ICANN81 | AGM – How it Works: Root Server Operations
Sunday, November 10, 2024 – 13:15 to 14:30 TRT

ULRICH WISSER: Hello and welcome to how it works, the root server operations. My name is Ulrich Wisser, and I'm the participation manager for the session. Please note session is being recorded and is governed by the ICANN expected standards of behavior and the ICANN community anti-harassment policy. During the session, questions or comments will only be read aloud if submitted within the Q&A pod. If you'd like to speak during the session, please raise your hand in Zoom. When called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use their physical microphones to speak. Please state your name for the record and speak at a reasonable pace. All are welcome to use the chat. Please note that private chats are only possible among panelists in the Zoom webinar format. Any message sent by the panelists or a standard attendee to another standard attendee will also be seen by the session host, co-host, and other panelists. I will now hand over the floor to Ken Renard.

KEN RENARD: Thank you. Good afternoon. My name is Ken Renard. I am one of the root server operators from the Army Research Lab in the US. And today we're going to present about the urge people to ask questions. It's going to be very interactive. We're glad that you're here. We're glad that we can share how things work. I do have a clicker. So briefly, what we're going to be talking about today is a few of DNS, because that helps us

explain how the root server system fits into everything. That's kind of the status of what the root server system looks like today. Quick explanation of Anycast, and then get into us, and then a little bit about their governance.

So we start off with the fact that the IP address is really the fundamental identifier on the internet, either with these numeric or strings of 1s and 0s at these addresses. So if you're connected to the internet, you've got an IP address. So the problem here is that we as humans are not very good at remembering these numbers. And these numbers can change a lot. And that's a good thing. So this was the reason originally to develop the DNS is to bring numbers and the fact that they change. The more modern part of that is that IP addresses are shared sometimes with multiple websites on the same IP address. And the complement to that is the fact that a single address, like a website, can actually be hosted on multiple IP addresses. So DNS is our friend. It makes things a lot easier for us and is used quite heavily. We have the namespace here. We're really using DNS to look up names into those IP addresses. So we have this hierarchical namespace that starts at the top there with a root. And underneath of that are all of our TLDs, or top level domains. These are country codes. These are gTLDs. Under that second level, under that third level, et cetera. So the probably most common use of DNS is to look up a name into an IP address. But there are other mappings, such as looking up mail servers for a domain. That's how mail works. There's often addresses where you use an IP address to look up a name that exists. There's also a lot of security information stored in the DNS that helps us implement web security. Looking at the DNS, great statement there at the bottom. DNS is a globally distributed, loosely coherent,

scalable, dynamic database. So think of it like the phone book. Spend some time with this. This is the process of looking up a name in the DNS. This is the classic diagram you'll see in almost every textbook. And it's useful for describing the basic things. I'm going to add a layer on top of it once we're done.

So looking at the process, on the left side there, you have your mobile device, you have your web browser, you have your computer, your refrigerator, your toaster, your light switch, anything that's connected to the internet is going to act in this role to look up names into IP addresses. So those names are looked up by your device with a recursive resolver. That's that middle box. These recursive resolvers are the workhorse of the global DNS. That's where most of the fun happens. Most of the action happens there in the recursive resolver. So your device is going to ask the local recursive resolver for the address of, here we have www.example.com. And that recursive resolver is going to go find out what that address is. We describe this in the process here. The first step is it's going to, where is .com? So the root server system is only that top right box. So it's going to ask the question. The root server system's going to say, hey, I don't know what www.example.com is, but I do know where to find the .com database. That's going to get back to the recursive resolver. Recursive resolver is going to say, now that it knows how to get to the .com database, it asks the same question. And .com pretty much responds the same way. I don't know this address, but I can tell you where example.com is. That gets back to the resolver, and then now the recursive resolver can ask the example.com database for the answer, and it gets that.

The next really important part of this is that along the way, that recursive resolver has remembered each of the answers that it has got. Your device asks again for that address for www.example.com. It already knows the answer, and it doesn't have to do this three-step process. If your device asks for q.example.com, it can skip the first step going to the root server system. It can skip the second step going to the .com server, and just goes to the last one. So there's incredible efficiencies built into this process. That information is cached or remembered for a certain time period. Specifically for the root, the values are 48 hours. So for all the times that you ask for something in .com, your recursive resolver has to go maximum—Another step of efficiency is that you're not the only one using this recursive resolver. You and tens or hundreds or thousands of your neighbors on the internet are using the same recursive resolver, so anything that someone else has looked up before may already be remembered. So your device could ask thousands of DNS questions, and never does it actually go out to the root server system for your transactions, because everything could be in memory. So incredible efficiencies there, and see here that the root server system, that root authoritative server, is very important because it's the start, but it's also not used all that often. So we'll go to the next.

A little bit more about caching and resolution. So every DNS record has a time to live value. That's how long that recursive resolver is allowed to remember that answer before it's deemed old enough that you have to ask again. Important piece is that the root servers only serve information about the top-level domains. It only knows where .com is. It does not know where example.com is. Just the top-level domains.

Right now there's 1,450 top-level domains. Very small database that we have.

