ICANN76 | CF – Root Server System Information Session
Tuesday, March 14, 2023 – 10:30 to 12:00 CUN

OZAN SAHIN: Hello and welcome to the Root Server System Information Meeting. My name's Ozan, and I will be the remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior.

During the session, questions or comments submitted in chat will only be read aloud if put in the proper form as noted in chat. I will read questions and comments aloud during the time set by the chair or moderator of the session.

If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record, and speak clearly at a reasonable pace. Mute your microphone when you're done speaking.

This session includes automated real-time transcription. Please note that this transcript is not official or authoritative. To read the real-time transcription, click on the Closed Caption button in the Zoom toolbar.

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

With that, I will hand the floor over to our RSSAC chair, Jeff Osborn, to kick us off.

JEFF OSBORN: Thank you, Ozan. Thank you, all of you, for attending online and locally. We appreciate your involvement, and we look forward to being able to shed a little light on this. So if you have any questions, feel free to look us up later. And I hope this will be informative.

I'm going to turn this over to co-chair, Ken Renard.

KEN RENARD: Thanks, Jeff. We can go to the next slide, the agenda. So welcome. This session is meant to give a technical background on how the Root Server System works. And we have several RSOs in the room. If we could raise our hands, those that are operating root servers.

In the past this presentation had been given by ICANN staff, and we thought it would feel a little bit more at home if it's actually the RSOs presenting this. So here we are.

The Root Server System is quite technical. You can really get down into some fancy ones and zeros discussions. So we're trying to target this at a specific level of detail. If it's too much, as us to slow down or ask questions. If it's too high level, feel free to go to sleep. We have one hour today, and our objective is to help everyone understand the Root Server System. Next slide. And then next.

So just looking at the DNS from a very high level, really stepping back to 10,000 feet. So the fundamental identifier on the Internet is the IP

address. So it's a string of ones and zeros. We typically see them as numbers. If you're connected to the Internet, you have an IP address. Okay, you may be connected indirectly and have certain types of IP addresses, but do you see in the lower right-hand corner those numbers down there—an IPv4 address and an IPv6 address? Just an example. And if we go to the next slide. So why are we doing this DNS at all?

So originally, 30-plus years ago, the problem was that IP addresses were just hard to remember. And that was before IPv6, before these super-long addresses. And then the fact that IP addresses change. They can change quite often, actually. And for the most part with DNS, we don't know we; don't care. It just happens.

So more modern reasons why are now that, especially when you're talking about websites, a single IP address may sort of have multiple names. And in the same vein, multiple IP addresses could actually serve the same name, or same logical names, or a human-centered entry point. That name gets translated in multiple ways to one to many, many to one. So it makes it a lot easier for us as end users to understand. Next slide.

So this is kind of the typical hierarchical example we see of the Domain Name System with a root at the top. And then underneath that, we have our top-level domains, and then second-level domains, and so forth. This sort of shows the authority levels or levels of ownership of the namespace. And so from the root, we can delegate top-level domains and delegate, underneath that, lower-level domains.

Please have a seat at the table.

So the most common use of the DNS is to map a name into an IP address. And you see that. www.example.org there equals 192.51.100, etc. But there are many other uses of DNS. So mappings of mail servers; obviously, IPv6 addresses; the inverse of looking up addresses into names. And there's also lots of security records that are in the DNS that allow us to enable security at other levels and to enable security within the DNS itself.

So we see the distribution here. The root is actually a pretty small zone. It's a very small set of names in the root compared to other parts of the DNS namespace, but it is critical. As the entry point into the DNS, it's critical in order to look up that top level. So in theory, everything depends on that root level being there and being available. So next slide.

This is one we'll talk most about with respect to visualizing the process here. So on the left side, you see your device—your iPhone. This could be your laptop. This could be your web server or a big supercomputer. Pretty much every computer on the Internet takes on this role, and we see this idea of a stub resolver. This is just usually an entry point in the system, like the OS. The operating system will sort of aggregate this DNS function so that applications don't have to implement the full DNS protocol.

We'll put a footnote on there for those that are interested later on.

So typically, when your device is connected to the Internet, you are assigned or given configuration to a recursive resolver. This could be a dynamic when you join the wireless network, or it could be pre-configured. This recursive resolver here in the middle is kind of the

workhorse of the DNS world. There are probably millions of these out there. They have the biggest job and do most of the actual work for the end user. But we'll walk through this.

And for the typical explanation here, we'll say that this recursive resolver has just been rebooted and there's nothing in the cache. Or we'll say we've rebooted the Internet and there's nothing there. So your application, your phone, your web browser—whatever it is—is going to look up www.example.com. And that recursive resolver has no a priori knowledge of any bit there, so the only thing it knows is the location of the root name servers.

I'm skipping a step, but basically what it's going to do is ask the root, as the first step, "What is the address of www.example.com?" And the root server says, "I don't know, but I do know where to find the .com name server. Go ask the .com name server."

So the recursive resolver recurses this or repeats this process. It goes to the .com name server, asks the same question, "What's www.example.com's address?" .com comes back with a similar answer, "I don't know, but here's the name server for example.com."

Finally, that recursive resolver goes to the authoritative name server for example.com, asks the same question, "What's the address of www.example.com," and it gets an answer which it then sends back to the device, the phone on the left, and we can go on.

Now the important thing to note here is that recursive resolver, it has gone out and asked three different questions to three different name servers. Every one of those answers that it has gotten, it caches, it

remembers for a certain amount of time. So the next time that you need to ask a question—"What's www2.example.com?"—it can skip the process of going to the root because it already knows the answer of where the .com name server is.

In fact since it's the same second-level domain, it can skip the second part where it goes to .com, and just go right to the example.com name server. So there's incredible efficiencies built into this recursive resolver and its caching mechanism.

So the other important note here is that this recursive resolver is used by many end users. Many could be hundreds, thousands, millions, maybe. I don't really know the scale of some of these larger recursive resolvers, but the effect is that for the millions or billions of queries that this recursive resolver is going to see over, let's say, a day—a 24-hour period—very few of the queries actually come to the root because of this caching.

So again, the size of the root zone is pretty small. So in theory, you would have only, say it's ... There's about 1400 or so, 1500 entries in the root zone. There may be only 1500 queries to the root zone every time period.

Now, each of the answers that comes back from these name servers on the right comes with a TTL, a time to live. So this is a recommendation to the recursive resolver of how long you can keep this record in the cache before it's a good idea to come back and ask the question again because it may have changed.

Now for the root zone in particular, those values are most likely 48 hours. So that recursive resolver can get an answer from the Root Server System and keep that answer for up to 48 hours before it has to ask again for that particular set of records. The recursive resolver could choose to do something different—less, maybe even more—with the recommendation from the authoritative name server. Let's go to the next slide.

