server exists, that it's serving the data that it should be serving, and so on. So there's quite a few checks on the data that goes into the root zone to make sure that it's valid and correct data.

UNIDENTIFIED SPEAKER: That's an IANA function, and between IANA and the root zone maintainer, we get the data. So a very valid question, and I think there's extensive documentation about how IANA serves that, but the important distinction is that that's outside of the root server system itself, where sometimes those get confused.

YAZID AKAHNHO: Thanks for your good questions. I will just ask you to keep the remaining ones. We need to read the two questions that are online and complete with the slides and get back to you guys. Okay, so the first online question is from Anupam, a follow-up on the earlier question on relationship between the root server operators and PCH. Does the instance of 1,700 instances include the PCH nodes as well?

UNIDENTIFIED SPEAKER: Yes. Yes, the PCH nodes are included in that count.

YAZID AKAHNHO: Cool. Another question from Herve. That's in French. Is ICANN doing anything to support the ISPs on BGP skills enhancement?

WES HARDAKER: So ICANN does have an educational section that does outreach for things like that. That is beyond the scope of our expertise. So I'm going to have to have you defer that question to a different portion of ICANN.

YAZID AKAHNHO: Cool. Thanks, Wes. So here we are talking about, of course, the root server operators and the question is related to ICANN. So coming from the technical engagement team, we do, as part of our technical engagement activities, we do talk about BGP, of course. We do talk about Anycast, definitely, but in relationship, of course, to DNS. But BGP is a global routing protocol. And sometimes we just divert this to the RIR, who has more skills on talking about routing and talking about all of those topics related to IP. But yes, we do talk about this in our DNS classes. Thank you. So back now to the presentation. We still have 30 minutes left. Wes?

WES HARDAKER: And there will be time for more questions at the end of this, too. So do hold your questions, and we'll be happy to answer as many as we can. And you can also go find us in the hallways. We're friendly people. So I am going to talk about the RSSAC and the RSSAC caucus. So what is the RSSAC? The RSSAC is the Root Server System Advisory Committee. And our purpose is to advise the ICANN community and the board on matters related to the root server system with respect to its operation, administration, security, and integrity.

Note that this is very narrow. The RSSAC does nothing except talk about the root server system. So we don't produce advice on the DNS. We

don't produce advice on routing. Unless it's directly related to the root server system. So it's a very narrow scope. We don't do policy development at all. So the PDPs, that's not us. We're purely an advisory committee.

So we primarily write our advice to the board, but we write it to anybody that we feel is good. Sometimes we write general RSSAC documents to just help explain things. The history of the root server system document, which is RSSAC 023, is a good one. It's not really advice in there. It's just the history of how the root server system was developed. It's a good one to answer a lot of the questions that some of you have had about the history of how it grew to this size.

All the root server operators are represented within RSSAC, but RSSAC itself does not actually talk about operational matters. Within the RSSAC, we don't talk about, hey, where does instances go, or how to do routing. We're very technical people as part of it, but that's not the job within the advisory committee of ICANN. We don't advise the operators on how they should run their service. We talk about how the whole system as a whole should function with respect to ICANN. Next.

We have two chairs right now, or a chair and a co-chair, I should say. We used to have two chairs, so I'm speaking from a historical viewpoint. So Jeff Osborn is our current chair. He is not here today, unfortunately. And he is from the Internet Systems Consortium RSO. And then you've already heard from Ken, who is the RSSAC vice chair. He is the U.S. Army DEVCOM Army Research Laboratory operator.

RSSAC is composed of some primary representatives from each of the root server operators. I'm the primary representative from my

university. We also have an alternate representative in case the primary isn't available to help make decisions. And we have incoming liaisons from some external bodies, and I'll talk about those more in a minute. And then we also have the RSSAC caucus, which is a body of volunteer subject matter experts. And that's where anybody that has expertise can kind of come help RSSAC craft our documents, which are usually technical in nature. And those members are confirmed by RSSAC based on a statement of interest.

