BARCELONA – How It Works: Root Server Operations
Saturday, October 20, 2018 – 15:15 to 16:45 CEST
ICANN63 | Barcelona, Spain

UNIDENTIFIED FEMALE: Good afternoon, everyone. Welcome to our How It Works Tutorial on Root Server Operations. Starting the session is Andrew from ICANN. Go ahead, Andrew. Thank you.

ANDREW MCCONACHIE: Is it on? There we go. So, I'm Andrew McConachie. I work for ICANN. I help support the Root Server System Advisory Committee, the RSSAC. This is the How It Work Tutorial session on the root server system. If that's not what you're interested in, you're in the wrong room. But, if that is what you're interested in, then please stay. This is ICANN 63. Welcome. Do I just call next slide? Can I call next slide? Thank you. So, this is the tutorial on the root server system. Next slide, please.

This is going to be a bit of what we're talking about today. I'm going to talk for the first half and then my colleague, Carlos Reyes, is going to talk for the second half. I'm going to be talking about the domain name system, a little overview on that. Then I'll go into the root server system today, a bit of its history and some of its features. I'll give a quick explanation of Anycast and how it's used in root server operations. Then I'll hand it over to my colleague, Carlos, who is going to talk about the RSSAC and some of the more recent activities of the RSSAC. Next slide. Great.

This is the overview of the DNS and root servers. Next slide, please.

So, a quick recap on identifiers in the Internet. IP addresses, what exactly are IP addresses and why are they important? They're the fundamental identifier on the Internet. Every device that wishes to speak on the Internet needs an IP address and all those connected to the Internet must have IPs.

It's a numerical label. It's kind of hard to remember which is one of the reasons why we have DNS, because people don't like remembering IPs. And there are two versions, IPv4 and IPv6. That example there we see on the top an IPv4 address, and on the bottom an IPv6 address. Next slide, please.

So, why do we want DNS? Well, based on the last slide, you kind of saw we don't really want to remember IP addresses, so we have DNS. First, we would host files which got really, really hard to maintain. So then somebody said, "Well, maybe we should distribute this." And thus, DNS.

So, the original problem was that IP addresses are hard to remember and they change a lot, so we can't just have host files. The modern problem of why we have DNS.

IP addresses can also be shared. There's this thing called network address translation and load balancers and also Anycast, which can result in sharing of IP addresses. And multiple IP addresses may serve as an entry point to a single service. So, which one to use? Next slide, please.

So, the domain name system is hierarchal. I think people are mainly familiar with this. Everyone kind of understands and everyone has kind of interacted with domain names a little bit, even if typing them into an address bar. But you may not have ever really thought about how the hierarchy of DNS breaks down. And there is a root zone at the top. That's what we call the top [bit]. Then, underneath that, you have what are called top-level domains. Dot-com is a very famous one, but of course you've got dot-edu, dot-mil, dot-uk, these examples here.

Then, underneath that, we have what's called the second level. Sometimes, they're called [2LDs]. They're just second-level domains. Then, of course, we also have the third level. Many people are familiar with www which is probably the most common third-level domain name that people would be familiar with.

There's an example.org mapping to an IP address once again highlighting the importance of DNS and highlighting names to IP addresses.

There's many other uses for DNS, such as mail servers, such as reverse lookups, and also for IPv6 addresses. So, DNS is not just used for mapping domain names to single IP addresses. There's a lot of other uses as well. Next slide.

So, this slide kind of maps the whole resolution of a name and it also goes into the signing and validation of that name, the whole DNSSEC process. I'm going to spend a little bit of time on this slide just walking through the whole process of how a name resolves.

It kind of starts on the right here where we have our user who's sitting at a computer and they really, really want to go to the web server at www.example.com so they go to their recursive name server which is what they have configured on their local computer. They go to their recursive name server and the recursive name server, for the sake of this example, knows nothing. It has nothing in its cache. Maybe it's just been turned on. It needs to start from the beginning.

So, it sees www.example.com and it goes, "Well, I don't know where dot-com is. So, I'm going to go to root name server because that's the only thing that's configured in the recursive name server. So, it goes to the root name server to find the location of dot-com, to find the NS records for dot-com.

It gets back the location of the name servers for dot-com and it also gets back a signature, a DNSSEC signature, which it can then validate. And we'll get into DNSSEC a bit more in this presentation but that's basically what happens in that step.

Now that it has the names or the location of the dot-com servers, here's the dot-com servers and it says, "Hey, where's example.com?" Again, it gets back a response with the location of example.com along with a signature. It validates the signature, makes sure it's good, and now it can go to the example.com name servers and it does the same thing.

It says, "Where is www.example.com?" And the example.com name servers come back with the location of www along with the signature. It then validates that signature and then it can finally go back to the user and say, "Here you go. This is the location of www.example.com.

And not only this is location, but hey, I validated it, too. Isn't that cool? So, yeah, that's a very basic introduction to DNS recursive resolving. Next slide, please.