So we've talked about most of this already. The recursive resolver caches those things, and you get incredible efficiency, such that these recursive resolvers very rarely come to the root zone or the Root Server System. We talked about the time to live, the time period which they can keep that record in its memory. And the root servers only have information about the top-level domain servers. So most of that question that was asked—www.example.com—that's way more specific than the root zone. We only know about those top-level domains and where their name servers are. So a lot of efficiency there in that caching recursive resolver. Thank you.

So the Root Server System has been around for almost 40 years now, and it has evolved quite a bit. It has grown and scaled in size—the number of physical machines in the anycast cloud. These are some of the additions that have been made to the DNS and specifically to the root zone or Root Server System. So DNSSEC was added.

I can't quote the actual date myself, but DNSSEC is the cryptographic signatures that provide authentication to the data. So regardless of which actual server you get this response from, you can verify that the

signature is correct and that this data is original and unchanged since it came from IANA.

So DNSSEC is very important in the accuracy of the responses that you get. So we encourage everyone that runs a recursive name server to enable DNSSEC validation.

So other enhancements are for privacy. So when we saw that example, the entire query www.example.com went to the root server. Well, the recursive knows that it's only going to get an answer of where the TLD is, so you can actually have the recursive resolver only ask that specific question. So that way the root zone doesn't see the entire query. It only sees the request for, "Where is .com?" And we encourage that as well.

So other privacy enhancements ... So that privacy enhancement is called query name minimization or QNAME minimization. Other privacy enhancements—DNS over TLS; DNS over HTTPS; and another one, DNS over QUIC—are out there. And that's actually encryption that applies to the connection between the end device over there on the left to the recursive name server, as well as from the recursive to the far right, the authoritative servers.

The other refinement to DNS that's been implemented in the Root Server System is anycast. So we'll see soon that there are 13 root server identities, but that does not mean that there are 13 servers in the world. We each have multiple servers, and they share a single IP address which greatly improves latency, resilience, and protection from denial-of-service attacks.

With that, I think I'll hand over to Liman if you're ready for the next set. Thanks.

LARS-JOHAN LIMAN: Thank you. My name is Lars-Johan Liman. I work for one of the Root Server Operators, Netnod, based in Stockholm, Sweden. And I am going to talk a little about the Root Server System itself and how it works between the Root Server Operators and what the anycast is that was mentioned on the previous slide here.

Can I ask to go back to slide 7? Because that's such a very central picture of how the DNS system works. At the very top right, we have the root authoritative server, and that's the piece of this picture that is operated by the Root Server Operators. Everything else is operated by other people, and we have no idea who they are, most of them, and we don't really care. The arrow going upwards towards the right is the incoming traffic to the root servers, and that's what we see and what we operate with. So when I continue to talk about the Root Server System, it's only that top right square of the system.

So if we can go back to slide #11 now. Thank you.

So first, some useful definitions. The DNS world is full of technical words and abbreviations that are very specific and identify various details of the system. And I keep comparing the DNS system to Lego pieces that you can put together as toys and play with, and you can build large structures.

DNS has a bit of the same properties. It has very small pieces that are very easy to understand, but using them and combining them, you can

build very large structures that become a bit complex when you try to look at them. But I'll try to identify some of these here.

So the root zone is actually a piece of a database. The DNS system is a large database that you can access using the DNS protocol. And it's publicly available. Anyone on the Internet can look up virtually any data in the DNS system, even from inside the room here.

So the root zone is a small piece of this puzzle. It's the part of the database that is stored in the root servers. And as Ken said, this is actually a rather small thing. It's a list of the top-level domains. They are 1,400. So you can imagine 1,400 text lines listing the various top-level domains. But in addition to that, it also lists which machines on the Internet act as authoritative servers for the specific top-level domain in question. And they all have their own set of such name servers. So the top right box has a list of all of the other ones that act as the next level down.

In addition to that, there is also ... Again, mentioned by Ken, there's the DNSSEC, the security signatures that are added to all of the records. So it actually is larger than 1,400 lines, but it's not 2 million lines. It's on the order of thousands of lines. And in computer lingo, that's a small database.

So the root zone is the content. That's the data inside the server. The Root Server System, often abbreviated RSS, is the combined set of servers. From all of the Root Server Operators, all together, we provide the Root Server System. So that's a set of over 1,500 servers over the entire Internet that provide this top-right corner service. You can talk to any one of them; you will have exactly the same answer, and we all

provide the same service. And if you put it together, you get the Root Server System.

That system is operated by 12 different organizations. We work very closely together to make sure that we provide exactly the same service; and when we make changes, we make them coordinated in a stable fashion. And the 12 organizations are denoted as Root Server Operators. So I work for one of them, Netnod. Ken works for another one. Jeff works for a third one. And together there are 12 of us.

A root server anycast instance. That's one unit of these 1,500. So Netnod has a root server instance in Helsinki in Finland. We have another root server instance in Thimpu in Bhutan. We have a third one in Asunción in Latin America. So each and every one of these 1,500 is an instance, and every instance is a complete root server. It has all the data needed to provide the service to respond to all root server questions that arrive at it. So they are all identical, so to speak.

The Root Server System Advisory Committee. That's a committee inside ICANN. It's been there from the inception of ICANN back in 1999. I was there for RSSAC meeting one. So were several of my friends. Most Root Server Operators have a very, very long experience of working with the DNS. I started to work with the DNS in 1989. I started to work with the Root Server System in Sweden in 1992 when it was six months after its deployment. And the same goes for many of my colleagues as to Root Server Operators.

ICANN has this committee, RSSAC, to advise the Board on matters that pertain to the Root Server System because the Board needs to make sure that it doesn't make policy decisions that influence the Root Server

System in a negative way. The Root Server System, as mentioned, has evolved over time rather drastically, but it needs to do so in the same pace as policy decisions are made and as technical changes happen as the network grows.

So Root Server Operators have to be involved in many fields. We are here at ICANN meetings. We are often at the IETF for technical standards: what's new in the DNS? Do we need to have new software? What do we need to adapt to?

We are at network operator meetings because ... How is the network growing? How are the operators interfacing with the root name servers? We have to be aware of all these details in order to evolve the system with its time. The machines that we operate now in 2023, nothing in common with what I did in 1992, 30 years ago.

So this advisory committee exists to advise the ICANN Board and the general community about the Root Server System, to help them and inform them to make good decisions that may influence the Root Server System. Next slide, please.

So the root zone versus the Root Server System. Again, the root zone is the content. It's the database content that people want to access using the DNS protocol. And it's the starting point for the DNS system which is hierarchical.

Sorry. Yes? A question back ...

UNIDENTIFIED MALE: If one of the root servers failed, is there an issue? Or is that why there's 12?

LARS-JOHAN LIMAN: There aren't 12. There are 1500.

UNIDENTIFIED MALE: I guess if one of the operators—

LARS-JOHAN LIMAN: Yeah. To begin with, we would have to make ... I would have to make a serious, serious, serious [fault] to take out the 70 nodes that Netnod happens to operate. And even if I did, you wouldn't notice. There's even a document that describes the fact, so to speak, that the significance of a single Root Server Operator is not a sensitive part of the system.

I see Jeff's hand up.