Talking about the liaisons first, we have eight liaisons, essentially. There's four incoming and four outgoing. We have an incoming liaison from IANA, from PTI, which is James Mitchell. We have a liaison to the root zone maintainer, and that's Duane Wessels, who spoke earlier in the back. We have a liaison from the IETF, which is also the Internet Architecture Board, provides that liaison. That's Daniel Migault. Then we also have a liaison from the ICANN's Security and Stability Advisory Committee, and that's Russ Mundy. And then we also have a number of outgoing liaisons, liaisons from RSSAC members to other elements within ICANN. So I am the liaison to the ICANN Board of Directors, so I sit on the Board of Directors. We have the liaison to the Customer Standing Committee, and that is actually Ken again. We have a liaison to the Root Zone Evolution Review Committee. That's RZERC. That's Daniel Migault. He's both an incoming liaison and an outgoing liaison. And then we have the liaison to the ICANN Nominating Committee, which is Hiro Hotta at the moment.

The RSSAC Caucus is an expert body of people that we try and get to help us. We don't want the root server operators to really be the only people thinking about the problems of the root server system, so we try

and build a body of experts that can come help us, and we currently have over 100 people in the RSSAC Caucus. They are DNS experts, or at least have a very strong background in it. They write public statements of interest of why they want to come help, what their background is, you know. And then we give them credit, so when they help write an RSSAC document, their name goes into the authorship list. And again, the goal is to really bring a diverse set of expertise to the publications that RSSAC writes. We don't want to write them just from the RSSAC itself. We want to write them from the expertise of the larger community. It also gives transparency, so everybody sees how these documents get written and the content that goes in them, and it's the framework for getting work done.

The root server system governance evolution is really describing sort of we're in a transition point, an inflection point, of how the root server system is being governed. Back in 2018, RSSAC published RSSAC 37, which is a fairly well-known document, where we did, as a group, get together and describe, here's how we think the future of the governance system for the root server system should happen. Because there is not actually a process for adding and removing root server operators. The last person that made changes unfortunately passed away, and we didn't sort of leave a follow-on and set of instructions for here's how to make decisions about this in the future. So ICANN's in the process now of coming up with a way of building a framework for having these discussions and for making changes to the structure of the root server system as a whole. That is being done within the root server system governance working group, which is the GWG. It is meeting all day one day. Later this week, we have a lot of meetings, very actively

under engagement right now. It's working on the same 11 principles that Ken talked about earlier. It's sort of our starting principle of this is how the root server system should function and this is how RSOs should. These are our guiding principles. It includes representation from not just RSO operators. It includes representation from many other stakeholders within ICANN. I'm going to get to that in one second. The eventual model will go through the ICANN publication process, so everybody here and everybody within the ICANN community will be able to comment on it, of course. They are open to observers, so you can come listen to the sometimes very dry meetings, but you are welcome to come listen anytime you want to to the GWG meetings that are open and recorded as well for later consumption.

One of the things that RSSAC 37 that started the whole process of what became the GWG, we spent a lot of time thinking about who are the stakeholders of the root server system, who actually are the end beneficiaries of it. That really is the ICANN community is the obvious one. The RSOs themselves are really stakeholders in it. Then the IETF and the IAB. I actually said on the IAB as well, but the IETF is the body that controls how the DNS protocol itself works and is designed. That's not within ICANN. That's actually within a completely different standards organization that defines how the DNS works.

The root server system governance working group, the GWG, is made up of members from a lot of different organizations. There's a liaison from IANA. I'll go through the liaisons on the blue on the right-hand side first. There's a liaison from the root zone maintainer, and there's two liaisons from the board, from the ICANN board. Then there are members from, there's a member from each of the RSOs. There's two members

from the registry stakeholders group, two members from the ccNSO, and then two members from the, appointed by the Internet Architecture Board on behalf of the IAB. So the IAB appoints them on behalf of the IETF, I should say.