Some modern refinements now to DNS. These are basically, these are not specific to the root server system, but they're things that help out DNS quite a bit. So DNS security, the DNSSEC extensions, we can now cryptographically sign the DNS data. Eliminates, or security people will say, reduce the risk of spoofing. And that intermediate server, that resolver, recursive resolver in the middle, has the ability to do the cryptographic check sums and verify that is actually correct. So it doesn't really matter where you get your DNS information from, because you can actually verify that it is correct and hasn't been tampered. There are also privacy enhancements. Now these are, the fact that you're asking for www.example.com, we saw that that question got sent to the root, it got sent to the .com servers, as well as that recursive. There could really be some sensitive information in there. Now from the root server system perspective, we do not see your device's address, we only see the address of that recursive resolver. So we don't [inaudible] to a specific machine or device, only the recursive resolver that they're using. There are other protocols like DOT, DNS over TLS, DNS over HTTS, DOQ is DNS over QUIC, that actually encrypt the transfer of the messages here. And those at this point operate between your device on the left side and that middle server, the recursive resolver. It remains to be seen exactly how they will be used in that second step.

Alright, a little bit about the root server system today. Just some definitions here, the root zone. That's actually the list of TLDs and their

name server addresses. So that's the content of the root zone. We as the root server system roll over that. We'll see that in a bit as well. The root server system is collectively the set of root server operators and all the infrastructure that each of the operators deploys that are the computers on the internet queries. The root server operator, you'll hear this a lot, the RSO, is actually the organization that maintains this and operates each of these root server—a little bit deeper. The root server Anycast instance is really one network location. So we have 12 root server operators, 13 root server entities, and over what, 1800 Anycast instances. So each operator on their IP address operates multiple computers all around the globe.

The Root Server System Advisory Committee is the ICANN advisory committee that is comprised of root server operator appointees and some liaisons. And we provide advice to the board and to the ICANN community about root server operations. To look at this a little bit further and kind of split out the idea of the root zone versus the root server system. Again, the root zone is the data. The root server system is the computers that operate on the internet that you can query and will send you a response and answer those questions. So very small part of the overall system. Again, very important, but these separations are really on purpose for security reasons, for separation of duties, so that clean roles and responsibilities.

So this looks a little bit more at root zone administration. So if you are a TLD manager request, you may have to change your name server address. That goes to IANA. That's where those things are changed. The actual database is updated. File does all the signatures, and that's

where the root server system begins. We get that table that's already signed, already locked in, and we just manage it. We send that file out to all of our computers, all your DNS resolvers on the right side query our system. So we're just a small part of the total root zone administration.

We see going back to 1984, so 40 years ago, the DNS was born, and there were four root servers at that point. And as the internet grew, the net in the U.S. as well as some others, there was the need to expand this to more. We had seven root servers. We see a lot of this growth happened in the 80s and 90s. So we have root servers, not only root server operators that are based around the world. Each of those operators have computers that exist all around the world. Over time, we respond to technical demands. This is a technical service, and we hold some principles that are very important that this system has to work there. It always has to work. So we're responding to technical demands as the technology, as the demand, the queries increase over time. Scaling issues. A long time ago, 13 root servers actually meant 13 physical machines. Now 13 root servers means over 1,800 machines. So Anycast has allowed us to scale the system in a new dimension that may serve a much larger community. If you're really interested in some more of the history, RSSAC 023 has some interesting things there.

So root server identifiers and their operators. These addresses are all public. We all operate on IPv4 and IPv6 addresses, so we've been doing that for a while. There is a website that you can go to now, root-servers.org, and there's a nice little map you can zoom in and look around and find these 1,800 instances where they're located. Now

that's not necessarily the same as where they are topologically located because the Internet is not exactly the same as geography, but you will find a really wide distribution of where these servers are. You don't need to have one physically close to you because it's the Internet. You can reach any one. You can, from here in Istanbul, you can talk to any of the root servers and they'll share that. It may route you to one that's close by. It may route you to one thousands of miles away, but you're going to get a response.

Developed a set of core operating principles in the RSSAC. This is probably too much to read. I'm sure you can look at this in a slide deck online. I want to call out just a few of them. These are all things that we hold near and dear to our hearts, and we're calling out principle number two that IANA is the source of the DNS root data. So that's where we get our data from, IANA via the root zone—One Internet, one root zone. If you ask any of those 1,800 servers on the Internet, they're looking at the exact same database.

Diversity is a strength of the overall root server system. So we have diversity of hardware, diversity of software, diversity of organizations. If there's a failure in a particular vendor of hardware, it's not going to affect the entire system. The same thing with software. If there's a bug in one particular software, it would only affect a small portion of the overall system. The same thing with our industries. We have government agencies, we have for-profit organizations, nonprofits, educational institutions. So the failure of any particular industry is not going to affect root server operations. RSOs, as root server operators, we do collaborate and engage a lot with each other and with the

stakeholder community. That's why we're here. We work quite closely together on some of the tech issues as we evolve and maintain the root server system. But we're also anonymous and autonomous and independent so that we don't control each other. We don't have any authority over each other, but we collectively do this work.

A few myths about the root server system that we hear quite a bit of control over where the internet traffic goes. That's certainly not true. Routers do that. We provide one part of the name resolution process. Queries are handled by a root server. We saw in the earlier slide about that recursive resolver and that caching. Very few DNS queries actually result in the root server system being contacted. The contents are controlled by IANA. We're just the computer operators. We're just the IT department of the root zone. All of the letters, they're lettered A through M, there's no specific meaning to those letters. They're just basically randomly assigned. On the earlier slide, this last myth about the root server instances receive the entire DNS query. If you query for www.example.com, that entire question comes to us usually. We see the recursive resolvers asking what's the address of www.example.com. We know that we're only going to be answering the where is .com question. There's actually a technology out there called QNAME minimization that allows the recursive resolver to reduce that query so that what we see as the root server system is not what's the address of www.example.com. We have a little bit more privacy preserving that your recursive resolver could implement.