So, as you saw in the last slide, the root servers only need to know who needs to be asked next. So, the only thing that's in the root zone is the location of the name servers for the TLDs. It's not [names] … Underneath those TLDs, just the names in those – the locations for the name servers of the TLDs.

Now, as we saw in the previous example, our recursive server in the previous example didn't have any information. It was starting from zero. Zilch. Its cache was empty. In most instances, that's almost never the case. There's a lot of caching that goes on in recursive resolvers. They remember the locations of these name servers and they time out over time based on TTLs and what not, but there's a tremendous amount of caching that goes on, so that the recursive resolvers don't have to always go back to the root all the time. They know the locations of these servers because they remember them from the last time they queried them.

So, what that means for the root servers is they actually don't get queried all that much, only when the recursive resolvers don't know the locations of the TLDs. Next slide.

So, that was a pretty … The last two slides were really kind of a basic introduction to DNS and DNSSEC. This is more about some modern enhancements to DNS, just so we can be familiar with some of the

terminology. DNSSEC, the DNS security extensions. You'll hear a lot of talk about that at ICANN.

What that is is DNSSEC is basically there's two sides to it. There's the signing side and there's the validation side. On the signing side, the names that are being returned in responses to queries are signed. They come along with signatures. Then, what the recursive resolver does is it takes those signatures, along with the name, and it validates them. And by validating them, what that means is that based upon the trust anchor of the DNSSEC root which is sometimes called – you might hear talk of the recent KSK rollover and whatnot. But, based upon the trust anchor configured in the recursive resolver, the recursive resolver is able to validate that response and know that it's the correct response. The purpose of this is to reduce the risk of spoofing so that people can't lie to the recursive resolver. That's the entire point of DNSSEC. It prevents lying to recursives.

There's also been some recent privacy enhancements to DNS, specifically DNS over TLS which is a relatively brand new protocol from the IETF. This helps preventing the leaking of query information to passive observers and this [inaudible] standards are being created to reduce this, which is true. There's active standards development in this area. It's an ongoing area of work for the IETF.

The third box here, Anycast. We'll have an explanation of Anycast later in this presentation, but basically what Anycast allows you to do is it allows multiple servers to share a single IP address. So, from the point of view of the recursive, there's one IP address but it's actually hitting

an entire constellation of servers and it doesn't know and it doesn't really care which server it's hitting. It is just using one IP address. So, you can think of Anycast is just many, many servers all using the same IP. That has some interesting effects. It really does vastly improve latency and resilience. It also helps protect against DDoS attacks. There's a whole section on Anycast which we'll get into later. Next slide.

So, there's this difference between the root zone, which is the data, and the root servers which serve the data. The root zone you can think of it as … You remember I talked about the starting point, the list of TLDs and name servers, the locations of the name servers for the TLDs. So, this is the data. This is the data that the root servers serve. This is managed by ICANN and it's per community policy. It's compiled and distributed by both the root zone maintainer and then distributed to the root server operators.

I just emphasize once again this is the database content in the root servers. This is not the root servers themselves. Then, on the other side, we have the root servers which are the ones which respond with data from the root servers. The root servers listen for the queries and they respond with answers from the root zone.

There are currently 13 identities and over 1,000 instances at physical locations worldwide and it's Anycast that allows there to be just 13 identities and over 1,000 servers.

The root servers have a purely technical role. They respond to queries. And each one is the responsibility of the root server operators. Next

slide. Can we just do next? Do we have a lot of those? Huh. I don't remember what those were exactly, but we'll just go on.

So, the root server system today and its features. This slide has a little bit of the history of the root server system. You see it beginning in 1983 with four addresses and then it grows to seven addresses in 1987 and then eight and then nine and then it stabilizes in 1998 with 13 identities.

These changes over time were largely responding to technical demands. These were all changes that took place before people were using Anycast and before there was really … Well, I don't think there was any threat of things like DDoS at the time, not in the Internet of the 90s. So, this was kind of a different time and things evolved.

Today, we again have 13 IPv4 and IPv6 address pairs which we generally call identities. We have over 1,000 international instances around the world. Next slide. Next slide. Sorry about that. I don't know what's going on there.

These are the root server identities today. We have 12 different operators and 13 identities. You see each one on the far left. We have the host name and on the far right we have the manager. In between the center column, there's an IPv4 address and an IPv6 address. These are the addresses that are going to be configured in what's called the root hints file which the recursive servers will have, so they can then find the root servers, either via IPv6 or via IPv4. Next slide.

So, this is the distribution. This is how the root zone gets distributed to the root server operators. On the far left, we have a TLD operator, so

this would be someone who manages a TLD and needs to make a change to where their TLD name servers are.

So, what they'll do is they have a change request that they put into the IANA function. IANA takes that change request, verifies that change request. And this could be something like their name servers, maybe the names of their name servers changed or maybe more likely the IP addresses of their name servers changed. Remember, these are the IP addresses of the name servers for the TLD. They change.

