SAN JUAN – How It Works: Root Server Operation
Monday, March 12, 2018 – 10:30 to 12:00 AST
ICANN61 | San Juan, Puerto Rico

UNIDENTIFIED MALE: Good morning. ICANN 61, How It Works: Root Server Operation, March 12.

CATHY PETERSEN: Good morning, everyone. Welcome to How It Works. We are running a little bit late after that wonderful opening ceremony, so please be patient. We will be starting shortly. Thank you.

UNIDENTIFIED MALE: Hey, folks. We're going to get started in about two or three minutes. We are starting 15 minutes late because the opening ceremony ran a bit over. But if you could be ready to start in two or three minutes, that would be great.

CATHY PETERSEN: Good morning again, everyone. Welcome to How It Works. This session, we will be talking about Root Server Operations. Thank you again for your patience. Andrew McConachie is our first presenter. Andrew?

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

ANDREW MCCONACHIE: Thank you. Hi, my name is Andrew McConachie. I work for ICANN Policy Support, supporting the RSSAC. I'm going to be talking about the root server system.

First off, a bit of an outline. We have four sections today: overview of the domain name system, the root server system today and its features. Then I'm going to hand it over to my colleague Steve Sheng, and he's going to give an explanation of Anycast and then talk about the RSSAC and some recent RSSAC activities.

After that, we'll have a question and answer period where some of the root server operators that are in the room come up here on stage to take your questions. So please save your questions until the end.

Let's start off with an overview of the domain name system and the root servers. A bit of a recap: what are IP addresses and how do they function as identifiers on the Internet? IP addresses are the fundamental identifier on the Internet, and all hosts connected to the Internet do need to have IP addresses. Whether that's IPv4 or IPv6 or you're operating through a NAT, you still need an IP address. IP addresses are a numerical label. They're not really all that human friendly. They're just numbers.

Why do we need DNS? Well, the original problem was, as I mentioned on the previous slide, IP addresses are hard to remember and they change a bunch. So the original problem with DNS was just having some human memorable names that we could map to IP addresses so we didn't have to remember IP addresses.

Those problems remain, but there are some more modern problems as well such as IP addresses may be shared and multiple IP addresses can map to a single service. So we have this modern problem of both the many-to-one and the one-to-many being added upon the original problem of just IP addresses being really hard to remember.

Now the domain name system is hierarchical. As you can see in the diagram, at the top we have a root. Beneath that, we have what are called top-level domains or TLDs. Some examples include .edu, .mil, .uk. Then beneath that we have what some people refer to as the second level and then the third level and on and on. These are named IP address mappings. That's what we're most family with, but there are other mappings as well such as mx records for mail servers, reverse records sometimes called PTR records which map from IP addresses back to names.

This slide is pretty complicated. I'll spend some time walking through it. This slide is meant to show the DNS resolution process, how a user experiences the DNS, what a user goes through, the various steps, the flow of what it means to interact with the DNS so that a user can resolve a domain name to an IP address and eventually get to a website.

We have a user over here on the right. They really want to go to www.example.com as a web server. The first thing they do is open up their web browser. That triggers a DNS request, and that DNS request is going to go to what's called a recursive name server. Assuming the recursive name server doesn't have anything in its cache and let's pretend for the sake of this demonstration that someone just turned this recursive name server on, it's cache is empty, it doesn't know anything. What does it do?

It just got a query for www.example.com. It has a bunch of work to do before it can back to the user with an answer. The first thing it does is it goes to the root and it says, "Here's www.example.com. Where is it?" The root says, "I don't know where that whole thing is, but I do know where .com is." So it returns to the recursive name server with the address of the .com name servers.

Then the recursive name server goes out to the .com name servers and says, "Where's www.example.com?" The .com name servers say, "I don't know where that whole thing is, but I do know where the example.com name server is. Here it is."

Then the recursive name server goes to the example.com name servers and says, "Where is www.example.com?" Finally, he gets the response he's looking for and the recursive name server is able to respond to the user with the address of www.example.com.

That's how a user will go through the entire process of resolving this domain name to an IP address and finally be able to visit the web page.

There's another thing I haven't talked about in this slide which is the security aspect, the DNSSEC or sometimes called DNS security aspect, which is that with each one of these questions between the recursive name server and each one of these authoritative name servers – the root name server, the .com name server, and the example.com name server – these answers that go back to the recursive name server from these authoritative servers are signed and the recursive name server is able to validate that they're the correct answer. That this answer hasn't been tampered with. That it hasn't been given by the wrong authoritative name server. This is the correct answer, and it's able to do this through DNSSEC.