JEFF OSBORN: Lars is correct. But to put it more tersely, if three of the Root Server Operator organizations disappeared from the face of the earth right now, I would defy you to figure out that it had happened. There literally is so minor a change in the function that it would be indistinguishable to anybody other than a scientist looking for it. This is a 1600-unit system operated by 1200 people all doing the identical task. And so the failure of any huge [swathe, it] doesn't actually happen.

I was going to save this line because I'm not a presenter today, but my boss has been involved with this for decades and is a great computer

scientist. And when I said to him, "I'm trying to model what a failure of the Root Server System would be like," he leaned back and thought about it and he said, "Ah. In the sixth month of nuclear winter after a devastating global war, you can imagine a bifurcated ..." and he went on to explain a root server failure.

We have much bigger problems in things like survival, radiation, and things like that before the system fails. It is literally robust to a degree that is difficult to model a failure. Every single instance on Earth by itself could manage this function.

WES HARDAKER: Oh. Let me correct one slip of your tongue there, Jeff. You did say 1200 Root operators. There are 1200 people. There are not 1200 people [inaudible].

JEFF OSBORN: There are 1600 instances operated by 12 organizations. I apologize.

WES HARDAKER: Right. Each of which have multiple people, but it doesn't add up to 1200.

JEFF OSBORN: Correct. My bad.

UNIDENTIFIED MALE: Thank you. I just had one more question, too. Can you talk about the number of TLDs and domain names and how that impacts the systems? And it's obviously very complicated and technical, but is there a point

when it would need to expand if there were so many more TLDs than there are today?

LARS-JOHAN LIMAN: Thank you. Of course, there's such a point because the size of a root server is not infinite. But, again, the system is so over-provisioned that it's a long stretch from today's system until we reach that point.

And we had a session with the ICANN Board yesterday where that specific question was asked because they are looking into the new round of gTLDs. And the question was said, "Will the system work with 1,000, or 10,000, or 100,000 new top-level domains." And personally, I'm fully confident that 100,000 new top-level domains will not impact the root system noticeably. And that's without touching the system that we have today.

Now, what I told the Board yesterday is that the important thing is not the number, but it's how fast. How quickly do you want to add new zones? Ten per week? A hundred per week? Two hundred per week? That's what we need to know so that we can expand the system in cooperation with the other things that are going on regarding a delegation.

And that's exactly why RSSAC exists because that's the dialogue between the Board and RSSAC and the IANA who does, you know, the hands-on delegation when it happens so that we can have information about what are the intents from the Board and from the other parts of ICANN and so that we can get feedback. "Oh, two million new per week? You can't do that for long, so please give us a little slower pace."

So that's, you know, the give and take that we have to have. But the numbers that were mentioned by the Board don't worry me in the least. The system is massive, and it's massive in so many directions— number of servers, the capacity per each server, the network connectivity for each server, the memory size in each server, and the CPU power of each [box]. Massive. So I go to sleep with a light heart every night.

WES HARDAKER: Ozan, we'll get your question in one second. A real quick follow-on to that. I'm on the ICANN Board. We have a ridiculous number of problems to work out about the next round of the gTLD system. They asked us, "When are you guys going to be ready?" And yesterday we told them, "We have been ready for a long time." Right? Actually, this is the one problem the Board does not have to worry about.

LARS-JOHAN LIMAN: And on top of that, I don't worry because I know that the administration that will eventually have to deal with the applications cannot pump out 200 top-level domains per week. So, not the problem.

Sorry. Yes, Ozan.

OZAN SAHIN: Thank you, Liman. We have a question in the Zoom room. So first, I'd like to check with Liman and presenters whether you'd like to address questions now. Or do you want to keep them until the end of the presentation?

LARS-JOHAN LIMAN: Keeping them to the end has the advantage that maybe we cover what's being asked in the rest of the presentation. That said, having a dialogue is always a good thing. But we also want to reach the end of the presentation, so we need to keep an eye on the time. So I suggest we take this question, but then we post them to the end of the presentation if that's okay with people.

OZAN SAHIN: Okay, so—

LARS-JOHAN LIMAN: I don't see any objection, so if, yeah.

OZAN SAHIN: The question is coming from Lütz Donnerhacke. And the question is, "How long will you survive a failure of the root key signing procedure?"

LARS-JOHAN LIMAN: Hello, Lütz. Thank you. That actually ties in back into a part of the system that the Root Server Operators do not operate. That is generating the zone file. But that part has ... And I look at—

WES HARDAKER: You might flip to the next slide which is where the diagram is that you're going to talk about.

LARS-JOHAN LIMAN: Yes. Yes, but it doesn't have the times in it. So there Duane when I want you. Thank you. Duane Wessels just stepped into the room. Because you probably have the timing details that I'm going to talk about at the top of your head.

So the signing, the key signing parties that take place every three months generate keys that can be used for three months. If the Root Zone Maintainer who generates the root zone continues to work, the Root Server System operators will be able to pick up zones from there. And that will just continue to work.

But when the key that's used for signing the zone runs out, that's when things start to go downhill. And I believe that you do ... I'm looking now at Duane from Verisign, the Root Zone Maintainer. You have a spare key that ... And you do these things in advance, so you have keys for, what, six months or so to run ... Do you remember?

DUANE WESSELS: Yeah. Sorry for jumping in here in the middle of the question. So Wes showed me the question. The question is how do you survive a failure of the key signing system?

LARS-JOHAN LIMAN: How long does the Root Server System survive a failure of the key signing process?

DUANE WESSELS: Yeah. So maybe as you were getting to, there are two key signing processes. Right? There's a KSK. There are signatures from the KSK

which happened every quarter. And so, you know, those are planned well in advance, and there's sort of lots of extra time. Failures of that is ... There's plenty of time to recover from that—months to recover from that. A failure of the zone signing key system ...

So first of all, there's quadruple redundancy. We've got four HSMs and duplicate data centers, and so on. The signatures from those are valid for two weeks, I believe. And those are resigned every day. So if there's a failure, we've got a couple of weeks to recover from that. So again, plenty of time and plenty of redundancy.

I don't know who asked the question, but hopefully that answers.

LARS-JOHAN LIMAN: Thank you. And sorry for [shooting] you into a bad spot. So the root zone is the content that we put into the root servers, and we do obtain that from the Root Zone Maintainer who in turn receives instructions from the IAN which is part of the ICANN structures. And that zone is compiled and distributed to the Root Server Operators from multiple points. There is redundancy in that, too. We collect the song content. A couple of times per day, we update our servers with fresh content. And we can download that from several different sites in different locations, different parts of the world. So there is good redundancy in that as well.

And that root zone, the content of the database that we put in our service, is actually publicly available. You can download that into your web browser and have a look at it if you want to. And it's exactly the data that we put in our servers. There's no fake.

So that's the answers we have to give out when people ask DNS queries to the root server. And the Root Server System is loaded with this data. It's distributed to all the 1500 servers—the same data in all of them. And the servers respond with that data.