Some recommendations on use of the root server system. If you operate a recursive resolver, you can do things like check the routes to your

instances. So you can do a trace route or something to find out where each of the instances are. And you typically want three or four that are sort of nearby. We don't have a precise definition of nearby. It's usually measures or latency. We recommend that you do use a DNSSEC validating resolver just in case anybody else was to [inaudible] you have the ability to verify that that information is correct. We recommend participating in the DNS technical communities. The RSSAC has a pretty limited membership of root server operators, but we also have an RSSAC caucus, which is open to DNS experts. And that's where a lot of our work in RSSAC, is done in this caucus. So if you would like to participate in that, you can site and look for the caucus. There's information on there how to request membership with that. And some root server operators actually will request for the ability to host an instance of their root server in your network somewhere. So the way to do that is talk to one of the RSSAC members here or root server operators here and send an email to ask-RSSAC@ICANN.org. That will get to us as well. Quick look at this. [inaudible] a lot of statistics about usage of the root server system. So this is just looking at queries overall, essentially an averaged queries per day. And we're talking about around 100 billion queries a day that we collectively see. Sounds like a lot. Most of them are garbage and misconfigured things, but we definitely see way overprovisioned. There's a lot of capacity there. Alright, with this I will turn it over to Liman to talk about Anycast.

LARS-JOHAN LIMAN: Hello, my name is Lars Liman. I work for Netnod, which operates one set of root servers. And I will try to talk a little about Anycast. So Anycast is

actually not related directly to DNS. It's a network trick that we use to be able to share the load between multiple servers for every operator of the root system. So all these 1800 servers that Ken mentioned, they are grouped in different size groups between the various operators. So Netnod operates roughly 75 or 80 of these, out of this 1800 total nodes. And all servers operated by the same operator, by the same IP addresses. It has two, one for IPv4 and one for IPv6. And normally you cannot have two computers on the network that have the same IP address. But Anycast is a trick that allows you to do that in the way that we have identical copies.

I will make an analogy here. If you travel around Europe with your mobile phone, you can dial... If you happen to end up in emergency situation where you need an ambulance, you can dial 112 from your phone. And you will get to an emergency response center that will help you to dispatch an ambulance to you where you are. If you are in Amsterdam, Netherlands, and dial 112, you will reach an emergency response center near Amsterdam. But if you travel to Rome in Italy and dial 112, your telephone call will not go to Amsterdam. It will go to a dispatch center in Rome or nearby Rome. And this is the same trick we use but in Internet terms. So if your resolver, the computer that assists you with looking up things on the Internet, if your resolver is located in Amsterdam and tries to reach a root name server on a certain IP address, it will reach a root name server in or near Amsterdam. If your server is located in Rome using the same IP address, it will reach a server in or near Rome. And these were only two examples. Now we have servers in 75 locations and the other operators have in maybe 50 or 200 locations also around the globe. And sometimes we have them in the

same location and most cases we have them in different locations. So there is this entire network of servers and there's always one close to you. So by using these specific numbers for reaching the root server, you can always reach one nearby.

So I will try to illustrate that with a few slides. But before that, I will also stress that this gives us a couple of interesting features. No, I'll go through the slides first actually. So next slide please. Thank you. On the Internet, Internet packets try to take the shortest path between two computers. And typically it will have to go through a number of routers or I will ... I slightly wrongly use the term network switches that kind of forwards the traffic between various links and then tries to forward the packet closer and closer to the destination. And in the normal case when you don't use Anycast, you go between two unique IP addresses. The sender address and the recipient address and you go towards the recipient address. So all the routers will work together to send your packets towards the end, the destination, the server that you want to reach. And DNS queries on the Net are normal Internet traffic. It's nothing special with DNS packets. So they follow the same principle. If you send the DNS query from the resolver to the left to a DNS server on the right, it will go through the network. Next slide please.

Now imagine that you have three destinations that are identical. They actually work the same way. They have the same IP address. They have the same database content. This is very important. And they use the same network, the same way to attach to the network. Packet from the source will try to reach the nearest one. And it actually works by the destinations telling the network, here I am, here I am, here's my IP

address. And you will go to the nearest one who you will listen to and hear as the strongest voice, so to speak. A good thing with this here is that different sources will end up at different places. So next slide please.

So the green source to the left will reach the server to the low left because that's nearest to that source. But the red source in the top right will reach the server to the top right because that's nearest to that source. So depending on where you are on the network, you will always go to the nearest destination for your DNS queries or this can be used in other contexts as well. It doesn't have to be just DNS. Of course, it's very important that these blue destinations are synchronized. So regardless of which one your query goes to, they receive the same answer. And that's the job of the root server operators to make sure that these are all well synchronized and that they have modern and correct data.

This has a lot of interesting features that you don't think of immediately. One is if you have volume attacks, what's called the distributed denial of service attacks, when you have large, large, large systems of small computers possibly, it could be your webcam at home, it could be your laptop, it can be your TV set-top box connected to the Internet and then possibly infected by a virus that will send attack traffic towards a certain destination. If you use Unicast, the older model, all that traffic will go to the same destination because there is only one destination. And that means that you focus all this attack traffic towards a certain target and that target will be overloaded and not be able to respond to queries properly. If you use Unicast with this

model, the set-top boxes in the red area will reach the red destination, sorry, the destination to the top right, whereas the TV cameras on the lower left will hit the server on the lower left. So you distribute the traffic between all these nodes. And if you have 1,800 nodes to attack, you have to have a pretty big network of web cameras and set up boxes and laptops to attack that, to overload them all. And only when you overload the entire system is when it stops working. So this way of scaling up massively to thousands of servers gives us a good resilience for this type of attack.