So if you have more questions, A, we're going to take questions in just a minute, and then B, there's a lot more information on the web about everything that we've just talked about today. RSSAC has a main webpage, of course, within the ICANN webpage system. We have a frequently asked questions list, because we get a lot of the same questions over and over, so it's a good way to even increase your knowledge, is just go read the frequently asked questions list. And if it's not answered there, we do have an email address, which is ask-RSSAC@icann.org, where we try and answer questions as they come in. And then for more information on the RSSAC caucus, there is, of course, a webpage about the RSSAC caucus, and then you can apply to it by sending email to the RSSAC-membership@icann.org if you are interested in helping us craft these documents. Again, we have a strong expectation of a technical background in how the DNS works, or at least how the root server system works, but we really do want a wide diversity of help in doing that. So I think that is the very last slide, so we can go back to questions. So questions that you didn't get answered earlier about technical content, or questions about RSSAC and the caucus and the membership and with the ICANN community.

YAZID AKAHNHO: Thank you so much, Wes. So back to the audience for questions now. Oh, okay, yeah.

NIKLAS: So this is Niklas from .se. Are there any statistics about how much of the traffic that hits the root zone that are IPv4 versus IPv6?

UNIDENTIFIED SPEAKER: Yes. There is a... I'm going to have to find it and type it up there. I think it's RSSAC002.root-servers.org. We publish those statistics in depth. Andrew did a great job putting that together. V4 versus V6, TCP versus UDP, among all the different operators. Lots of interesting statistics if you're a computer geek like us. That's published and you're welcome to take a look at those.

WES HARDAKER: Specifically, one of our RSSAC documents, 002, it's technically our third document because 000 was the first, but actually describes the format of all that data and how all the operators publish it so that you can go get it on a regular basis and analyze it. There's archives on GitHub that you can pull it all in. So if you don't like our analysis or graphs of it, you're welcome to create your own. If you think something's missing, we also continually revamp that document as well. I think we're up to version 2 where, again, 0 was actually the first.

UNIDENTIFIED SPEAKER: I wanted to know if there is a list somewhere for the Anycast operators. Is there a list somewhere?

LARS-JOHAN LIMAN: I would like to oppose a bit to the term Anycast operator. There is no such thing. There are root server operators and some of us operate our own routers that connect to the internet and we do the BGP announcements there. Some other root server operators cooperate with other groups and PCH is one such example but there may be others who do the routing but they don't necessarily determine... Anycast operator is nothing that fits into our model here, but you can operate with someone who provides a number of sites for you and you can cooperate with a company that says, oh yes, they will provide 25 sites and myself I will provide another 25 sites and they will all be part of Anycast cluster of machines and sites that work together. So would you please try to not use that term, because sometimes these terms hook into people's memories and then we have to spend a lot of time trying to denounce this and say no, that's not how it works.

WES HARDAKER: I think one way to think about it is all of the root server operators, we control our routes from all of our instances and we have partnerships with organizations that are hosting us physically and with the networks. So even in one of my sites, I often have routes to a bunch of other different places. So all of my advertisements are Anycast but I'm the one actually doing them so there's not an Anycast provider, it's me controlling the routes for our particular set of instances.

YAZID AKAHNHO: I think it's also maybe time to remind the audience that there is a lexicon published by the RSSAC, RSSAC 026 I think, which provide a broader terminology about the root server system and it's also publicly

available. Thank you. Any other question from the Zoom so far? No other question? Okay. Thank you.

ROSEMARY SINCLAIR: I'm interested in the relationship between RSSAC and IETF and how those processes and that knowledge sharing I guess works. Thank you.

WES HARDAKER: I'd be happy to speak to that. I'm integrally involved in all sides. So the IETF is where the protocol is actually created. It's actually even how the root server system is technically is defined in the IETF. ICANN maintains the IANA database of who the root server operators are, but how it works is entirely defined within the IETF. There are many people that play in both worlds. I attend all of the IETF working groups as does my colleagues here too because we care about how the DNS protocol works even though it's not part of ICANN. And so they define when you need to send a question to the internet this is how you do it and you put this zero before this one and it gets very, very technical.