And the Root Server System, there are 13 identities that you can talk to. So from every location on the Internet, on the globe, you can talk to a selection of 13 different root servers. You're able to reach 13 different root servers. And you can choose to use either IP version 4, the older version of the Internet Protocol, or IP version 6, the more modern, version because they're both served by the Root Server System.

The role of the Root Server System is purely technical. The Root Server Operators do not mess with the content of the zone. We don't check anything. We download it, we put it in our server, and we take our hands off.

So any policy regarding top-level domains, which top-level domains should exist and be reachable through the Root Server System, that is out of our hands. That's part of the policy process inside ICANN. The Root Server Operators do not tamper with the data. And it's very easy for each and every one to verify that the data received is correct because it's signed. It's signed in the process of generating the zone root zone file.

And you can verify that everything is correct—that the Root Server Operator didn't tamper and fix with the file, DNS packets traveled across the network without being interfered with, and the signatures should be okay. Next slide, please. There we go.

This is an attempt to do a picture of how the root zone is created and moved into the actual instances—the 1500 servers that provide the service. So at the left side, you have TLD operations which is ... Sorry. That's TLD operations. That's your TLD. That's your registry for your top-level domain. And it may want to request a change to the root zone. It does so by contacting the IANA function which is part of the ICANN structures.

The IANA function will validate that this is a valid request coming from the right source, that doesn't have any technical issues; and then request the change to be implemented by the Root Zone Maintainer. This is a function operated by Verisign, and Verisign will generate a new root zone and add to that unnecessary signatures and then make it available to the Root Server Operators. That's the small white square in the middle. It's depicted as only one white square, but there's actually a number of places on the Internet. They copy it to distribution servers that all Root Server Operators can reach.

And we Root Server Operators pick that zone file up on a regular basis. We get a notification saying there's new data, and we immediately pick that up. And we have, then, our own distribution system. Each Root Server Operator works independently, so we have our own distribution systems to spread it to our part of these 1500 servers. Netnod operates 70, 75. So we copy the zone file from Verisign, and then we put it into our 75 servers.

And the other Root Server Operators do something along the same lines. And that happens very quickly. When we receive an update, it's

probably within a minute, I would say, that 90% of all of the 1500 instances are updated. It's really, really quick.

And on the more right-hand side, you'll see the RS bubbles. That's the individual instances that actually receive the DNS queries from the DNS resolvers on the far-hand right.

So the initiative to make a change comes from the TLD registry which interacts with IANA—validation, that is okay—down to Root Zone Maintainer, distributed to several operators, and then further down to the instances and made available. Next slide, please.

The Root Server System is actually fairly old. It's as old as the DNS system because it's been a vital part of the DNS system since the inception. That goes back to 1984, actually. Back then we had only four root servers, and you will see that a number increases slowly.

In the beginning of the year 1991, there were eight root servers, and they were all located inside the United States. And in 1991, in July, the first one outside the United States was deployed. And that was actually ours in Stockholm. And then it grew slowly until 1998 when we were up to 13 which is the current number.

What we started to do then was to deploy more copies of the same servers we already had. So in 2002—or maybe actually earlier—in Japan ... In the late 1990s, anycast started to be used—and I will talk about that in a minute—which made it possible to deploy more servers than identities. And that gave us the factor, the tool we needed to deploy large number of servers. Instead of deploying—going from 13 to 14 to 15 to 16, we could go for 13 to 50 to 200 to 500 and make the

deployment much more quickly and much more broad than we had before. Next slide, please.

This is just a data set. This is just a list of the 13 identities that you can reach, and the IP addresses—both the older IP version 4, the more modern IP version 6—and also the organization that's operating each of these identities. And you will see here that there's actually a large variety in type of organizations. There are commercial companies like Verisign. There are universities like University of Southern California and the University of Maryland.

There are companies that are non-profit and work for the good of the Internet. That would be, for instance, the Internet Systems Consortium and, to some degree, Netnod. We're kind of a mix of commercial and non-commercial. And there are government-operated servers, and there are community-operated servers like the RIPE NCC. And there's also one operated by ICANN themselves.

Now, each of these organizations has a very different organizational structure and very different decision paths inside—how the organization is governed. And that's actually a strength because that means that we are not susceptible to the same type of organizational threats.

And the most important property of the Root Server System is stability. We must continue to operate this system for the Internet to continue to work. So therefore, we want to have very diverse systems in everything. I shouldn't do anything like Hans Petter here or like Wes over there because if we do things in the same way, that's the single point of failure that we don't want to have. Next slide, please.

This is just a picture from a web page called root service.org, and that's a general gateway for information about the Root Server System. We reserve operators, we provide information to that system. It's a common system for all, and it has a map where we put in all of the instances, all the 1500 instances. You can see in a clickable map where they are located.

And this map is very much zoomed out, but if you use Zoom in to that map, it will change and you can see where the individual servers are. But you can see there are hundreds of servers in each part of the world. And altogether, it's more than 1500. So I encourage you to have a look at that page. At the bottom of the page, you'll see the address there, the URL. Next slide, please.

So the Root Server Operators operate under a number of principles. We've agreed on those inside the work RSSAC here, and they pertain to the operation and how we shall relate to the rest of the Internet community.

And you can see that the first principle here is that we very much support the notion that Internet needs a single, unique public namespace. Without that, we will have a fragmented network in no time, and it's essential for Internet that we keep it all together.

There's a good number of principles here. Some of them pertain to how we provide the service, and some pertain to how we relate to other organizations. These are noted in, actually, two documents, but one is easier to access. So if you look at RSSAC and the principles, you will probably be guided to that document. Next slide, please.

So I would like to focus a bit on a few of these principles. For instance, #2 and 4. Number 2 is very important. It's absolutely paramount that we serve exactly the same data because if we don't, we have a fragmentation of the network. So all of the Root Server Operators fully and 100% agree that IANA is the source of information. That's the data that we distribute.

And #4 speaks to what I mentioned, just, which is the diversity aspect. The Root Server Operators are very diverse. We have different types of organizations. We're in different countries. We have different types of hardware and software and different operating procedures. And all of that is a strength because that means that we have very few single points of failure where all of the Root Server Operators are susceptible and risk problems from a single thing. Next slide, please.

The Root Server Operators also need to collaborate and engage with their stakeholder community, and that's what we do right now, for instance. We try to inform. We try to be accessible. We do take part in meetings, ICANN meetings here. That's a very obvious thing, but we all are also ... Not all of us always, but you will find Root Server Operators at network operators meetings or DNS specialist meetings or Internet standardization meetings. And even governance meetings with ...

I am somewhat regularly involved with interacting with the European Union and the European Commission on how they do networking, regulation, and so on. And equally, from other operators, they are involved in how the Internet works. So you will find us in many places, And we hope to talk to you if you have issues with the root service.

Also, the Root Server Operators need to be autonomous and independent. And that also actually connects to the diversity aspect because we need to be autonomous and independent in order to create this diversity. If we are not independent, that means that someone has a grip on the entire system, and that will destroy the diversity that is so, so important for this system. Next slide, please.