Another feature that you don't realize immediately is that if we were to remove the server on the top right and remove it entirely from the network, the routers in between will be notified of this. There is no longer a server at the top right. So the traffic will then be forwarded to the next nearest node, which is the lower right. And that happens all automatically when you stop your border gateway protocol, your BGP announcements, saying this node is no longer here, automatically the entire network of routers will switch and say, oh, we have to take this route instead, then go this way. And it just seamlessly will continue to work. Of course, the server in the lower right will receive a little higher volume of traffic, but that's not a problem because they are all massively over-provisioned. Each server is dimensioned times the traffic it already gets in a normal case. So that is not a problem. And the feature here is that as the root server operator, I can take our node in London, in England, out of service, but you won't notice. Everything will continue to work just fine, and we can do maintenance. We can install a new operating system, or we can change a hard drive that is broken, or we can do things with our computer, and then we just switch it back

on again when we get things working. Oh, now the node is back here. There is a shorter path, and the traffic will start to go back to the server that is re-enabled. So this gives us a lot of flexibility in doing maintenance on these servers. If you go back 25 years in time when there were only 13 machines, if we were to take down our single machine, that took away a little of the resilience in the root server system as a whole, and it could be noted if you really carefully looked for it, but today it's so much more flexible and so much more easy to do this. You don't even have to tell each other that we do maintenance because it automatically comes together again in the routing system and just continues to work. I think this is where I hand over to Wes to talk a little about the root server system advisory committee and its caucus. Thanks.

WES HARDAKER: Thank you very much, Liman. My name is Wes Hardaker. I am the root server operator for the University of Southern California's Information Sciences Institute, and I have some other roles than I can, and I'll get to you in a minute on the appropriate slide, but we are going to talk next about what RSSAC is and the RSSAC caucus. So finishing the technical topics a little bit and moving more into the administrative side.

So what is RSSAC? Its purpose is to advise the ICANN community and board on matters related to the operation, administration, security, and integrity of the Internet's root server system. That's very carefully worded, and if you look at it, it's a very narrow scope. We don't advise about the DNS. We don't give advice on anything else related to ICANN.

We only give advice related to the root server system, and that's it. Not IANA, not any of the other parts of ICANN.

So given that, what does RSSAC do and not do? First, you know, we're committed to produce advice, like I just indicated, primarily to the board. Almost all of our advice is sort of directed to the board, but generally we can provide guidance to the community. We have some documents that have suggested that the community look into follow-on research or follow-on studies, for example. Root server operators are represented inside of RSSAC, but RSSAC does not itself do anything operational. We actually don't discuss how anything that, like Liman just described, was related to Anycast or, you know, related to how we host our services. That is not part of RSSAC. RSSAC is really just the policy portion of the root server system.

Our leadership is composed of Jeff Osborn, who is the RSSAC chair and sitting opposite me, and RSSAC vice chair, who you already heard from earlier, which is Ken Renard. And the organization itself is composed of a primary representative from each root server operator, one representative and an alternate representative. So as I think Ken mentioned earlier, there are 13 root server operators or 12 root server operator organizations, which means that there are 24 representatives in total in RSSAC from the RSOs. And then we have incoming liaisons from other bodies. I'll talk more about that on the next slide.

And then the RSSAC caucus, which I'm going to describe in detail, I think Liman already talked about it a little bit, or Ken did actually, excuse me. But it's a volunteer body of subject matter experts, and that's actually where we do most of our work. The RSSAC itself does not actually

author many [documents these days.] We have moved almost all of that for purposes of having more people help us. I will get to that in a minute. And the members of the RSSAC caucus are confirmed by the RSSAC based on their statements of interest. So they submit statements of interest [inaudible] on a later slide as well.

So we do have liaison relationships with a number of other organizations. We have incoming liaisons, people that help us with inside of RSSAC. One is from the IANA functions operator, PTI. One is from the route zone maintainer, Verisign, who is also, of course, a route server operator themselves. So they have two roles within RSSAC. We have a liaison from the Internet Architecture Board from the IETF. And we have a liaison from the Security and Stability Advisory Committee, or SSAC. We also have outgoing liaisons. So we have an outgoing liaison to the ICANN board. I actually serve in that role currently. And so I sit on the ICANN board, and we have outgoing liaisons to the ICANN Nominating Committee, the Customer Standing Committee, and the Route Zone Evolution Review Committee. The last one is RZERC, and that is another committee devoted to considering how changes to the route server system might work. It is different than RSSAC, which is more about how the system works. RZERC is responsible for how changes might occur if we needed to make technical changes.

So what is the RSSAC caucus? I mentioned it earlier. It has over 100 current people in it that are all DNS experts. They are people that well understand the technology. That is who we turn to when we want to do some deep thinking about how can we improve the DNS or how can we study it further with respect to the route server system. And they all

have submitted public statements of interest that indicate their background with the DNS, why they are interested in participating in the RSSAC caucus, and what their expertise level is. And most importantly, they get public credit for each of their individual work. So in all of our RSSAC documents, if you look at the bottom, we have an authors section that is basically everybody in the RSSAC caucus that contributed to do it. And you will find that the number of people that contribute to a particular document is often quite extensive because we get a lot of help and we value their opinions greatly.