Then, the IANA function will prepare the root zone and send it to the root zone maintainer who then distributes it to the various root server operators. Then, here on the far right, we have the DNS resolvers. We talked about them as recursive resolvers and you can see them sending queries and getting responses to the different root servers operated by the different root server operators.

At the top here, we have a distinction being made between RZERC and RSSAC. RZERC is the Root Zone Evolution Review Committee and the RSSAC is the Root Server System Advisory Committee. You can see it's meant to show the responsibilities of each committee. So, one is about the zone and the data and the other one is about the serving and responding to queries. Next slide.

So, some futures of root server operators. There's a fair amount of diversity within the different operators. They have different organizational structures. They have different histories and when they became root server operators or how they operate. They use different hardware and software. This might mean that they use different

operating systems or they use different versions of software packages or different software packages, so that there's a certain amount of resiliency through diversity. They also have different funding models. They receive funding through different ways. Some of them are non-profits. Some of them are for-profits. But, throughout that diversity, they have some shared best practices and high physical system security, over-provisioning of capacity in case of DDoS and professional and trusted staff, just best current practices. Next slide.

These root server operators cooperate through a lot of industry meetings, through a lot of meetings such as this one, ICANN meetings. Also at IETF meetings or NOG meetings such as RIPE or NANOG or research communities such as DNS-OARC. They use Internet-based collaboration tools to communicate so that if there are issues they can communicate with one another about operational things. They share a goal of transparency. And they coordinate through [inaudible] infrastructure changes and emergencies. They use things like mailing lists, just your standard communication tools that anyone would use to communicate and respond to issues.

They do host periodic activities to support emergency response capabilities. Again, they coordinate through these established Internet bodies such as ICANN, RSSAC, IETF, and DNS-OARC. Next slide.

Okay. These are some common myths that people have about root servers and the root server system. The first myth is that root servers control where Internet traffic goes. This is not really true. Routers control where Internet traffic goes. What root servers do is they tell you

the addresses of the name servers for TLDs and that's mainly it. Its routers distributed across the entire Internet on basically the ends of every link that decide where packets actually go and I think there was a presentation here previously by Alain Durant of OCTO from ICANN kind of explaining how routing works.

Second myth. Most DNS queries are handled by root server. Not really. Because of caching specifically, recursive resolvers don't have to look up the locations of name servers and root servers all that often because they cache them. They remember the responses, so they don't have to query that much.

Administration of the root zone and service provision are the same thing. No. We've seen that the administration of the root zone is handled by different community, different group, than handles the service provisioning. Remember that slide between the RZERC and the RSSAC.

The root server identities have special meaning. No, they don't. They're all just pretty much the same. The letters that are assigned to them, they're letters.

There are only 13 root servers. No, there are over 1,000, but there are 13 technical identities. With Anycast, remember you can have many, many servers answering on one IP address.

The root server operators conduct operations independently, and actually they coordinate quite a bit. They do all operate their own servers, but there's a lot of coordination that goes into that.

The final myth, root server operators only receive the TLD portion of a query. Actually, they receive the entire query, although there is some recent work called [queue name anonymization] which is trying to change that and that has been deployed a little bit, but that's still, we'll say, in deployment. It's not done yet. So, for now, usually the entire query gets sent to the root servers. Next slide.

Okay. So, now, we're onto the explanation of Anycast and how Anycast works. Next slide.

So, first off, we have Unicast and Unicast is what people generally think of when they think about traffic, fooling around the Internet. You've got one source and you've got one destination. Packets from all the sources go to the same destination. So, you've got a destination IP address. It has one server attached to it and all the traffic sent to that IP address goes to that server.

Now, the problem with this in DDoS – we've talked about DDoS earlier. But the problem with DDoS is that all the attack traffic will also go to that server and so you have one server taking a whole lot of traffic and it might just crash into that.

So, what Anycast does is you have multiple instances serve the same data to all sources, so you have multiple instances behind one IP address. There's one identifier, but there's a lot of different servers.

What this does, it does two important things. One is that it impacts latency, so sources get the data faster. The source is able to, based on

routing, get to the server that is closest to it, topologically, which usually relates to closest to it geographically or physically.

What this also means is that DDoS attack traffic will go to the closest instance. It will be sink holed, as some people will say. So, if there's a DDoS attack, the attackers, because they're all sending traffic to the same IP address and there's many, many servers answering that IP address, the attack traffic, as well as legitimate traffic, will go to the closest destination from the source. So this should hopefully distribute both the attack traffic and the legitimate traffic. Next slide.

So, here's digging into this a bit more. We see Unicast where the traffic takes the shortest route to a single destination. Next slide.

And here's Anycast. Now, you see we have multiple destinations. They're shown in the blue. Then the source again is shown in I think that's kind of an orange-ish green, but I'm color blind and kind of stupid, so maybe it's green. I don't know. But, the destinations are blue.

We see the intermediate routing policies of these other circles don't really matter so much. All that matters is that there's a destination that's really close to the source and that gets all the traffic. Next slide.