There is a good number of myths around the Root Server System working out around there from mouth to mouth, and we would like to try to kill off a few of those. One is that people tend to believe that root servers control where Internet traffic goes. It does not at all. The root servers are part of the set of traffic signs, but we do not regulate the roads. We just, so to speak, tell people where to drive. So what controls the Internet traffic, where the packets go, that's routers on the network. And they interact in a totally different system than ours.

Next, "Most DNS queries are handled by a root server." That is not true at all. Ken spoke to that earlier. The root servers are actually queried very seldom, and most of the queries are handled by the caching resolvers often through their caches so that they can respond very quickly without talking to a root name server.

"Administration of the root zone and service provision are the same thing." No, it's not. Again, the root zone is created by IANA through the Root Zone Maintainer, and the operation of the servers is the job of the Root Server Operators. So we have a very strict division of labor there, and the point where that touches is where the root zone is handed over to the Root Server Operators from this distribution service.

"The root server identifiers have special meaning." No, they don't. The A-root server is not better than the B-root server is not better than the C-root server. They are all identical. They all give exactly the same answers, and they all have exactly the same properties. So the A-root server is not special in any way.

"There are only 13 root servers." Well, that was true in 1997; it's not true anymore. We've tried to hammer that in that we now have more than 1,500 servers across the globe. There are 13 identifiers, but that means that from every point on the network, you can reach 13 different servers. But your friend on a different part of the globe can reach a different set of 13 servers. So that is simply not true anymore.

"The Root Server Operators conduct operations independently." Well, we are autonomous, but we have very close collaboration when it regards operations. We have regular meetings where all of the Root Server Operators participate. We have emergency channels to communicate if something should go really bad. I know Wes has the telephone number to my bedside telephone, and that goes for all of us. So we can actually reach each other really, really quick. And we have alarm systems that would go out to all of our cell phones, and we have ticket systems to manage things. So this is a well-coordinated system.

And the last, "Root server instances receive the entire DNS query." That's not always true. It's true in some cases, but not always. That goes back to Ken's ... He mentioned this QNAME minimization. In traditional old-time DNS, the entire query will reach the root server if need be. Typically, the recursive resolver will handle the query. But if it needs to talk to a root server and it's an old system, then it will send the entire

query. If it's a modern system, it will use QNAME, and we will only see the TLD part of the query. Next slide, please.

So, The Root Server System and Your Network. It's good practice to make sure that you have access to root servers. You essentially only need one because if they are all identical, for resilience and redundancy, it's probably a good idea to be able to reach three or four in your near vicinity. And using network technology, you can actually reach out to the root servers and see how far away they are in terms of milliseconds—how long time does it take for the query to a root server to come back? And you'll probably want to have a look and check that your network has reasonable access to a root service. And if you don't, please come and talk to us because we want to know. We want the service to be good. If it's not good for you, let us know that, and we can have a look at why it's not good and see if we can do something about it.

Do use a DNSSEC validating resolver? That's the part that checks the DNSSEC signatures. If you don't check them, they're worthless. We provide them, and please check them because that tells you that you have the correct data from the root zone that no one has tampered with.

Do participate in DNS technical communities. If you are interested in DNS and want to understand how it works, do engage and reach out because that's a very good way to understand better, to get a network of people to talk to. And that can be very, very useful when things go wrong.

And if you are interested in hosting one of these instances near you, come and talk to us. We are mostly open to deploying new servers. There are obviously restrictions for how and when and what we can do, but we want to know that there is a request for service somewhere so that we can take that into our plans and see if it's feasible. It's not always feasible, but in some cases it is.

And just to finish off this part—next slide, please—a traffic picture of the combined traffic to the root servers over time. You can see it starts in 2017, and the source for this data is statistics that we all provide. The Root Server Operators provide query statistics so that you can pick this up.

If you want to look at the traffic volume to the root server that Netnod provides, you can go to our websites and pick up the data and generate these graphs. And by doing that for all 12 organizations, you can combine this into this picture. And you will see that this is a large number. I guess it's ... I would guess—and please help me here— that the number on the Y axis going vertical here where it says "queries" is 150 billion queries per day, but that's a guess. Next slide, please.

I have a few slides here on the concept of anycast. So next slide, please.

We've been talking about that when we step from 13 servers and started to expand the numbers greatly. The traditional way, going back to the 1990s is called unicast, and that's you have a single server with a unique IP address and anyone in the world who wants to talk to that IP address will send a packet to that very machine. So that's what we had in back in 1998 when we had 13 servers on the entire Internet that provided this service.

That means also that if this server is attacked by a denial-of-service attack where you send a humongous amount of packets to the same server so that it's totally overloaded with traffic that would hit this single machine. And in the end of the 1990s and early 2000s, we started to discuss expanding this and using a trick, essentially, which is called anycast.

Anycast means that you have multiple servers in different locations. Each and every one has the same IP address. That's not supposed to work, but it actually does. And I like to compare this with the emergency number that you can dial on your phone. In Europe, you will dial 112. If you're in Britain, you'll dial 999. Or if you are in the U.S., you will dial 911 to call for an ambulance or fire, help, or the police, or something.

When you dial that number ... And let's say you're in the USA. If you're in New York and dial that number, you will be talking to a service center, the call center for emergency calls that's located in the New York area. But if you fly over to Los Angeles and dial the same emergency number, you will be talking to a call center in the Los Angeles area. And that's exactly what anycast is about, but on the Internet. So next slide, please.

In the old unicast model, everyone on the Internet, being the source in the picture, would reach the same destination. There was only one destination to reach, so the 911 number would have only one call center in the U.S. and possibly located in Chicago in the middle. What do I know? But that's not what we want to have today. So next slide, please.

So what we do is that we put up several blue blobs here of several destinations, several call centers—or in this case instances—for a certain IP address. And the source will reach the nearest one. So when

you send out your query, the network—the gray and green blobs in the middle—they will actually talk about who's the nearest one and send your traffic.

Again, this is the routers sending the traffic, and they will automatically select between them which is the nearest one and send the traffic there. And my last slide. Next slide, please.

And this means if you have a denial-of-service attack where people who are rather bad, bad people, bad computers, send large volumes of traffic, enormous amounts of packets to the same destination. An attacker, the red blob here, will only reach the nearest instance. So they may affect the top right corner destination, but it will not affect the bottom left destination.

Again, you can compare this to the call center. If you're in New York and 2 million people start to call the emergency number, that call center will be overwhelmed. But the call center in Los Angeles has no problem, so the people in Los Angeles, they see no effect of that DDoS attack. So the call center analogy kind of works here, too.

That's the end of my part. I would like, here, to hand over to Wes [and you].

LARS-JOHAN LIMAN: All right.

JEFF OSBORN: Wes?

WES HARDAKER: Yes.

JEFF OSBORN: Can you indulge me for a minute? I'm not sure what the protocol is, but there's something going on in the chat that I really feel like needs to get addressed before it—

WES HARDAKER: Okay.