The purpose of the caucus is really, as I mentioned, it's a bunch of DNS expertise who bring diverse experience to the publications we produce to make sure that we think about it fully and deeply as opposed to just a narrow scope that maybe a couple of authors wouldn't be able to provide that diversity. We also provide transparency on who does the work. The caucus effort is all generally open. And it's really a framework for getting things done within the RSSAC.

So next I'm going to talk about the route server system governance evolution. So this is less about RSSAC except that RSSAC sort of started it. So what happened was when we assigned the 12th route server operator a long time ago, over 20 years ago, there was no procedure left for adding or removing them if need be. So how do we change who is doing this? And in June of 2018, the RSSAC decided to sort of bootstrap that effort to figure out what are we going to do for the future? How are we going to solve this problem? So we produced RSSAC 37, which was an approach to the governance of the route server system, including the ability to add and remove RSOs to the RSS. We also produced a bunch

of principles that I think somebody referred to already about how we operate and the ethos we provide to the root server system, things like we don't filter traffic and we answer all queries that come to us regardless.

After that document was published, the ICANN board acted on it through the guidance, through the recommendations in RSSAC 38, the organization was directed to create the root server system governance working group. That is going on now. There's five meetings later in this week that are open if you like, and in that we're developing a governance model that preserves the 11 guiding principles that we documented in RSSAC 37, including diversity, collaboration, and independence, and it includes representatives not just from the root server system stakeholders, but other entities as well as which we'll get to in a slide in a moment. And eventually this will be an ICANN document that will go to the ICANN public process for comment as well, so once we get done documenting the output of the GWG, the ICANN community will have the opportunity to comment on the results of that work as well.

So the stakeholders of the root server system, we defined actually in RSSAC 37, and it's of course the entire ICANN community has a vested interest in the root server system, as does the root server operators. We maintain the service, and so clearly we have an interest in making sure that it continues to operate, as does the IETF and the Internet Architecture Board, so the Internet Engineering course is where the DNS specifications are actually maintained, written, and developed, where changes go. That was last week in Dublin. There are many participants

of the IETF that are here this week, and they clearly have a vested interest in the root server system since they're the ones that maintain the DNS as a whole. So in the root server system governance working group, there are a number of different people that represent different organizations, so there's one member from each of the root server operators, there's two members from the registry stakeholders group, there are two members from the ccNSO, and two members from the Internet Architecture Board that I mentioned a minute ago, and then we have some liaisons as well. The blue circles on the right, the three blue circles, are the two liaisons from the ICANN board, one liaison from the root zone maintainer, and one liaison from IANA.

So with that, that actually brings us to the conclusion of all of this presentation. We will take questions in a minute, but a couple of resources for you. For more information about RSSAC itself, we do have a webpage, of course, within ICANN, and within that webpage we have a frequently asked questions list. It is designed to be easy reading, and so if you want to know more about the root server system, it's actually a great thing to just read top to bottom because it fills in a lot of common misconceptions. We also have an email list, ask-RSSAC@ICANN.org, if you have questions about the root server system or you want to know about resources or some people write us to ask about how to deploy a root server system near their country or region, typically we will point you to the FAQ for that because that's one of the things [inaudible] For more information on the caucus, especially if you are interested in joining the caucus, we have a webpage as well as there is a RSSAC-membership@ICANN.org. If you're interested in being a member, we will ask you for a statement of interest and then that will

go to the RSSAC caucus membership committee to review. As I mentioned before, we are typically looking for DNS experts or some other area of expertise that you can benefit the caucus in helping to produce our documents. With that, are there any questions at all? We sometimes break in the middle of this and we didn't. Questions about any of it, whether technical, whether about the administrative side on the end

ULRICH WISSER: So currently there's no questions in the Q&A pod. Just a reminder, please raise your hand in the Zoom room or write your questions in the Q&A pod. I will read out them. And if you're in the room, come forward to a mic and, okay, please go ahead. Good morning, everybody.

OMAR SHURAN: My question is just an informative question regarding the Anycast. Is it the same as ICANN managed root servers, the IMRS, or not?

WES HARDEKER: Anycast is a generic routing technology and all of the root server operators make use of it. A resolver would still say I want to get to IMRS or I want to get to USC ISI's server and then you will be redirected when you try and get to that particular root server operator's infrastructure based on where you are. So all 12, there were time periods where some of the root server operators were using it. I will be honest, we were the last ones to switch to using the Anycast technology. My predecessor believed that maybe it would be better if we delayed for a very long time

to make sure that it was a stable mechanism and we now know it's very stable. Not just the root server system uses it.

ULRICH WISSER: Well, there are no questions in the Q&A pod. More questions from the room?

WARREN KAMURI: So why do the root server operators do what they do? I mean, it sounds like there's a fair bit of cost and stuff. Why are you spending your money and time doing this?

WES HARDEKER: That is an excellent question, Warren, who happens to be in the RSSAC caucus and knows the answer. But we do appreciate the question. The reality is we've been doing it for a long time and we do it entirely for the benefit of the Internet, period. It is not cheap. We do not get paid for it. My organization, I have budget battles every year for how much we can devote toward buying new infrastructure and things like that. It's not a cheap process. So, Liman?