Here we see a DDoS attack going on in an Anycast cluster. Again, we see this source on the left going to a destination that's close to it and we see the DDoS attackers traffic going to a DDoS sink. So, the traffic from the DoS doesn't impact the user at the source. They go to different places.

Now, you could have a situation where both the DDoS traffic and the legitimate traffic are going to the same destination, but one of the things which makes distributed denial of service attacks so terrible is that they're distributed and there are many, many, many sources. So, instead of the person who's controlling that distributed denial of service attack being able to send it to all one place and crash the server, he kind of is forced to distribute that traffic across many different destinations, so hopefully more destinations will remain available, so that legitimate traffic can get through. Next slide.

So, if you're interested in hosting a root server and if you're running a network, what you want to have is you do want to have three to four instances nearby, and by instances, that means Anycast instances. Then, of course, you increase your [peering] connections and you might even want to host a root server instance yourself.

The second bullet, RFC 7707 technology. This is an ability to host the root zone on a local loopback. This is mainly for increasing caching. You definitely want to turn on DNSSEC validation in resolvers. This ensures you're getting unmodified IANA data. You're getting the correct data.

If you're interested, you can participate and contribute to the RSSAC caucus. Carlos, my colleague, will be talking a bit more about what the RSSAC caucus is and what it is they do. But technical advice is created there. But Carlos will talk more about that in a little bit. Next slide.

Now I'm going to hand it over to Carlos, so he can tell you about RSSAC and recent RSSAC activities.

CARLOS REYES: Thanks, Andrew. Hi, everyone. My name is Carlos Reyes and with Andrew and my colleague, Mario and Steve, we support the work of the Root Server System Advisory Committee. I think Andrew did a good job of giving you an overview of the root server system and how it functions. Within ICANN, obviously you're at an ICANN meeting – within ICANN that work and the advice to the board and the community about the root server system happens within the Root Server System Advisory Committee, the RSSAC. So, I'm going to give a brief overview about that. Keep in mind that the RSSAC does have an information session on Tuesday morning, so if you want more details about some of the items I'm covering here, please join us then. This is just a very high-level overview to complement the content that Andrew just presented.

So, what is the RSSAC? I basically just covered this, but the RSSAC is the Root Server System Advisory Committee and they advise the ICANN board and the ICANN community on matters relating to the operation, administration, security, and integrity of the root server system.

If you look at the ICANN bylaws, there are three other advisory committees charged with different mandates, but the RSSAC mandate is a very narrow scope. Andrew explained a little bit about that when he was ascribing the split between what the Root Zone Evolution Review Committee and what the RSSAC handle in terms of their mission. This is a very narrow scope at ICANN and the RSSAC tries to keep its focus on its work to that mission.

What does RSSAC do? The RSSAC, obviously it's in its name, is a committee and they produce advice. This advice goes to the board, also the community, and then anyone else who participates in DNS operations.

The root server operators, there are currently 12 organizations. They participate in RSSAC through representatives. Within that, it's important to note that RSSAC does not advise on operational matters and it also doesn't get involved in individual operations of the root server operators. In that sense, this speaks to what Andrew mentioned earlier about the diversity – for example, operations and funding and the different models that every operator has. That diversity contributes to the resiliency of the overall root server system.

So, this is RSSAC. If you look at the ICANN Board of Directors and the whole multi-stakeholder model, the RSSAC, as I mentioned, is one of the advisory committees and they also appoint a liaison to the ICANN board.

So, let's talk a little bit about how the RSSAC is organized. As I mentioned, the root server operators appoint representatives, so there are 12 representatives, and then there are also 12 alternates, one from each RSO.

The RSSAC also has liaisons. These are liaisons either from other organizations to the RSSAC or the RSSAC also sends its own liaisons to either ICANN groups or external groups as well.

Then, the RSSAC caucus. This is a relatively new group in terms of ICANN history. The caucus came into being in 2014 and this is a dedicated pool of technical experts that inform the work of the RSSAC. I'll describe the caucus in a little more detail shortly, but currently there are 100 members. So, if you're interested, I'll talk a little bit about the membership process for that.

Current RSSAC co-chairs, Brad Verd and Tripti Sinha, representing the University of Maryland and Verisign, two of the operators. These are some of the liaisons. So, I'll briefly mention some of them. The IANA functions operator – that is an inward-facing liaison, so the PTI (Public Technical Identifiers) which is the current IANA functions operator, has a liaison to RSSAC. There are a few other inward-facing liaisons. We have the root zone maintainer. Andrew described that earlier. That's currently Verisign. We have a liaison from the Internet Architecture Board, the SSAC which is one of the advisory committees here at ICANN. Then, outward liaisons, RSSAC appoints a liaison to the ICANN board, to the Nominating Committee, to the Customer Standing Committee which oversees the performance of PTI, and then also the Root Zone Evolution Review Committee.