That's the domain name system resolution process. Like we saw in the previous slide, the root servers only know who needs to be asked next. They only have the addresses of the TLD name servers like .com, .net, and .org. However, they typically don't ask that often.

In the previous example I gave, we had this hypothetical situation where the recursive name server had just been turned on and it didn't have anything in its cache. Well, this is pretty rare. Recursive name servers have pretty extensive caching, and the vast majority of queries that go to recursive name servers are answered out of the cache. That means that there's a lot fewer queries to the root than what you might originally think.

Some modern refinements to DNS. I already talked about DNSSEC also sometimes called DNS security or security extensions. The point of DNSSEC is to sign responses that go to the recursive name servers so that the recursive name server can validate them. By validation, I mean he can ensure that it's the correct response because it has been signed by a key through cryptography, so he's able to ensure that the response is correct.

There's also been privacy enhancements, and these are still being heavily worked on at the Internet Engineering Task Force. Something like DNS over transport layer security, which will secure the transmission of the query over the wire and ensure that the prying eyes can't view it. Those are still very much in active development.

Another modern refinement to DNS is Anycast. Anycast is used heavily by the root server operators. Anycast basically does two important things. It allows multiple servers to share a single IP address, and it protects against DDoS attacks. My colleague Steve Sheng will be talking later in more depth on Anycast and how it does that.

The root zone versus the root servers. The root zone is the data that the root servers serve. You can think of this like the root zone is the starting point. It's the list of TLDs and name servers. It's the top of the hierarchy or tree. It's managed by ICANN per community policy. It's compiled and distributed by the root zone maintainer to all root server operators. Again, it's the database content of the root server operators. It's the data that the root servers serve.

On the other hand, the root servers respond with data from the root zone. Currently, there are 13 identities and over 900 instances at many different physical locations worldwide. The root servers are purely a technical role. They serve the root zone data. Each of these Anycast clouds that the root server operators run are their own responsibility.

Delving a little bit more into what a root server operator is, there are 12 different professional engineering groups that are focused on the reliability and stability of the service, accessibility for all Internet users. They're professional and they cooperate with one another while also acting independently. They're a diverse group of organizations. By that I mean they're technically, organizationally, geographically diverse.

However, operators not involved in policy making and they're not involved in data modifications. They just serve the root zone data. They are involved in careful operation of the service, of serving that data and evaluating and deploying new technical modifications – so new standards that might come out of the Internet Engineering Task Force – and making sure that the service remains stable, robust, and reachable by all users across the Internet.

That was a bit of a background on the DNS, a little bit technical but probably not too technical. Now we're going to go into the root server system today and some of its features.

The growth of the root server system. This slide shows a bit of the history of the numbers of root servers, root server identities there have been over the years since the 1980s. You can see that it has progressed.

Now since 1998, we have 13 different identities. These changes have mainly been responding to technical demands as well as scaling issues. Nowadays, scaling issues are really solved by

Anycast. Anycast is just a wonderful tool in the tool chest of root server operators to deal with scaling issues.

Root servers today are all operating IPv6 and IPv4, so there are 13 IPv4 and IPv6 address pairs. Again, there are over 900 individual instances.

These are some of the foundational principles of the root server system. There are five of them. It's important that that root server system provide a stable, reliable, and resilient platform for the DNS; that it operates for the common good of all the Internet; that the IANA is the source of the DNS root data – that's the root zone data; and that architectural changes be made based on the results of technical evaluation and demonstrated technical need; and that technical operation and expectations of the DNS are defined by the Internet Engineering Task Force.

If you're interested in a bit more of the history of the root server system, you can download and read RSSAC023, the History of the Root Server System, from the RSSAC website.

These are the root server operators today. We see that there are 13 identities. The host names of them are listed on the left. You see in the middle column the IP addresses. There's both IPv4 and IPv6 for all of them. Each one of these IPv4 and IPv6 addresses, well at least all the IPv4 addresses result in an Anycast cloud. So behind those IP addresses there are many,

many servers – over 900 at this point, but it's growing all the time. At the last ICANN when I gave this presentation, I was saying there are over 800 instances and now I'm saying there are over 900 instances. So it's constantly growing.