JEFF OSBORN: —sits there for too long. Dwayne, you walked in in the middle of a question, which was completely unfair.

DUANE WESSELS: It was.

JEFF OSBORN: But if you want to dial back to about 11:01 local time in the chat, there's a misleading or misunderstood statement being made that I really can't leave hanging there. And if you'd, like I'll read it.

It was the response to what you said, and the note here is, "Thank you. The answer is: We are at most two weeks away from the full breakdown. And the second part of the answer is if we have a global problem which brings down travel, we have at most three months to solve the problem. We scratched this limit during Corona."

I will argue that's not just misleading, but it would be bad if that weren't addressed, I think. So, again, I apologize, Wes.

DUANE WESSELS: Thanks, Jeff. So as I understood, the question was how would you deal with potential failure of the key signing process. And like I said, we sign those. We generate the signatures with a two-week signature validity, so that gives us two weeks to remediate the problem. I think that's very different than a full breakdown of the Internet, the root zone. Right? I mean, we feel very comfortable that two weeks is plenty of time to address any problems that might arise. Does that clarify ... Does that ...

JEFF OSBORN: I believe so. We'll—

DUANE WESSELS: Yeah, does that—

JEFF OSBORN: —work through the chat.

DUANE WESSELS: Yeah, I—

JEFF OSBORN: But thank you very much for interrupting, and I appreciate it. Wes as well.

DUANE WESSELS: Yeah, right. So as Wes is sort of hinting at here, so certainly ... Again, speaking for the root zone maintainer role, we have recovery plans and backup plans, and so we ... The signature validity period of two weeks gives us that two-week period to act on those plans and implement the procedures that we would if there was an unexpected failure in signing. So thanks.

JEFF OSBORN: Again, to put it tersely, it's difficult to imagine not fixing such a simple problem in two weeks with so many people in so many places and so many resources to do it.

WES HARDAKER: Yeah. I think the last point I was whispering into Duane's ear is that two weeks assumes that you haven't executed any of your backups or any of your other plans. That is, you know how long you would have to notice. So we have plenty and plenty of redundant backup plans.

All right. My name is Wes Hardaker. I'm over here in the corner, for those that can't see me. I'm from the University of Southern California which is actually where the DNS was invented way back in the day. And I'm going to talk about RSSAC and the RSSAC Caucus. So this is really where ICANN kind of comes into play. Go ahead and go to the next slide.

The Root Server Operators all operate with ICANN, and we do that primarily through the Root Server System Advisory Committee. And its purpose is to advise ICANN and the Board and other elements of ICANN about how the Root Server System works and answer any questions. As an example that came up earlier today, there's an RSSAC publication

that directly answers a question from during the development of SubPro of, "Can you guys handle this many more new TLDs being added?" And we wrote back formally in the ICANN way and said, "Yes. Not a problem. Don't worry about it," and suggested our limits going on for the future. So we answer any questions that come our way.

But this is a very narrow scope. We do not delve into ... We deliberately do not delve into policies that might affect whether TLDs should be added or what they might look like and things like that. That is up to the ICANN multistakeholder system. We only provide advice related to what can the Root Server System handle in all of those policies. Next slide, please.

So RSSAC, you know, is a committee. We produce advice. We'll talk a lot more about that in a minute. And it's primarily ... Like everything in ICANN, most advice goes to the Board and then gets reassimilated to parties of interest. But we write to multiple places, as I said.

There are 12 Root Server Operators, as has already been discussed. And every Root Server Operator is represented in RSSAC. And RSSAC is only on the policy ICANN side. It is not an operational body. We don't talk about how the servers work and changes to protocols and stuff like that. That is handled in a technical forum that is also represented by all members of the Root Server Operators. And we do that on a regular basis, too.

We actually meet three times a year. Well, we meet monthly as RSSAC, and we meet three times a year in person as technical operators as well. We actually communicate seemingly nonstop about things that are happening. Next slide, please.

These are the leaders of the RSSAC. You've met two of them already, the two that are on this slide. So the RSSAC chair is Jeff Osborn who was just talking a moment ago, and the vice chair is Ken who is over there sitting next to him. Next slide.

So RSSAC is composed of two representatives from each Root Server Operator. So we even have redundancy in our membership. If the primary member is not available, there's always an alternate that is able to step in and vote or do whatever else is needed within the ICANN context. And then we have liaisons from other bodies that also participate in ICANN. And I'll talk more about those on the next slide.

We also have the RSSAC Caucus, and I'll describe this in greater detail in a minute, too. This is a body of volunteer subject-matter experts, people that come that want to help out the creation of especially technical advice that comes out of RSSAC. And anybody can become a member of the RSSAC Caucus. All you have to do is submit a safe Statement of Interest, and what your background is, and how you're going to contribute. So we encourage anybody that has an interest in Root Server System Policy, to come help us out. I'll talk more about that in a minute. Next slide, please.

So RSSAC, besides representatives from each of the Root Server Operators, also has a number of incoming liaisons—people from other organizations that participate so we can make sure that we're well-coordinated with lots of other organizations out there. That includes the IANA Functions Operator—you know, the people responsible for actually the creation of the root zone and what goes in it—the Root

Zone Maintainer, Verisign, which you just heard from Duane a minute ago sitting to my right.

The Internet Architecture Board. Daniel Migault is the incoming liaison from the Internet Architecture Board which is really a liaison from the entire IETF. He's responsible for making sure communication happens with the standards body that produces the DNS protocol.

I'm actually also a representative ... I'm on the Internet Architecture Board, but I try and let Daniel be the official representative for [inaudible] cases.

And then also the Security and Stability Advisory Committee. Russ Mundy is our incoming liaison from that. That is another advisory committee within ICANN that looks at all security-related discussions within ICANN. And we have a very close relationship with them. We meet them every ICANN and discuss things with them.

Outgoing liaisons. There's an outgoing liaison from RSSAC to the ICANN Board. I am actually that Board of Directors member, recently elected. We have an ICANN Nominating Committee, so we have representatives. Some of these are actually from the Caucus. They're not actually necessarily Root Server Operators that go as outgoing liaisons.

We have an outgoing liaison to the Customer Standing Committee and an outgoing liaison to the Root Zone Evolution Review Committee as well, trying to stay tightly integrated with everything. Next slide, please.

So I'm going to talk about the Caucus next, and I'm going to reverse the order on these slides. I realized when I was reviewing them that it actually makes more sense to read the second part first.

What is it? Why does it exist? Right? It is a bunch of DNS experts who bring your expertise to talk about the things that might affect the Root Server System. Almost all RSSAC publications these days are actually not written by RSSAC proper, but are written by the RSSAC Caucus. And then the RSSAC actually formally publishes them.

We did this deliberately to try and bring in more and more outside expertise into the documents that we write. So it's not just written by the 12 RSOs. It's written by the entire ICANN community since anybody can come and participate. It gives us transparency. Our mailing list is open, and it's archived. And it's a real framework for getting stuff done. And we do this through work parties that are formed to either create new documents or update past ones.