The IETF has many liaisonships with other organizations including ICANN. The Internet Architecture Board is one of the leadership bodies that sits above the Internet Engineering Task Corps which is the IETF. And the Internet Architecture Board is responsible for appointing those liaisonships. So there is a couple of appointments from the IAB, the Internet Architecture Board, to bodies within ICANN. Including we have an appointment to RSSAC which is Daniel Migault that I mentioned earlier. We have two appointments to the GWG which is Jim Reid and Geoff Huston. And I think there's an appointment to the ICANN board

actually which is Harald. There's an appointment to SSAC. I don't know. So there's enough that we try and keep a very strong technical liaisonship because it's critical that both sides understand what's happening in both communities. Yeah, please.

LARS-JOHAN LIMAN: I would like to add that I, as Wes, participated in just about all IETF meetings not only to gain information and knowledge about the DNS, what's going on there. It's actually the fact that the root server operators have technical coordination meetings three times per year at the IETFs. These are not public meetings. That's where we synchronize so that our systems work or function in the same fashion everywhere and make sure that everything is adapted to and when we have new things we take them on board in a careful manner and stability, stability, stability is the mantra that we have. But I also go to the IETFs to pick up what's going on in other parts of the internet technology space. How is the web evolving? How does that affect the DNS service that I provide? And it's important for us to be part of all these discussions so that we understand and can adapt our systems so they continue to function well. And also in a few cases to say no, that's probably not a good idea because you didn't think of how the root servers will react to that. But it's a very tight and very well functioning interaction between the DNS technical community here. Here we interact with the policy side of things. In the IETF we interact with the technology things. So it's all very tight.

WES HARDAKER: I should say that the Anycast definition is actually an IETF document. That has nothing to do with ICANN. It actually came as well as the routing and everything else. A lot of the same people exist in both. We wear very different clothes. I wear a very nice shirt here and a very nice jacket and I go there and I wear a T-shirt. So it's a very different set of community.

YAZID AKAHNHO: Same for the transaction signature and the zone MD as well, right?

WES HARDAKER: Yeah, exactly.

YAZID AKAHNHO: Other questions? So far no new questions from the Zoom. Okay.

UNIDENTIFIED SPEAKER: When people talk about fragmentation of the internet, they often talk about people creating their own roots and so forth. I wanted to know the view from the RSSAC on this issue.

WES HARDAKER: I was just having this conversation earlier today too. I would say that everybody in almost everywhere, both on the technical side within the IETF as well as within ICANN believes that end users find it easier when one name maps to one thing, right? You don't have a conflict where one person is using one name and another person is using another name. The IAB, the Internet Architecture Board, actually has a document. I

think it's document 2826 that actually talks about this. It's actually quite old. It talks about the need for a globally unique namespace. When we have other naming systems that are coming in and maybe they want to do something different or maybe they want to do something better, we really need to work hard to find a way to interoperate, right? I personally don't care if they use a completely different protocol to figure out where they need to go. If they don't use the DNS, that's great. What I worry about is, is there going to be a conflict in those names where, you know, one name that looks like a DNS name actually is used for something else and it gets somewhere different than a DNS name that gets somewhere? So one of the things that we as the root server operators have to deal with are when systems believe that they can look it up something with a different name and that other naming system doesn't exist in that computer, then the system actually thinks everything's in the DNS and it gets sent to us and we can't answer it because it's not part of the DNS, right? They sent us a request that is not part of the DNS.

So the high-level answer is we want one naming system. It gets pretty chaotic for end users if they type a name into their web browser and they tell their friend to type a name into the web browser and they don't go to the same place. Chaos that way is [inaudible]. So I think everybody in RSSAC would back me up on that. I know everybody in the IAB would back me up on that and I know everybody in the IETF would. There are some people that may not fall to that same belief that they don't mind having two different names.

LARS-JOHAN LIMAN: And from the RSSAC perspective in our principles, our guiding principles, we serve the IANA root zone. That's what we're here for. So recognizing that those other namespaces might exist, but our job is the IANA root zone.

ROSEMARY SINCLAIR: Just on that mention of principles, we seemed to glide over that slide very, very quickly. I wonder if we could just go back to it given that we've got just a few minutes and go through them.