Here's a bit of a view of the root servers today. This is from the website root-servers.org. It's just an overview of where the root servers are. It's not particularly precise. It doesn't tell you that there are exactly, for instance, seven root server instances in Madagascar. It's a graphic. It's kind of interesting. You can drill down further if you go to the website. You can even look at the actual cities that each instance is in under each one of the operators. This is a very broad and general overview but if you go to the website, you can really drill down and get some interesting information out of that.

This is the root zone management structure. This is how the root zone data, the root zone, gets to the root servers. Let's say that you're a TLD operator, you need to make a change to the root zone. Maybe your NS records change. Maybe your glue changes. You need to change some of the information associated with some of the records for your TLD.

So you'll go to IANA and make that change, and then that change will be passed on to the root zone maintainer who is currently Verisign. Then I think two times a day they'll distribute that change, they'll distribute the entire root zone to the root server operators. Then the root server operators are responsible for getting that out to their whole Anycast cloud and then serving or responding to the queries that come in from all of the recursive resolvers.

Some of the features of root server operators: there's a diversity of organizational structure, their operational history is sometimes different, they use different hardware and software. They'll use different hardware platforms as well as different software platforms. That helps in terms of security because there's strong correlation between better security and increased diversity. They've also got different kinds of funding models. They're different kinds of organizations and they get their funding in different ways.

They do, however, have some shared best practices: strong physical security, they overprovision their capacity to deal with DDoS attacks and as well to deal with spikes in traffic, and they all have a professional and trusted staff.

They cooperate through various industry meetings in the community. ICANN is one of them but also the Internet Engineering Task Force, NOGs such as NANOG or RIPE, DNS-OARC which is an operational and research group. They also use

Internet based collaboration tools, and they're transparent in their operations.

They also coordinate to prepare for emergencies in order to protect the infrastructure in case of catastrophic emergencies or other types of emergencies. They have periodic activities to support emergency response. I can't read that last bullet because it's cut off.

Responses to an evolving Internet. As the Internet evolves, new requirements are placed on the DNS system. Over time, the root server operators have adopted IPv6, Anycast, DNSSEC. IDNs are also mentioned here because there are a lot of IDNs in the root zone. The important thing is to increase the robustness, the responsiveness, and the resilience. Again, there are over 900 Anycast instances deployed today.

Some of the myths that people might have, some of the misconceptions that folks might have about the root server system. Myth one, root servers control where Internet traffic goes. That's not entirely true. In fact, it's not really true at all because it's a myth. Routers actually control where the Internet traffic goes. I think maybe this myth might have come about because DNS maps names to IP addresses, but ultimately it's routers based on IP addresses that control where the packets go.

Another myth is that most DNS queries are handled by a root server. As we saw in the example, that might be true if the cache of recursive DNS servers is empty, but that's very rarely true. So most DNS servers are not handled by a root server. Most of them are handled out of the cache of the recursive.

Administration of the root zone and service provisioning are the same thing is another myth. That's not true. The diagram I think I showed previously about the division of responsibilities and how a change makes its way through to the root servers, there are different parties involved there.

Another myth is that the root server identities have special meaning. They really don't. Or that there are only 13 root servers. There are over 900.

Another myth would be that the root server operators conduct operations independently. While they are independent organizations, there's a lot of coordination and cooperation that goes on to ensure the stable service of the root server system as a whole.

The final myth that the root server operators only receive the TLD portion of a query, well, that's not actually true. They receive the entire query. This is just the way DNS works. There is some work in the Internet Engineering Task Force to change this. If you're interested, the keyword there is QNAME minimization and you can go read about this work where the root servers might only receive the top portion of the query.

Now I'm going to hand it over to my colleague Steve Sheng who will take you through the last two sections starting with Anycast.

STEVE SHENG: Thank you, Andrew. My name is Steve Sheng. I'm also a policy staff supporting the RSSAC. I'll give an explanation of Anycast and also RSSAC and its activities.

Anycast is a routing and addressing term. There are two terms here: Unicast and Anycast. There are important differences. In Unicast, packets or datagrams from sources all go to the same destination and a single instance serves all the sources. So in the event of a denial of service attack, all the attack traffic goes to that single instance. That's Unicast.