LARS-JOHAN LIMAN: If I can continue to fill in there. This goes back to days when the Internet was something completely different. Today the Internet carries a lot of business and a lot of traffic. But when the DNS was invented and when the need for root servers appeared, when the DNS was invented, we started to need root servers that could provide the top level, the

topmost level of the DNS, very small. And I will admit, as Wes did, that when Netnod's root server operator was started, not by Netnod because it's a spinoff of that first, at the Royal Institute of Technology in Stockholm, it was a workstation, a computer this size, sitting at a desk in an office. That was how unimportant it was back then. Our systems today is something completely different. It's absolutely maintained in a professional way. We have tens of staff that attend to this and it has grown over 30 years into this and we have gradually followed on. But in the beginning it was just a simple question from the IANA, can you do this for us? Sure. And we've realized that this continues as a very important service to the internet community and we want to provide that.

A benefit from, if we go back to Ken's part of the presentation, that we have very diverse organizations. As you mentioned Ken, we have government agencies, we have for-profit, we have not-for-profit, we have academic institutions. We all have different financial foundations for how we operate this. So Netnod finds money to provide this service in our way, whereas the University of Southern California has a different way, [inaudible] agency has a third way and the commercial entity has a fourth way. So even if the financial systems were to change in a significant way, not all root service operators would be affected immediately by the same change. So it's also a type of diversity which is very important for the longevity, for maintaining this service over a long period of time. Thanks.

WES HARDEKER: Any other questions?

KEN RENARD: I just feel the need to throw in another answer to the answer that's already been made because Wes works with the university, I'm the president of a non-profit company and it's gotten to the point where on the internet when you say you do something for free, I literally will have people throw that in my face and say Facebook doesn't charge me either. I want to make it really, really clear. My organization writes software that the internet requires to operate and we give it away. Our entire expenses are paid for by people who come back and say, I need some help, can I buy support services from you to help with the free software that I didn't pay you for? We have an arguable million and a half users of our BIND software around the world. About 100 of them pay us any money for the help. In addition, we take about 20% of the money that we make from supporting BIND and other open source software and we spend it on operating a root server system. So when we say we don't charge for this, I don't mean and then we make the money selling all of your private information. There is no private information passing hands. There is no information passing hands. There is no content passing hands. All that we do is hand out the address of an authoritative server that knows where the TLDs are. So when we say we are a non-profit, not making money, that is the truest statement you will ever hear. This is a genuine non-profit, genuinely not making money. We literally do this and if you're cynical, maybe you don't want to hear it. We literally do this because I operate a company of 19 people from 19 different countries who all genuinely believe that a single name source internet is valuable enough to go make less money working for a non-profit. We really are committed to our mission.

HANS PETTER HOLEN: If there are no more questions, I can provide another answer because since we are 12 operators, we all have different answers, right? I'm operating K-root, or my organization is. Don't let me play with the service anymore. We're a membership organization. I have around 20,000 members who are paying for the operation of K-root along with other services. If you wonder what it costs, you can look at our annual financial reports or activity plan and budget to see how much this costs. If you divide it by the number of members and you can say that they chip in like 50 euros each a year for this. There are many different ways of financing this which is a strength of the system that we all have different ways. It would be very strange if everybody will run into a financial crisis at the same time.

ULRICH WISSER: Well, we have a question in the Q&A pod. How is RSSAC's relationship with other groups and I can especially non-commercial?

WES HARDEKER: So RSSAC has meetings with other constituencies that have sort of direct ties with us or direct need to talk to us because as I think Ken pointed out earlier, we serve the IANA root, right? We have pledged that the source of the data that we get is from IANA. Most of the I can constituencies that are representing TLDs, for example, regardless of whether it's a ccNSO or the GNSO, they communicate directly with ICANN to get those TLD functionality change. We communicate with IANA. That's our primary focus for actually operational. We have

meetings with the Security Stability Advisory Committee on a regular basis. We have meetings with the ICANN board on a regular basis. We're meeting this week with the ASO. So we meet as needed. And we've had meetings, I think, with ALAC sometime recently in the past as well. We met with a GAC this morning. So it's as needed. We typically don't like to hold meetings unless there actually is a need. So we're actually not meeting with the ICANN board this time. We have for the last three or four times. There are no direct immediate need between two organizations. We don't, you know, waste people's time by making them sit in a meeting with no content.

LARS-JOHAN LIMAN: But I would like to stress that interaction is always welcome. There are issues which you believe involve the roots of operators. Please come and talk to us. And if that leads to the need to have a meeting, we are quite open to set up these meetings. But we don't conduct regular basis unless there are issues on the table that we need to discuss. As Wes said, there are enough meetings in this context. We don't need to add them just for the sake of meetings. But we really, really welcome interaction with other groups and with individuals. If you have questions regarding the root server system, stop me in the corridor and ask. You'll find me. I'm two meters tall. I'm easy to find.

KEN RENARD: There's 12 root server operators, root server organizations. Seven of them are in the room right now, so we can raise our hands. We're here.

ULRICH WISSER: Well, another question in the Q&A pod. Can you tell me more about the Anycast operators and what is the relationship with RSSAC, if any?

LARS-JOHAN LIMAN: The Anycast operators. That's us. We operate the Anycast network. And mind you, we don't have to do anything special with the routers on the on the Internet. So. In the case of Netnod, we operate 75 servers in different locations, and that's our servers. We operate them. We ship them to allocation. We have some people help us to mount them in a rack and connect them with the necessary cables to network and power and whatnot. But then we log into the server and we do the operation and the server will announce in the relationship—we have to negotiate in. A relationship, it will announce to the next router upstream towards the network. It will announce that now Netnod's root server is here with this IP address. Please bring traffic to me. So there's no need to have a specific Anycast operator. Various root server operators operate their own subset of these 1800 servers. And each of these server installations, it can be one more than one physical box. It can be a couple of servers and a local router for the for that root server operator. That cluster, that combination will announce to the nearest router that I'm here. Send traffic my way. And that's all that's needed. Thanks.