So, the caucus. As I mentioned, there are 100 members. There is an application process where applicants submit a statement of interest. All of these are available online and they describe their qualifications and their motivation to join the caucus as well as any relevant experience that they may have.

The caucus, as I mentioned, is a group of DNS experts or root server system experts and this is really to inform the work of the RSSAC. Because the membership of RSSAC is limited to the representatives from the root server operators and the liaisons, a few years ago when RSSAC was restructuring itself, there was really an interest in expanding the expertise available to contribute to RSSAC work and the caucus was that mechanism.

So, if you're interested in applying for the caucus – as I mentioned, there is an applications process – if you send an e-mail to that address, I am the support staff for the membership committee, I will review your note and then if you'd like to move ahead, I'll send you the template for the statement of interest and then that will be added to the queue. I'm happy to talk to anyone afterward if you have any specific questions, but there is a membership committee that reviews application and then makes recommendations to the RSSAC.

Recent publications. These are the three most recent documents that the RSSAC published. As I mentioned, the RSSAC will go into more detail on Tuesday morning, but I'll just highlight maybe one of these. So, RSSAC 41, that was an advisory on organizational reviews. As part of the accountability and transparency mechanisms at ICANN, every supporting organization and advisory committee undergoes an organizational review periodically and RSSAC just concluded that process. Along the way, they learned a few things that they wanted to share with the ICANN board and the ICANN community, so the group documented its experience and published an advisory where it makes

recommendations to the ICANN board about how to improve that process.

Two current work parties in the caucus, these just launched within the last two months. The way the RSSAC currently decides how to pursue work, it basically polls the caucus. About a year ago, we launched a survey with the caucus to see what interests they had and potential work items. We received that feedback. The RSSAC then prioritized that list with feedback from the caucus and then launched two work parties. So, these are the newest ones. The statements of work are available online so you can learn about the scope of the work and the timeline and the different work party leaders and shepherds for those efforts.

So, on this slide, I'm going to stand over here to the side so I can read it, but this covers some efforts that both the RSSAC and the RSOs have taken to increase transparency. Obviously, within the RSSAC – I'll work my way from the left to the right – the RSSAC was, as I mentioned, reviewed in 2008. That was the first organizational review and that resulted in some structural changes to the RSSAC and the founding of the caucus. And as I mentioned, the caucus is this broader pool to inform the work of the RSSAC.

All minutes of meetings are published as well as reports from any RSSAC workshops. So, there's documentation about what RSSAC discussed and any decisions that were achieved throughout those meetings or workshops. We have public calendars for RSSAC and caucus. As staff, we maintain a calendar of various work party calls or RSSAC calls and those are public as well.

Obviously, public meetings. About a few months ago at the last ICANN meeting in Panama, the RSSAC agreed to, by default, have all of its work sessions open for observation. I think a few of you may have been in some of the work sessions this morning. That's relatively new for the group, but it is an ongoing effort to increase transparency and promote the work of RSSAC and ensure that the community is aware of its discussions.

Meetings with other ICANN community groups. This is some sort of standard within ICANN meetings. You'll see a lot of joint meetings. The board, for example, meets with pretty much every single community group at an ICANN meeting. The RSSAC will meet with the board Tuesday. But the RSSAC also meets with other groups, either in closed or open fora. It's depending on the agenda and the preferences of the other group. At this meeting, there's an open joint meeting with the At-Large Advisory Committee and that's one of the four advisory committees and that advisory committee is charged with voicing the interests of the Internet end users within the ICANN community.

Tutorials like this one and liaison relationships that I covered earlier. Also, operational procedures. Basically, how the RSSAC operates and how it administers its business. All of that is published and the RSSAC actually revises that document every year based on its experience with it. We're in that process now.

RSOs. Agendas for Root Ops meetings are published. Also, one of the first publications that the revitalized RSSAC published in 2014 was on measurements for the root server system and every RSO committed to

publish its data and those are all on the various RSO websites. That's part of their efforts on transparency. Obviously, RSOs participate [inaudible] in RSSAC. There's a root servers website and then every RSO has its own website as well.

Then, the root server operators do collaborate on reports, on events related to the root server system. About a year and a half ago the root server operators agreed that if any questions were provided to the RSSAC about the root server system the RSSAC could be a vehicle to communicate with the root server operators.

So, if you have any questions about RSSAC, there is a website. Also, we do have an inbox where you can send questions about RSSAC. We monitor that. Again, caucus webpage, and if you're interested in applying, there is an application process and I'm happy to explain that. That is it. I'll pause here for questions. I'm happy to take questions about the RSSAC or the caucus. Andrew can take questions about his presentation. We also have members of the RSSAC. How about if everyone just raises their hand? I don't think we [inaudible] for the room. We have a handful of members from the RSSAC around the room. Are there any caucus members as well? Very good. Great. So, there are a few resources here in the room if you'd like to ask any questions. Thank you very much. Please state your name and your affiliation.