Whereas, in Anycast the multiple instances serve the same data to all sources. These multiple instances have the same IP address, and the intermediate routing policies determine the destination based on the source. This means the source gets the data faster, gets to the closer destination, and the DDoS attack traffic is sent to the closest instance.

Let me illustrate that with a diagram. Here the illustration of Unicast, you see a source and a destination that is identified.

The destination is single instance, and the traffic takes the shortest route to the single destination.

Here in Anycast, you see the three destinations in blue. These destinations all advertise the same IP address and the routing policies determine the closest, where that source to destination. This means the path from the source to destination is shortened and data is delivered more quickly.

How does this help the denial of service attacks? In a denial of service attack, an attacker attacks the destination. But because it’s Anycasted, the traffic only goes to the closest link. So therefore, maybe one of the destination links is overwhelmed but the other destination still serves traffic.

One of the questions we get from these tutorial sessions is the root server system and your networks. Some of you are network operators, some may operate recursive instances. If you’re a network operator, you want to have three or four nearby instances. This gets the instances closer to you and in some instances it will reduce the roundtrip time.

I think in addition to that, you also want to increase your peering connections and peering arrangements. Sometimes you see you may have a root instance near you, but the traffic still travels around the world to get to you. That’s because of the peering connections and peering arrangements. So that's also an important factor.

If you are a recursive resolver operator, to increase caching you might consider deploying RFC7706 technology. This is running a copy of the root zone on a loopback address. The benefit of that is sometimes it reduces the privacy risk from the recursive resolver to the root server from the prying eye.

It's obviously important to turn on DNSSEC validation in resolvers. That ensures that you're getting the unmodified IANA data all the way, as Andrew mentioned, through the correct data.

And finally, I think this is because we're at ICANN, we invite technical experts and others to participate and contribute to the RSSAC Caucus. That's where the technical advice of RSSAC is generated and created.

With that, let me give a quick update or an overview of RSSAC and recent RSSAC activities. RSSAC stands for the Root Server System Advisory Communicate. It's chartered to advise the ICANN and community and board on matters relating to the operation, administration, security, and integrity of the Internet's root server system. Please note, this is a very narrow scope for this advisory committee.

One important distinction that is often conflated, especially within ICANN, is RSSAC is a committee that produces advice, primarily to the board, but also to other ICANN bodies and organizations that involve themselves in the overall DNS business.

However, the root server operators are represented inside RSSAC. But it's very important to note RSSAC does not involve itself in operational matters. So I think this is a very important distinction not to conflate those two entities.

Within the overall ICANN governance structure is one of the four advisory committees, and it sits in the ICANN ecosystem there.

Within the organization, the RSSAC is composed of the appointed representatives or the root server operators, and each one can appoint an alternate to the RSSAC. It also has liaisons to the root zone management partners and key technical organizations.

The RSSAC Caucus that I mentioned earlier is a body of volunteer subject matter experts. Their members are confirmed by the RSSAC based on a statement of interest.

The current RSSAC chairs are Brad from Verisign and Tripti from University of Maryland. Brad and Tripti, are you in the room? You can raise your hand. That's Brad. I know Tripti, but she may have just stepped out momentarily.

Within the RSSAC it has several liaisons. One is the liaison from the IANA functions operator, the root zone maintainer. Andrew showed that diagram, the root zone management flow, so IANA and the root zone maintainer. They are the two critical entities here.

RSSAC also has liaisons from the Internet Architecture Board. That provides architectural guidance to the ISOC and IETF on Internet architectural matters.

Within ICANN, RSSAC has liaisons to the Security and Stability Advisory Committee, the ICANN board, nominating committee, the Customer Standing Committee which is the committee set up to look at the performance of the IANA function that's currently performed by PTI.

And finally the Root Zone Evolution Review Committee, this is a committee set up as part of the IANA transition to look at architectural issues for the evolution of the root zone.

The RSSAC Caucus currently has 88 technical experts. Their statements of interest are public. In any RSSAC publication for any of the Caucus members that contribute or lead that work, they are acknowledged at the end of each report. So there's public credit for the individual work.

They bring diverse expertise to the publications. And the Caucus has transparency of who does the work. The mailing list is open, so you can view the archives. There's also a framework for getting things done.

If you're interested to join the RSSAC Caucus, the correct e-mail to apply rssac-membership@icann.org.