ULRICH WISSER: Is there any legal or technical requirement that root operators have to meet in daily basis other than uptime?

[HANS PETTER HOLEN:] So we tend to say that, you know, the Internet is not regulated or any government jurisdiction. And that is true for the policy processes of ICANN or the RIRs or anybody else. But each and one of the operators have some legal entity in some country and are subject to the laws of that country. RIPE NCC is in Europe and we have to adhere to European legislation. My colleagues here are in the US. They have to adhere to US legislation. Those legislations do not in detail regulate how we operate. But there is new legislation coming in putting new requirements on operation of critical infrastructure. And this is happening at different pace in different countries. So it's an exciting time to be.

LARS-JOHAN LIMAN: I would like to add because there was also a question of technical specifications. And yes, there are a couple of technical specifications. We, for instance, we have to follow the DNS protocol that's specified by the IETF, the Internet Engineering Task Force in a certain technical standard. And there is also a document that describes expectations of how the root servers should respond to DNS queries coming in. And that's published by the Internet Architecture Board in cooperation with the root server operators. But the RSSAC caucus, as we've talked about, have produced a document with service level expectations, which has been published or is about to be published very soon. If it isn't already published, it will be very soon, where there is a number of expectations set on the roots of operators that we are expected to respond to with technical limitations, response times. And between release of new data to the data is updated at the end nodes and all kinds of such things. So yes, there are technical expectations.

WES HARDEKER: RSSAC 047, I think, is what you're referring to. I will say, adding a little bit more, we can't just monitor whether our devices are up or not. You'd be pretty impressed, I think, with how to manage any network, regardless of whether it's a root server, or a DNS server, or a file server, or a web server. The number of graphs we have to create, and the number of alerts that we look for, it's not just whether it's up, is it acting too slow, is it being overwhelmed with traffic, that's what really wakes me up quickly. Is the disk running out of room? There's a lot of little things that go into running infrastructure, and you need to know immediately when something fails. And when something new comes along that we did not think about monitoring before, that's the first thing I do, is write a new thing to monitor, so that I have a co-worker that excels at creating very interesting graphs, and I look at them all day long.

LARS-JOHAN LIMAN: There's also a whole lot of experience. I mentioned we have probably 10 people operating this, and there are numbers of operators at every root server operator. It's not just the few people in the room here, we are the people that deal with it and stuff, so again, I very seldom touch the technical details of the root servers. But the experience of the root server operator with running DNS servers is probably measured in centuries. If you put a whole lot of experience of running DNS, of looking at DNS content, DNS traffic flows, packet content, update procedures, monitoring, network traffic engineering, all of that is part of operating

a root server. And we all have that expertise in our respective organizations, so it's not just the few people around the table here.

[YUSUF:] My name is Yusuf, I attend from [Marmara University] in Istanbul. My question is, how do you understand there's enough instance in a region, for example in Istanbul, there are seven instances. How do you know, is it enough, whose responsibility to check this?

HANS PETTER HOLEN: Though I can only speak for one of the root server operators, K-root, RIPE NCC also operates worldwide measurement networks with more than 13,000 measurement probes. And we use that measurement network and other sources to look at the latency from where the probes are in the networks to see that we provide good service at all parts. Now this can be used by all the root server operators, some use other measurements in order to see whether they think the service they provide is good. If you have special interests in this, RIPE NCC also takes open applications for hosted nodes, so that you can get in touch with us to get one on your local ISP for instance. If you believe that there should be more distribution, you can contact us and discuss that. And the other root server operators may have other ways of doing this, and that's part of the beauty with diversity. We don't have one common way of approaching this, we have 12 different ways of approaching this, so that even if my way isn't right, then somebody else may be better.

WES HARDEKER: We should note that the internet worked just fine before Anycast, and there was 13 servers. So you actually don't really need one near you, as we have indicated before, caching makes your ISP's DNS server fast. It's not because you have a root server near you, very, very few questions actually get to the root, so the need to get to the root is not critical.

[TAHIR:] Hello, Tahir from Istanbul Technical University. I would like to ask one question about, in the Anycast we have equivalent servers distributed around. But what about the data or the information they have? Do you have any synchronization problems, or how often do you need to synchronize everything? Do you have online offline procedures for these?

KEN RENARD: Yes, all those servers need to be loosely synchronized. The root zone gets updated maybe one to three times on a typical day. We do measure and publish the measurements of how fast that happens to, whatever, RSSAC 002. So typically a root zone is published, and now it's out on all infrastructure. It can take seconds to send it all out. But it varies if an instance is in a remote location that doesn't have good connectivity, it's going to take a while to get it there. So we rarely have any issues with that synchronization. It does happen. But there are DNS protocol parameters that say if your version of this zone is so much old, stop serving it. So we want to avoid the situation where the data is just too old, especially if the DNS signatures are no longer valid because they have a validity time as well.

WES HARDEKER: It's worth noting that every record in the root zone, you are allowed to memorize for two days before you ask again. So keep in mind that actually all of the records are essentially promised to be valid for two days.

LARS-JOHAN LIMAN: And again, Anycast plays in here because if we, through the monitoring systems, we keep an eye on all the systems, and one parameter we check is, is it updated properly. If we notice that it's not updated properly and we cannot fix it quickly, we turn it off. We turn that node off and the traffic will no longer come there. It will go to a different node, which is updated, and we can fix the problem while it's lying. And then we switch it back on again when we get things working. So again, Anycast works to our advantage and yours.