Aside from administrative documents, as I said, pretty much all technical documents produced by RSSAC come through the Caucus process. It consists of ... We actually have over 100 experts at the moment. Not all of them are as active as others. It depends on the topic that they're interested in. Some people are very active on some of our work parties and less active on others because each of us have different interests. It's required that they do a Statement of Interest, which is a common framework in ICANN for stating why you're interested in being on it and what your background work is.

And then the nice thing about it is that if you help contribute to one of the RSSAC publications, we do add the names of everybody that helped contribute to the bottom of the RSSAC document. So you can go look in any of the RSSAC publications and see a list that it's not just the

operators that are in there. It's actually the entire ICANN community that helps contribute to these. Next slide, please.

So we're going to dive into the Root Server System Governance Evolution, and we'll jump immediately to the slide after that. So back in 2016-ish, we got together is RSSAC members and decided the person responsible—Jon Postel—for actually helping create ICANN in the very beginning, when he passed away, he had given the last of the Root Server Operator assignments out to various organizations. WIDE was one of the last ones, some of the later. RIPE was one of the last ones. And there was no procedure for who was going to take over that process.

So the 12 operators today have been operating for a very long time since he passed away 20 years ago. And there was no process for what do we do if we need a new one or if we need to—if we technically need a new one or if we need to remove one for some reason or another.

So we decided we should fix that, and we have an internal process for how we negotiate and communicate between the RSOs and an existing governance framework. But we really wanted to start opening this up to two external frameworks, so we started a document that produced RSSAC037. In 2018 is when it was published. And it was a goal. It was sort of a framework for, okay, how are we going to do this with a way that has membership of all of the different ICANN stakeholders to bring them together to add or remove ourselves, for example, or make any changes to the Root Server System that might need to be made at a policy level at least.

And here is actually the principles that Liman was talking about earlier. Eleven stated principles are included in there of how we exist today. From that, the ICANN Board directed that a new working group be created which has an even broader set of participants in it called the Root Server Governance Working Group. It has met here four times this week, and there are two more meetings to go later today.

We have had two workshops, I think, and we're going to meet in person in another workshop in a couple of weeks because it takes a lot of discussion to get this done. And we're developing a governance model that, based on these 11 principles, we'll create a governance model that allows the ICANN multistakeholder system to better govern the Root Server System in a way that is more inclusive.

As I said, we've always had mechanisms for making sure that the Root Server System is stable and secure, not only technically but operationally and policy-wise as well. Let's go on to the next slide for now.

It is a public forum. I'll talk about the membership here in a minute, which is why I skipped that line. The stakeholders that we identified in RSSAC037 are three. Right? One, obviously the entire ICANN community. And the nice thing is that the ICANN community pretty much incorporates everybody in the world. So you're all a member of the ICANN community and a stakeholder of the Root Server System. The RSOs that operate it obviously should have a say in making sure that things work well and are produced in a way that makes sense.

And the IETF is the Internet Engineering Task Force, which is represented through the Internet Architecture Board as a liaison role.

And they're the people that produce the DNS protocol, so they should have a say in making sure that the Root Server System adapts to their vision of how the DNS works. Next slide, please.

So the Root Server System Governance Working Group, on behalf of helping out those stakeholders come up with a model that's going to work from a governance point of view, is composed of many different members. And it includes people from inside ICANN like from the Board, and from the Root Zone Maintainer, and from IANA. It includes members from the Registries Stakeholder Group, for example; one member from each of the RSOs; members from the ccNSO; and, again, two members appointed by the Internet Architecture Board on behalf of the IETF. Next slide.

I tried to rush through this a little bit because I really want to make sure that we have a little bit more time for questions. I know there's a couple of them queued up, and I'll let Ozan speak to those in a minute.

But as a last point of information, if you want more information, the main RSSAC page, of course, we have a list of all of our publications and a lot more about what RSSAC is. We have a Frequently Asked Questions because we get a lot of them. So we do repeat them. So it's a good place to read and get a little bit of background material if you haven't read it yet. And then there's actually a general mailing list for asking questions, and that's ask-rssac@icann.org if you have questions about the Root Server System at all.

As far as the Caucus goes that I mentioned, the RSSAC Caucus, there's a main webpage for that; as well as if you want to apply to become a member and help us out, we'd greatly encourage it. And you can write

to RSSAC-membership@icann.org, and ICANN staff will help walk you through the process of submitting your Statement of Interest to become an RSSAC Caucus member. And again, we greatly encourage that.

So with that, I think, Ozan, I'll turn over to you. I believe that there's some questions in the Zoom chat that we may want to start with.

OZAN SAHIN: Sure. Thank you, Wes. We currently have three questions in the Zoom room—two from Abdullah Qamar, so I'll start with Abdullah's first question. "Is there any blockchain or DAO application on root server or any case study?"

WES HARDAKER: So I'll answer that because I've done work in this area in terms of academic research. I work for a university, so I do a lot of academic research associated with the DNS in general.

The DNS doesn't need a blockchain per se. One of the advantages of having 1500 root server instances around the world is that it's already distributed. So the benefit of doing blockchain-like technology for doing distribution isn't well as needed.

The IETF has produced a document called RFC 8806 which allows you to run a local copy of the root zone, or it tells you how to do it in your house or locally to your environment, if you wish, as well. There's a lot of work by various blockchain organizations such as the Ethereum blockchain. The GNU DNS effort has been recently pushed forward.

There was an older one called Namecoin which I don't think it's used as much anymore.

And in the IETF, they are currently producing an RFC coming out to be soon that will make a new special reserved TLD called .alt for where to house these alternate namespaces. You have to remember that the blockchain-related names are not part of the DNS. They're very, very different. And so how do you create a name that doesn't conflict with ICANN registries and things like that? And so the way to do that that the IETF is recommending us to put them under .alt. So that's coming out soon.

The ICANN Board just got notified of that document yesterday. They knew about it beforehand, but they got officially [notified] about it yesterday, and it's going through the IETF last call process right now.

OZAN SAHIN: Thank you, Wes. So the second part of Abdullah's question, "What are the current research areas in root server, and is there any e-mail list for academia researchers." Does your answer also cover that?

WES HARDAKER: I didn't cover it very much. You're welcome to reach out to me. It's my last name @isi.edu. For anybody that wants to write me. We do a lot of academic research related to the DNS in general. The RSSAC Caucus has actually been discussing the last six months of creating a list of all of the academic publications that have occurred related to the DNS root. It turns out to be over 100 long. I've lost track of how many we've put

together so far. That work hasn't concluded, so we haven't actually produced that list formally in a [referenceable] place yet.

But you're welcome to rate me personally if you want, and I'll help tie in the academic side of the world.

KEN RENARD: For those interested in research, the RSSAC and all of the Root Server Operators publish data that is described in RSSAC002. That information is readily available as well as, I think, almost all Root Server Operators participate in the OARC DITL, the day-in-the-life.