UNIDENTIFIED MALE: [inaudible] from APNIC but I speak on behalf of myself only. I have a question about the root server. We have 13 root servers, right? So, can the number be increased to maybe 14, to 26? This is the question.

CARLOS REYES: Thanks. Who would like to take that from the RSSAC? Fred?

FRED BAKER: So, Fred Baker, ISC. Yes, the number can be changed. We can cut it in half. The question that we're usually asked is, "But my favorite company wants to add one." Yeah, that can be done."

If you think about the process by which you do the priming message, that tends to be the controlling factor as far as the number of root servers out there. Now that we're carrying IPv6 addresses and DNSSEC signatures and so on and so forth, the priming process actually takes several messages. So, if we added a root server operator, that would extend that process by some amount and so on and so forth.

Yes, it can be done. There are issues. The reason for the number 13 is historical. Once upon a time, DNS messages were 512 bytes and we were carrying IPv4 addresses only and 13 fit in that size package. That has been superseded.

That said, the next question is, "So, do we need another one?" Is there some location in the world that isn't being served well? If so, what would be the best way to solve that? Jumping in and just adding another root server operator, you want to have a good argument for doing so. So, there would be some discussion and such about that. But yeah, we in the RSSAC, one of the things that we're looking at and RSSAC 37 discusses this in terms of evolution of the model for the

RSSAC, looks at the question of how we might designate a new one and what's involved.

UNIDENTIFIED MALE: Okay, a follow-up. Is that, possible, though in [inaudible]? Is ICANN considered like [inaudible] or does any procedure apply for a new root server?

FRED BAKER: Well, that would be the designation and removal function that's discussed in RSSAC 037 and that hasn't quite been instantiated. We have had several entities that have indicated that they'd be willing to provide the services.

I think you want to go back to, "So, what's the argument for adding a new operator?" Is there something that isn't being done? What's the value?

CARLOS REYES: Thanks, Fred. Any other questions? Yes, you can use a table mic there.

RUVENI WAQANITOGA: Thank you. My name is Ruveni Waquanitoga from fellowship. I would like to just find out which root server will normally respond to the queries if answers are not resolved locally on the authoritative DNS servers?

CARLOS REYES: Andrew or anyone from RSSAC? Fred again?

FRED BAKER: So, I want to be sure I understand your question. You're asking which root server would respond to the queries?

RUVENI WAQANITOGA: [inaudible] root servers, right?

FRED BAKER: The root server that you ask is the one that's going to respond. We have actually 13 pairs of addresses, so you would pick one of those addresses and send a packet to that server and it will respond.

RUVENI WAQANITOGA: It would query the root server dynamically?

FRED BAKER: Yeah.

RUVENI WAQANITOGA: So, which one of [inaudible] would respond?

FRED BAKER: The one you send a packet to. You will pick the address that you send the packet to. Your resolver will. The server that you ask the question of will respond.

RUVENI WAQANITOGA: So, any of the root servers can respond to the query?

FRED BAKER: Well, the one you ask. You send a packet to an address. That address goes to a server. That server responds.

WES HARDAKER: First off, there's thousands of instances, as Andrew was saying earlier. The reality is that most resolving software out there today is fairly intelligent at trying to pick the best of all of those thousands of instances, and between Anycast actually helping you pick – helping the resolvers pick – and you will pick the closest ones to you. So, exactly how many is dependent on which resolver software is in use, but most of them will sort of prioritize what addresses to ask, based on the latency as it measures it.

So, as it starts out with the priming query that Fred was talking about earlier, it's going to record how far is it to each one and the result will be is it will pick the one that was actually the closest latency. It usually juggles – my colleague, actually, at ISI, has done a study to see exactly how often they go back to recheck. But, even over time, the resolver actually gets intelligence through measuring again repeatedly. So, it will adjust over time, but generally it will pick the top two to four, somewhere in there.

So, that's why, one of the things Andrew said earlier was you need two to five instances sort of near you, near your ISP and more than that really doesn't help you that much.

BRAD VERD: Brad Verd, co-chair, RSSAC. I just want to add to that two things. One is the two things that Wes just pointed out was neither of those we have any say in – we being the root operators or the root server system. Your resolver tells you, basically gives you – is going to be asking the roots. And when Wes said the nearest ones to you, that's topologically nearest to you. It doesn't necessarily physically the closest one to you. It's closest to you via the network that you're on and how it's configured. Again, out of the control of the root server system or the root operations.

CARLOS REYES: Thank you Fred, Wes, and Brad. Yes, go ahead.

UNIDENTIFIED MALE: My name is [inaudible]. I am from CITC Saudi Arabia. I believe that root servers receive millions of queries. Is it possible to have this data available for analyzing it and maybe benefit from top query to questions, queries, or existing, non-existing queries, DNSSEC-enabled queries, etc.?