Here are some of the recent RSSAC publications. RSSAC has a publication series. That's the numbering. Currently, they are on 31. The last [free] publication is RSSAC029, which describes the outcome of the their 2017 workshop in October. RSSAC030 is a statement on entries in DNS root sources. And RSSAC031 is a response to the GNSO PDP subsequent procedures. This is about subsequent procedure for creating new TLDs. The RSSAC response touches on a topic of scaling the root.

The RSSAC will have a public session this week. Please attend to hear more in-depth details of these publications.

The current work, there are two: Harmonization of Anonymization Procedures for Data Collecting. The RSSAC published RSSAC002 and the root server operators implement that to publish statistics about the root servers and the root server system. There's [an effort underway] to look at anonymization procedure for some of those data. The other work party is the Packet Sizes and DNS.

Since the restructure of RSSAC in 2013, transparency is one important goal is trying to improve, and they have made lots of progress on that by establishing the Caucus, by publishing the minutes and workshop reports so that the ICANN community can understand the current status of the work and the reports, the workshops.

The is a public RSSAC and Caucus calendar with all the various work party meetings. At each ICANN meeting, the RSSAC holds public sessions. We have the tutorials, and the liaison relationships make sure the information flows to the key organizations.

Finally, the RSSAC has [codified] the operational procedures that define how RSSAC operates. That's also on the website. I think we're on the third revision.

The root server operators also take steps to improve transparency. The root-ops agendas are published for the IETF meetings. Every operator publishes RSSAC002 statistics. They are participating in RSSAC. There is a public web page, a single web page, and from that web page you can go to individual operator web pages. They collaborate on reports on major events. For example, the DDoS attack events last year. And the RSSAC acts as a gateway to channel these questions to root operators, and they will answer those questions.

For more information, here is a link to the RSSAC web page. Any general questions you can send to this e-mail address. The link to the Caucus and membership is listed here.

Finally, I want to bring to your attention that the RSSAC has recently published on its website a frequently asked questions and answers. I think it's a list of 25 questions frequently asked. Some of those are generated from these sessions. So those are very helpful in understanding the RSSAC.

With that, I think we come to the end of the presentation. We have some RSSAC members here. I would like to invite them to come up here to the stage to introduce themselves, and also I'll moderate a question and answer session. So with that, may I invite the RSSAC members to come to the stage.

If you have any questions, please raise your hand and I will identify you. So let's do that. May I ask first the RSSAC members to introduce yourselves starting from Fred?

FRED BAKER: Fred Baker, ISC.

JOHN CRAIN: John Crain with ICANN.

KAVEH RANIBAR: Kaveh Ranibar, RIPE NCC.

BRAD VERD: Brad Verd with Verisign.

LARS-JOHAN LIMAN: Lars-Johan Liman, Netnod.

STEVE SHENG: Thank you. We'll begin with an online question. Cathy?

CATHY PETERSEN: We have an online question from Jose de la Cruz. The question is, "Are there plans to expand the entities to more than 13?"

STEVE SHENG: "Are there plans to expand the entities to more than 13?" Who would like to take that question?

KAVEH RANIBAR: I can start. First of all, technically it should be possible to expand. That's my personal take. But I think that the real question is why should we expand the number of identifiers? Because those letters, basically, are identifiers. But from a technology standpoint if you look at the current situation, adding nodes or adding letters will not have a significant or a visible technical difference. So the first question is what problem are you trying to solve with adding new identifiers? That was my take.

BRAD VERD: I think I will add that this is a recurring question. We get this quite often, and I think the answer is RSSAC is looking at not only adding but maybe removing some. Maybe 13 isn't the right number. Maybe it's less; maybe it's more. We don't have that answer. It is one of the things that is on our work list to address. But as Kaveh said, our goal would be looking to solve a technical issue with the answer. Thanks.

STEVE SHENG: Thank you, Kaveh and Brad. With that, I open to questions on the floor. The gentleman in the front.

CATHY PETERSEN: If you could please – just a reminder – please state your name and affiliation if you have any. Thank you.

ABDULKARIM OLOYEDE: Thank you very much. My name is Abdulkarim from Nigeria. I'm a first-time ICANN Fellow. My question is about the root servers because each of the 13 root servers are probably duplicated somewhere around the world with the same IP address. So if there's a problem with one of the duplications, how can you differentiate since they have the same IP address? So how is that going to be located? Thank you.