So we publish essentially a 24-hour period or even more of queries and responses to a root server, usually with anonymization. And those are places that you can go if you want to do research on the data. So the RSSAC002 data—Andrew put together a great page. Go to that root-servers.org website, and I think the data is linked in there somewhere. It's called RSSAC002. Thanks.

WES HARDAKER: Yeah. Very good point. So DNS-OARC is where a lot of research is done even at early stages, and the website for that is dns-oarc.net. And that's where Ken was talking about the data that's housed.

All right. You have one more, Ozan, I think.

OZAN SAHIN: Yeah. And we have just received another question. That seems as a follow-up from Jean. So I'll first go to this one, and then go to Jeng-Wei Lin's question.

So Jean's question is, "About that upcoming IETF RFC on .alt, does Wes if it will be discussed in IETF 116?"

WES HARDAKER: So the DNSOP Working Group at IETF 116 would be where to go to discuss it. It is already passed Working Group Last Call, so the DNSOP Working Group has already said that the current version of the document is sufficient. So it's now in IETF-wide last call, but I have no doubt it will be discussed at least somewhat in IETF 116 in the hallway, if not in DNSOP itself.

OZAN SAHIN: Thank you, Wes. The other question we have in the Zoom room is from Jeng-Wei Lin, who is an ICANN76 Fellow. And he says, "I believe in the RSS robustness. However, it is different from its [availability to a particular] set of users. Is there any study or report or what we can learn from scenarios such as Tonga volcanic eruption in 2022? Thanks very much."

WES HARDAKER: I was trying not to hog the microphone, but this is Wes Hardaker again. There has been a lot of study. A couple of things to note. RIPE has the RIPE Atlas probes that measure the health of the Root Server System from all around the world with their RIPE Atlas probes that are

distributed that people can plug into their networks. And that's a great place to do your own research to figure out how far away or what availability [inaudible].

My colleague John Heidemann at ISI actually recently published a paper about measuring each of the Root Server Operators identifiers and whether they're reachable or not, and actually concluded by leaps and bounds that all of them are available from everywhere, that there's almost no break in connectivity. And he found a couple of places, and it's almost always the local network's problem not a Root Server Operator problem. But he's done some amazing research that's [inaudible].

There's a couple of peninsulas that you can get from one place to another, but you have to go through a weird root or something like that from a technical level. But everything's available, basically.

KEN RENARD: Yeah. There's always a last mile problem if you're on a remote and there's a single or very small number point of failure. You may not be able to get to incredibly large parts of the Internet if your local path fails. Now, there's a couple of ways to solve this. We could put root server instances in everybody's home, but that doesn't scale well. That's pretty expensive. Another way that was mentioned was—what Wes said—the RFC 8806 where you can get copies of the root zone in your local recursive resolver. So that can help for that being cut off from the world scenario at least a little bit.

I'll end the question, but I had another comment later on.

OZAN SAHIN: So we have one more question in the Zoom room coming from Wataru Ohgai. And the question is, "Can young students who have fewer experience compared to others be eligible for Caucus membership?"

JEFF OSBORN: I'm actually the chair of the RSSAC Caucus Membership Committee, so I can speak to this. We have, as has been stated, over 100 people currently, and we had consciously made the decision over the last two years that with so many people willing to help, we've raised the bar a bit. And after a lot of discussion, we're accepting membership from people who have actual hands-on DNS experience because it's very difficult to operate something as an educational thing when what you're looking for is broad expertise from a range of areas.

And what we found is when you're trying to have a fairly deep discussion and people are stopping to ask much more introductory questions, it makes the process difficult. So if somebody has a proposal about increasing the amount of outreach and education that we're able to do—there are various parts of ICANN in my organization, other people's places—that would be interesting. But I think the educational opportunities within the ICANN Caucus have been deemed to not be helpful.

So we're really looking for somebody who has spent a lot of time in the DNS structure in your home country, or workplace, or university. But the bar is such that without a lot of years of doing this, you'll get a polite note from me or Ozan saying, "This probably isn't a good fit. Can you

explain further if you have something you bring to this that you think is valuable?" But this is not an entry educational group. Thanks.

OZAN SAHIN: There's nothing to add. I guess we're out of time. Okay. Over to you, Jeff and ... Please go ahead.

DANIEL MONDRAGON: This is Daniel from NIC Costa Rica. We own the TLD. Is the RSSAC group planning to use a stronger algorithm like elliptic curve, for example, for signing zones in a future major rollover [inaudible]? Will it affect the TLD operations?

KEN RENARD: Just one thing, and this goes for the discussion on the key failure scenario question that was brought up earlier. Just a distinction between the Root Server System and the other parts of the root zone. Again, the Root Server System is the delivery that the servers on the Internet that do this. The questions about the key, the cryptography failures in that are outside of the Root Server System but very closely related.

And then I'll hand over to Duane or Wes.

DUANE WESSELS: Thanks, Ken, for that preface. So ICANN has recently announced that there's a design team that's beginning to study an algorithm role for the

root zone. So if you'd like, I can point you to that information. But it's in the design phase at this point. Yeah.

WES HARDAKER: One last bit of information. I will be presenting at the DNS Security Workshop the measurement study that Victor Duchovny and I take on that actually measures the algorithm usage in all of DNSSEC. That will be Wednesday, later in Wednesday. I forget the exact time slot, but we actually do measurements. So elliptic curve is actually definitely one of the top leading signing things in everywhere else except the root, but there's a number of different algorithms.

One last point of order before we close. I have to run off to talk to some students because I do work for a university. I have a bunch of business cards. If anybody has questions that they want to write me about later, I'd be happy to hand one as I kind of walk out the door.

OZAN SAHIN: Jeff, we have a hand from Brad.

BRAD VERD: Yeah, Jeff. Can I add just a response to that?

JEFF OSBORN: Brad, of course. Go ahead.

BRAD VERD: In response to that question, I think it's important to note that while you've been given an answer, I think it's important to note that—Ken

kind of touched on it—the RSS, the Root Server System and the RSOs in particularly have no influence over the keys that are used in the signing of the root zone. That is all done with IANA and not with RSSAC or the RSOs or the RSS. So I think it's imperative upon us to get that message across.

JEFF OSBORN: Thanks, Brad. Good point. Lars.

LARS-JOHAN LIMAN: Thank you. I want to just put in two things here. The first one is, all RSSAC meetings and these Governance Working Group meetings, they are all open to observers. So if you are interested in the work that's going on, you are welcome to dial in to these often. They are telephone or Zoom meetings. Please come and listen and learn more.

And the second one is that we are, I think—I dare speak for the Root Server Operators—we are interested in talking to you. So if you have more questions, just stop me in the corridor. I may be big, but I don't bite. I promise. Thanks.

KEN RENARD: Yeah. And a closing comment. If there's any feedback on this session—too much detail, not enough—I'd love to hear it in the room afterwards. Thanks.

OZAN SAHIN: Thank you, everyone, for joining today. I guess we are adjourned. Please stop the recording.

[END OF TRANSCRIPTION]