[JUNAID MIYAJI:] Hi, my name is Junaid Miyaji, and I'm from Bangladesh, and this is the first time I'm attending ICANN meeting physically. So I have a question. A few days back, we have faced an issue in our country, which was an Internet shutdown. But in that time, I found that we have been disconnected from the whole Internet, but our IX network was working fine. But in that time, we was not able to resolve any DNS, because without DNS being connected on the Internet, it is not being able to resolve any domain names. So do you have any options or something like that to make root zone works with IX networks, I mean, the local

Internet, or is there any other way, like an individual like me can cache the records of root zone, for particularly our countries? Thank you.

LARS-JOHAN LIMAN: Yes, yes, and yes. So to resolve a normal DNS name, you need not only a root server, you need the following levels going down as well. And then it explodes in the number of alternatives very quickly, so you need to cache a lot of data. That said, the root zone is actually public information, so you can download that. I can show you how to do that later. And there is a document from the IETF that describes how to operate that zone in your local resolver. You run a copy of that root zone inside your resolver, so it never has to talk to a root name server. And then you rely on yourself only. And you can copy that zone regularly and do what's called a zone transfer from some open nodes. Operators have nodes that will allow you to copy that zone automatically. And if your network is shut down, so you cannot copy a new version, it will continue to use the old version for, I believe, up to a week or thereabouts. So, yes, there are ways to get around this problem. I do believe, though, that during this blackout, that the root servers in Bangladesh might have still worked, but the next hop down, may be the com server or the... Sorry, I don't remember the top-level domain for Bangladesh. But, again, the steps below might not work, because then the servers might be outside Bangladesh for the next levels down, and we can't help you with that.

WES HARDEKER: There are actually six root servers in Bangladesh, so they should have stayed working, most likely.

[JUNAID MIYAJI:] I'm sorry, but unfortunately, I have tried everything, the way you have mentioned that, to make DNS servers myself and caching the root records and making them resolver. Unfortunately, that didn't work for me, and that is why I have asked that question. Also, during the shutdown, we have... I have tried to find a thousand ways to make it work at Anycast, but it didn't work at all. And also, for the records to be resolved the way you have mentioned, that didn't work for me either. So I believe that I might have made some technical mistakes or something like that. So is there any way if I can be able to connect some technical persons or technical things to sort these things out from...

WES HARDEKER: So there are multiple things that you're referring to. So there are two parts, right? Imagine you're trying to look up the name of a website, like for example, and that website is running and I'll pick Antarctica, not in a country, right? If a region is not on the internet, could be from a natural disaster like a hurricane or whatever, whether it's political or not, there's no way to get there, right? The web server is in Antarctica, so even if you could look it up, that doesn't work, right? Now, the other aspect is if the DNS server, let's say you're looking up example.com, if the DNS server for example.com is in Antarctica and you're cut off from it for whatever reason, a cable cut, it doesn't matter. Your ISP may start at the root and you'll get it because it's right there, but the root server cannot fix your access outside an environment. The root server is not able to fix that for you.

[JUNAID MIYAJI:] Thank you, but I guess I have missed some point. I am acknowledged about that. What I'm talking about the IX networks, we have some servers in our local networks, like I have my servers in Bangladesh, in Dhaka. So it should be able to resolve the websites that are hosted in Dhaka, not outside Bangladesh, but would like to resolve the domain names that are hosted within Bangladesh. And those websites particularly was working with IX networks, and we was being able to resolve them through IP directly, but not with the domain names we would be able to resolve using the root zones. That is what I was asking about.

KARL REUSS: We partner with a PCH packet clearinghouse to provide a lot of our Anycast nodes. We partnered with them and run our servers out of their infrastructure and they help lots of countries set up exchange points and host their own internal DNS for their top level domain. This is the kind of thing where you need a lot of things to be set up and working right, and it is hard to test them unless everything is broken. You need to go through lots of troubleshooting to make sure it is set up right. We can work with you. I know we are in Bangladesh. We can look for some of what you saw and try to see if we can see anything that is not set up properly. There are several pieces and the root server is just one of them. The root server may be working fine, but there is still a lot that can be wrong.

[HANS PETTER HOLEN:] If I may add to that, I think this is a very interesting problem that only some very few countries have actually tested large scale. How can you keep the infrastructure in the country running if you lose connectivity to the rest of the world? Some countries have tested that and there are suspicions of why they are doing that. Some say it is to cut off their countries. Some say it is to defend themselves. That testing in itself has been seen as controversial. I think we as an engineering community should get together and look into this and see how we can make as much as possible of a subset of the network working in terms of crisis. Because we have countries in the world now that have armed conflicts going on and we have natural disasters and whatever. Whatever we can do to understand this and make sure that the internet stays resilient, I think that is worth to go through. But as I said, the root servers are a very small part in here. I think this is more for those of us that are also engaged in other activities in DNS.

KARL REUSS: It looks like we do have a lot of services through PCH in Bangladesh. So it feels like a lot of the pieces are there the way they should be and I think they are just not fully set up the way they need to be.

WES HARDEKER: Your problem unfortunately is not the root server. It is more like .bd or some of the other infrastructures. But it would take analysis to get there. Any other questions?

ULRICH WISSER: We think we are at the overtime already. Thank you everybody for coming.

[END OF TRANSCRIPTION]