FRED BAKER: This is a question really about how Anycast works, which there was some discussion in the previous presentations. The fundamental thing there is routing. You have each of these servers not only performing the service of responding to requests and here is the translation, whatever it is, but also engaged with BGP with the ISPs or the IXPs that they're associated with and announcing their address.

So when a request goes from somewhere to the address, routing will direct it to the instance of the server that's topologically closest. Now should one of the servers go down, should routing be lost, should bad things happen in some way, then that address from that place is withdrawn from BGP. BGP is no longer routing toward that, and there are going to be other instances around that are offering the same address. So routing will now take the packet to some other place. That's just standardly the way routing works.

In the worst case, let's imagine – and I don't know why this would happen – but imagine that the address was no longer available in routing. It just didn't exist. One of the reasons that we have 13 root server operators is so that the application requesting that, the resolver in somebody's computer, can pick one of the other addresses and go ask somebody else. So there are two levels of backup there.

STEVE SHENG: Thank you. Liman?

LARS-JOHAN LIMAN: I'd like to fill in that when you use Anycast, every server has two IP addresses. One of the IP addresses is the same for all computers, and that's the one that's used for the DNS traffic. In addition, every server has a unique separate address. That's used by the operators to reach from behind, so to speak, to do service and administration.

STEVE SHENG: Thank you. John?

JOHN CRAIN: Thank you. I think you're also questioning how you can tell which one you're talking to. There's actually a DNS query for a TXT query. You have to type CHAOS HOSTNAME.BIND and there's a couple of other variants where every instance actually has a name that you can query in the DNS which will tell you which one it is. I can show you that on the computer later if you want to know that.

STEVE SHENG: Thank you for that question and the comprehensive answer. Gentleman in the back?

UNIDENTIFIED MALE: [inaudible] from India. [inaudible] security [inaudible] DNSSEC. Can you just tell us in which countries that [DNSSEC] has been completely implemented and any problems that have occurred during implementation.

STEVE SHENG: Question on DNSSEC deployment. Anyone? I think there's a DNSSEC workshop on Wednesday. At the beginning of that workshop, they also show the deployment numbers across the globe. That would be one session where you can find those numbers.

BRAD VERD: That's a little bit out of scope for RSSAC. If there's a different way to rephrase your question and have it attributable to the root, maybe we can try to address it.

KAVEH RANIBAR: Basically just to clarify what we publish as root server operators, we get a signed zone file. That's the root zone which is signed. So it has passed, RSSAC or root server operators job basically starts after that when there is already a signed root zone file and we basically distribute that file. So from our point of view, we just distribute a signed root zone. And we make sure the integrity of the file we get is preserved, and we make sure that we keep it when we distribute the file or the content.

STEVE SHENG: Thank you for that. Any other questions? Gentleman over there?

TARAU BAUIA: Tarau Bauia from Kiribati. I have one question. When we deploy the DNSSEC, will there be a problem with the subdomains [or say] under the .com that have no keys or are not yet changed to the DNSSEC? Will that be a problem?

STEVE SHENG: Again, this is also a DNSSEC that maybe better fitted for Wednesday's workshop. That's my take. So I would invite you to that workshop. Please come to me, and I'll let you know the details of that workshop. Now I move to an online question, and then you're next.

CATHY PETERSEN: I have another question from Jose de la Cruz. The question is, "Who can participate in RSSAC?"

STEVE SHENG: Thank you.

BRAD VERD: The RSSAC has a Caucus which currently is over 80 members of subject matter experts. All of our work parties come from the Caucus, are sponsored by the Caucus. There is a – I think it's up here on the screen – the Caucus membership if you're interested you can send to the e-mail address. We have a membership committee that reviews the application. You have to give an SOI, basically a statement of interest. And then you're part of the Caucus and you are part of the solution.

STEVE SHENG: Thank you, Brad, and thank you, Jose, for that question. This gentleman in the front?

UNIDENTIFIED MALE: Steve? Could I?

STEVE SHENG: Oh, go ahead.

UNIDENTIFIED MALE: Just to add to that, just to reiterate, most of the actual technical work of RSSAC is done via the Caucus. So if you're a member of the Caucus, basically you are doing the actual work. So RSSAC, as you saw in the slides, the 12 organizations, the 13 operators, basically do most of the administration work.