LARS-JOHAN LIMAN: Yeah, so one of the things I'll point out is that those principles were part of one document. We thought that they were so important we actually split them out and discussed them further in a second document, which is at the bottom of that page. RSSAC 58 is the document. Page 17. Yeah. There. 55 is these 11 principles spelled out with text. So these are principles that we found ourselves are guiding how we operate and cooperate between the root zone operators. And you will find in here that cooperation and stability are very, very important things, and that's how we see our mission. We don't have to run the latest software or feature in the DNS. The paramount thing is that this continues to work and operate. And we try to create as many aspects of resiliency into the system as we possibly can. And you will see a couple of them hidden behind the various labels here where one is diversity. That's kind of our motto. Diversity is important. We are diverse organizations. I work for a small company in Sweden. It's a private company owned by a foundation. You work for the U.S. Army. What do these organizations have in common? Almost nothing, except that we operate root servers.

So if the Swedish legal system would break down, Ken would continue to operate, and you wouldn't notice if we totally disappeared as an operator from the network, because that's how different the system is. We also operate different operating systems. We have different hardware. We have different DNS software. We have different governing principles in our corporations. We have different legal legislations around us. So we try to be different in every way we can, except one. And that's how we provide the technical service, because you shouldn't be able to tell whether you talk to my server or Ken's server. You should receive exactly the same response, but provided in different ways that create this resiliency. So it's not just technical resiliency. It's also administrative, legislative, financial. I provide the money for the service that we provide at NetNod in one way, and you do it in a totally different way. So even the source of money is diverse. That's how we try to create this stable network. And this diversity only works well if we are transparent about what we do. So if you read these principles, you will realize that this all boils down to stability and resilience. That's really what we want to create. And to some degree, performance. But that's actually part of resilience, that we can withstand attack and provide a servie that is quite sufficient and hopefully very good for the entire Internet community to use.

WES HARDAKER: I was going to add one more thing because I was reading through these, and I just noticed that many of the questions you've asked today are answered in the principles. Principle number one, to remain a global network, the Internet requires a globally global network. And that directly is talking about some of the things. Principle number six. The

IETF defines the technical operation of the DNS protocol. Everything that we've talked about here today really comes straight out of these principles in many ways. And RSSAC 55, it goes through a lot more text than we have in remaining time to explain them. Good question. Thank you for that.

YAZID AKAHNHO: Yeah, I think we're getting very good questions. Any other questions from the room? Online so far? No new questions online?

LARS-JOHAN LIMAN: Can I just remind people that a lot of Root Server operators are here at this meeting and again, we are friendly people. We actually like when you come and ask us questions because this is our passion. We want to tell you about it. So please come to us and ask questions afterwards. And if you come up with a question tomorrow afternoon, grab me in the hallway or Wes or Ken or Karl or any other one. We are here to communicate with you. And a good way to do that is if you have questions.

YAZID AKAHNHO: Cool. Thank you, Liman. Let's wrap up. What could be the future of the Root Server system?

KARL REUSS: Yeah, yeah. I wasn't ready for that one. I think just the new governance system is probably the largest thing coming up on our roadmap. Slowly, we will evolve.

YAZID AKAHNHO: Thanks, Karl.

WES HARDAKER: I think Karl nailed it, right? That's the important thing. Probably the second thing is the interoperability with any other naming systems. I think that that will become more prevalent as time comes on in the near future as well.

KEN RENARD: I'll state the obvious. We're just going to keep working. We're just going to keep answering queries and do that reliably and stable.

LARS-JOHAN LIMAN: Yes, and I will add that the new round of gTLDs, it will fit into the system and we need to be part of the discussions around that so we understand what the future requirements on the Root Servers will be so that we can build the system to handle that. If we can only maintain that dialogue in a good way, there won't be any problems.

WES HARDAKER: I'll end with one quick comment about the next round. The Root Server operators have said we've been ready for years. Whenever ICANN gets done with the process point of view, we're already ready. We have been for years.

YAZID AKAHNHO: Thank you so much. We are at the end of the session. Thanks, everyone, for those who are in the room as well as those who are online. Thanks for your questions, but also thanks very much for our presenters. I think they really deserve a round of applause if you don't mind. And thanks for the technical support. With that said, we are at the end of the session and we can close. Thank you.

[END OF TRANSCRIPTION]