WES HARDAKER: Good question. Actually, there's an organization called DNS-OARC and they do annual collections of data to all of the root servers for typically a two to three day period. In fact, there was one that just happened while the KSK was being changed for DNSSEC and I think it was four days, actually, or three and a half. So, all of that data is available to be analyzed if you are a DNS-OARC member. It's a small cost in order to do that. You can actually go in and look at all of the data to all of the root servers, so it is possible to do that and the right way to do that is through the Day in the Life of the Internet Project at DNS-OARC.

UNIDENTIFIED MALE: So, can I query about queries originating from my country and see what is the top and what is the valid, invalid responses [inaudible]?

WES HARDAKER: Yes. That data is available. Brad just remined me that also available is the RSSAC 002 data which shows you all of the instances and how many queries they're receiving and that's actually available for every day of the year. But yeah, if you become a member of DNS-OARC and actually go look at those, yes, you can get information about a lot of things.

Now, some of the root server operators anonymize their data, so there's some level of anonymization, but that's documented so you can figure out. A lot of times, those are down to only [/24s] or something which is sufficient for even doing the analysis in terms of what comes from your country. So, yes.

BRAD VERD: And if I could just add the queries … Your specific question was can I figure out what queries are coming from my country. You need to figure out what your resolvers are because you might be using a resolver that's not in your country and not know it. We will know in those metrics and the IPs that are anonymized and that [/24] goes to the resolver, the people asking us the questions. So, you need to take that into account. It's not as black and white as you might think or hope it is. There's some research that needs to happen on your side to figure out where your traffic is coming from type of thing.

CARLOS REYES: Any other questions? Another call for questions. Your mic, please.

RUVENI WAQANITOGA: Ruveni, Wquanitoga, fellowship. Just a general query. What could cause some DNS queries to fail or not to respond?

CARLOS REYES: Could you repeat that?

RUVENI WAQANITOGA: What could have caused queries to give no answers or failed?

BRAD VERD: I'm not sure where to begin with that question. Any number of things could happen to have a DNS response fail. I don't know if you're looking for a specific scenario or something, but if you're using UDP, UDP is not

[inaudible] and sometimes those packets get lost over the wire. That's just normal UDP. If there is … I'm just running through any type of scenarios. There's a number of different answers to that question. All I can say is that all the root servers are available, meaning all thousand-plus. They all answer identically and you can validate the root with the DNSSEC key at all of them. They're all the same. They all serve the exact same data. It could be any number of things that caused your failure. It could be a resolver problem. It could be a network problem local. It could be a network problem somewhere around the globe. It could be any number of things. I'm not sure how to answer that question other than they all answer the exact same data, so if you tried any one of the instances, you'd get the same exact answer.

RUVENI WAQANITOGA: I would like to add to the question because we had an instance with a [by nine] we had [four by nine] servers in our ISPs and they failed to respond to two websites just recently. Just the A record for the website. But, [inaudible] it was giving the correct answers to part of our customers.

BRAD VERD: Okay. So, if I have your scenario correct, I'm going to try to interpret what I heard, you had some customers that could reach a couple webs and a couple that couldn't. Is that fair?

RUVENI WAQANITOGA: Yes. From a [by nine] DNS.

BRAD VERD: Sure. You have [by nine] but it's a resolver, correct?

RUVENI WAQANITOGA: Correct.

BRAD VERD: Okay. It could be a configuration problem on your resolvers. It could be, like I said, a network problem. It could be a problem with the website. It has nothing to do with … Because right now when you say website, you're no longer talking to the root. Just like Andrew when he talked about how the DNS hierarchy works, if you're talking to a dot-com website or a CC website, it might be a problem with their DNS. There's so many variables involved in that question of what are the possible failures. Don't know. We'd have to sit down and walk through it. I know that really didn't answer your question. It just gives you the idea that this is not a black and white answer because it's a very complicated system to try to address that because everybody is going to have a different issue or different type of issue.

JEFF OSBORN: I'm Jeff Osborn with ISC. We developed BIND and we also run the F root. If you're getting it from one implementation of DNS and not the other one, it's certainly not the root server that's the problem. I mean, as you well know, this is all complicated but it's difficult to imagine the rot

server was the problem if you were getting it working with one implementation and not with the other. That's pretty much …

RUVENI WAQANITOGA: Sorry. I'm not saying that the root server is the problem. I was trying to find out what's the root cause. What could have been the root cause?

FRED BAKER: Fred Baker, also from ISC. We can't answer what the root cause is because we don't have all of the different variables to trace it down with. In order to discuss that, we kind of need to go back with you to get the different information about the queries and different places and whether this is a consistent problem or there was an event, which there's any number of explanations for.

If you want, you can go to the ISC.org webpage. There's a support website. Put a ticket in there and ask for some help in resolving that.

RUVENI WAQANITOGA: Thank you.

CARLOS REYES: Any other questions?