When we receive a question or when advice is needed, we form a work party in the Caucus. And all of us – the RSSAC committee members – are also part of the RSSAC Caucus. So if we want to also be part of the solution, we will also join the work party. But for each question or advice that is received, we form a work party and basically the work is done within the Caucus so you will be part of that RSSAC.

BRAD VERD: And I'll add that the work is also attributable to the people in the Caucus who do the work. So it isn't that the Caucus does the work, writes the papers, and other people get credit for it. If you are a contributor, you get full attribution.

STEVE SHENG: Thank you. Go ahead.

ABDULKARIM OLOYEDE: Okay, thank you. I want to find out how often does RSSAC meet, and how often does the Caucus meet?

BRAD VERD: The RSSAC meets monthly. We have monthly calls where there are minutes taken and we cover issues. Those calls are public and can be – no?

UNIDENTIFIED MALE: The minutes are.

BRAD VERD: The minutes are public. I'm sorry. The minutes are public. Also, RSSAC meets here at the ICANN meetings, and then also RSSAC for the last couple years we've been doing two workshops a year as we have been working through some evolution work that you can hear about if you come to our RSSAC public meeting here.

As far as the Caucus goes, the Caucus works online. Work parties happen all the time. Work within the work parties is ongoing, depending on the work party itself. There might be weekly calls or biweekly. It just depends on the workload that's actually happening.

The Caucus meetings themselves are here at the AGM. And this is, by the way, determined by the Caucus. So the Caucus asked for when they meet, and they came to the conclusion that we meet at the AGM meeting of ICANN, and we meet at every other IETF meeting. Which I think an easier way to say that is it's the even-numbered IETFs is when the Caucus meets at those meetings.

STEVE SHENG: Thank you. And the next meeting for Caucus is at the IETF 102 in Montreal. Thank you. Any other questions? Gentleman on the left?

BONNIE MTENGWA: Okay, thank you. My name is Bonnie Mtengwa from the Telecoms Regulatory in Zimbabwe. I've got a question. We are keen to have one of the root servers in Zimbabwe, but does ICANN assist interested ones to have a root server or maybe it's dependent on the country negotiating with the root server operators alone or maybe there is assistance that is given by ICANN or maybe there are some requirements that must be met first?

STEVE SHENG: Thank you for that question on interest in hosting a root server instance. Liman?

LARS-JOHAN LIMAN: Several, if not all, but at least most root server operators have Anycast clouds and are willing to engage in discussions regarding where to place them and where to host them. It's not the discussion between a country and a root server operator. It's between a specific host, hosting organization. Many times it's an exchange point for Internet traffic or it's a large Internet service provider or something.

You are right that there are certain requirements that need to be met, mostly technical and financial requirements. We are working on compiling a list of contact points, but I would say come up and talk to any root server operator and we will try to explain we see the relationship and what the requirements are from us and others will do the same from their side. There are definitely openings for having a root [name] server in your environment if we can find a way to meet the requirements because it has to work, so there are requirements, yes.

STEVE SHENG: Thank you, Liman. Any other questions? Gentleman in the front again.

ABDULKARIM OLOYEDE: I've just been thinking that, yes, the DNS the root server is an important part of the Internet and the working of RSSAC and the root server operators seems to be very open. We talk about bad guys trying to attack the root servers. How do you filter that out? Where this person has an ulterior motive since anybody can join, anybody can be part of the meeting, anybody can contribute. So how do you filter that out? Thank you.

STEVE SHENG: Question on how to filter out bad guys [from the work with RSSAC].

KAVEH RANIBAR: That's a good question, a complex question [also] because there are multiple facets to that. But one of the things I think, and I'm speaking only for RIPE NCC but I think most root server operators if not all share that sentiment, that we cannot guarantee security of roots via obscurity. So we are very open, not only on how we work, but DNS by design is open. You can get a lot of information about instances, where they're hosted and all of that. In many cases we publish them but even if we don't, via DNS with some basic knowledge it's easy to figure it out.

So the system is open. Via the amount of capacity that we have from a technological standpoint, we basically try to guarantee to be able to answer every single query. And, yes, as a root server operator, we also a lot of let's call them illegitimate queries for no reason or for reason of attack. But in total, we have enough capacity to be able to [stand out] and still serve the right queries and the good queries.

STEVE SHENG: John?