KENNETH HERMAN: Hi, good afternoon. My name is Ken Herman. I'm an independent consultant. You gave a pretty good description of the system and how

it works and the process behind it, the sort of machinery that it works with with the advisory committee. Can you speak a little bit about some of the challenges that the advisory committee might face during the course of its work? It seems to work pretty well. I know it works very well. We use it all the time. But, what are you faced from year to year? What sort of challenges? Are there scaling issues that need to be addressed? Are there technical standards that are emerging that might present some challenges to the root server operation? Is security continuing to be a problem? I'm interested to know a little bit more about the kinds of challenges that the advisory committee might be facing. Perhaps, some examples in the past, but what are you seeing in the future?

CARLOS REYES: Thank you.

BRAD VERD: First of all, clarification. This question has nothing to do with the advisory committee at all. This is an operational question. RSSAC, as stated earlier, gives advice to the board on the root server system. Your question is very … It's a great question directed strictly at operations. The answer to all of them is yes, yes, yes, and yes. Lots of challenges. Everybody is focused on security. Everybody is focused on capacity, trying to make sure we have enough for the security events. All of those concerns are … And I'm only speaking for myself as an operator for Verisign. I'm not speaking for the other operators here. But we all face very much the same types of problems or challenges and I like to use

"challenges" because they're addressable, right? Problems are hard to solve. Challenges we can work at.

Scaling, always an issue. Always trying to have enough scale to handle not just the current load because the current load really is pretty nominal, but its those crowd events, those large events that happen, that we need to be able to … We can give answers back, so they don't get lost and people can get to where they need to go. Security is on top of everybody's conversations every day nowadays. It's a huge effort to stay in front of that. It's a huge expense. So, yeah, huge amount of resources are spent on all of those topics when it comes to the operations. Don't know if that answered your question. Anything more specific?

KENNETH HERMAN: Well, from an operations point of view, I can imagine that these kinds of things happen. The advisory committee, I understand that there's separation of duties there. So, the operators I'm sure get together and talk about all those challenges.

So, I'm interested of course in learning more about how those are and what they address, but my particular question really was I guess directed more at how the advisory committee works. It meets on a regular basis and I'm interested to know what sorts of issues are emerging that require its consultation, what's coming from the board perhaps. And maybe there are other sessions that will happen this week that we can dig a little deeper, but I was hoping to get a sense as to, in

its work of the advisory committee, where does it see the challenges emerging?

BRAD VERD: So, two quick answers to that, I think. First and foremost is that's a question for the RSSAC. This is not an RSSAC meeting. I would recommend coming to one of the RSSAC meetings. I think the public session is Tuesday morning, so there will be plenty of opportunity to ask that specific question there and it might even be addressed before the meeting is over.

To answer your question directly, I would recommendation going and reading RSSAC 037 which has been published. And it is, essentially, it's a proposal of a governance model that addresses these longstanding challenges that have been out there with the root server system and it is a proposal on how we would go about doing that. It's a large document. It's 60 some-odd pages long. We took a lot of time on it. We were very careful and we think … We took all these outstanding questions and challenges and said, "Okay, let's sit down and work through it." And that was our [inaudible]. So, I would recommend reading that.

CARLOS REYES: Was there another question here? Sure. I think there's a mic there.

[JASON]: Hello, I am [inaudible], or you can call me Jason. I'm from [inaudible]. My question is while reading the RSSAC 037, I learned that in section 5.4, performance monitoring and measurement function, and that reminds me of one of the topics that last ICANN 62 mentioned which is ITHI, Identifier Technology Health Indicator. I wondered if these two are connected because they are all about performance measuring and there's a section about root server performance in that. I wonder if it will be used or adopted in RSSAC.

BRAD VERD: Yeah. We're looking around the room, see if OCTO can try to address the … ITHI, I remember the discussion. That was an effort put forward by OCTO to go and monitor some of the health of the root server system. There's been a number of documents, both published by SSAC, by RSSAC and the [CDAR] report that came out a couple of years ago that stated that more monitoring was needed. This was advice to the ICANN board on the health of the system. So, that's what that is.

RSSAC 037, we tried to … It's a governance proposal, but it covers a whole bunch of stuff. It's a full model, beginning to end, and a piece of that model was the performance monitoring and measurements function. Sorry, lots of acronyms. Certainly, those things that need to be monitored by the PMMF have yet to be defined. That's part of … We put forward the proposal. We're now waiting for the board to come back with how to implement it, and when it's implemented, one of those things that would have to be implemented are what the measurements are and what the metrics are.

So, if you're a member of the caucus – if you're not, please join. If you are, that is all upcoming work that RSSAC will be taking on [inaudible] all of those things. But they're not directly related. It's a little different, but it doesn't mean they won't blend together in the end.

CARLOS REYES: Thank you, Brad. Any other questions? Going once, going twice. Alright. We can go ahead and end this session early. I'll stick around if anyone has questions about the caucus, membership process, and then feel free to ask anyone in the room about the root server system or RSO operations. Thank you very much for coming. There's a repeat session of this tutorial if you missed part of it or if you'd like to take it again on Monday.

CATHY PETERSEN: Monday at 5:00 PM.

CARLOS REYES: Thank you, Cathy. Thanks!

[END OF TRANSCRIPTION]