JOHN CRAIN: These are professionally run networks, so all of the operators have skilled engineers and security people and we take the integrity of our systems very seriously. That’s why when you host an instance, for example, there are requirements and some of those are about who can access the machines and how, etc. So security is something we take very seriously. But as Kaveh just said, the DNS by design and to meet it’s purpose is very open. I think that’s just the nature of the protocol, if you like.

STEVE SHENG: Thank you. Brad?

BRAD VERD: I think I’ll I think more specifically to your question. I think you were with RSSAC the Caucus being open and people being able to join, what is there there to prevent a bad guy from joining and trying to do something malicious. Is that kind of where you were going with the question?

I think, yes, that is a risk. We want to be open. We want to be transparent, and we want people’s multiple different points of view so that we come up with the best possible solution for any technical problem that’s put in front of us. As a co-chair, I hope that there are enough checks and balances with the people in there with good intent that we would be able to identify people with malicious intent and try to work with them to figure out what's going on. But as of yet, I'm not aware of anything quite like that happening, but it is a risk.

STEVE SHENG: Thank you. Next question? Gentleman in the back?

UNIDENTIFIED MALE: Hello. My name is [inaudible] and I'm from India. I go back to a point which you made earlier in your presentation that the traffic route is not determined by the root server but by the routers. I would like to ask you, especially with regard to the fact which you pointed out the RSSAC002 requirements which are made upon the root server operators to publish their root server statistics, so now if I want to determine what is the total Internet traffic originating out of a particular location or out of a particular country, then how can I extrapolate that or how can I go about measuring that from some of the statistics which are available readily online which are available open source? Thank you.

STEVE SHENG: [Thanks for the question].

BRAD VERD: I think the short answer is you can't. Root DNS traffic should not be used as a measurement of total Internet traffic.

UNIDENTIFIED MALE: No, sorry to interrupt you, but what I want to also ask you is how can I do a [inaudible] kind of estimation or some kind of approximation, you know a fair estimate, how can I use the DNS traffic as a measure of arriving at [inaudible]?

UNIDENTIFIED MALE: In general and what Brad is saying, basically to be able even to have a useful estimate, you cannot use DNS for that. DNS is not a platform for that, but there are other measurement techniques. For example, if you want, look up Google MLAP. Based on [torrent] traffic that they see, they try to estimate the [rest of] traffic for a country or a region.

So there are other [products], but DNS is really not the right platform for that. The main reason for that, first of all the content doesn't go through DNS. But the second part is, what you get at DNS at any level, especially at root, a lot of that is cached by the resolvers and efficiency of that cache is not visible to us. So it's basically impossible even to have useful estimates even for the DNS traffic based on that.

UNIDENTIFIED MALE: Let me add one thing to that. When you go to a root system and the RSSAC002 statistics say I got so many IPv6 requests, so many UDP, so many whatever, they're asking about requests to the root. That's people trying to find the .coms and the .nets and whatever of the world. They're not particularly looking at particular locations in any sense or even individual companies. They're looking at registries. So it is just the wrong data.

STEVE SHENG: Thanks. Any other questions? Do we have any questions online?

CATHY PETERSEN: No questions online.

ABDULKARIM OLOYEDE: In terms of RSSAC and organizing capacity building, do you guys do anything like that? Capacity building for probably developing countries or people who are interested? Because a lot of times if for example I'm interested, I might never work on a root server in my life, or I might be interested in what you guys and sometimes it might be too technical for me because I don't do it in my daily life but I want to know more. Thank you.

KAVEH RANIBAR: I will give mine because I don't know if I understood the question properly. First of all, let me thank you as a first-time attendee and a Fellow. So thank you very much for being so much interactive. It's very much encouraged.

About the capacity building, actually each individual root operator has or might have their own plan. For example, I'm talking for RIPE NCC. We are an RIR, a regional Internet registry, for Europe, Middle East, and Central Asia. But what we do, not only within our region but also for the rest of the world including Africa and the AP region, we are working with other RIRs. For example, [inaudible] we have an MoU with AfriNIC for Africa or APNIC for Asia Pacific, and specifically let me use the example from Africa.

What we do in Africa is AfriNIC has formed an alliance with ISOC Africa, and they are actually bringing funding and talk to operators and talk to interested parties and we have actually hosted a few nodes through that fund and that capacity which was built via ISOC Africa and [our RIR].

So this is the methodology that we at RIPE NCC have selected. Others have different methods and different ways of reaching out to regions. So you have to check each individual root operator.

And just to mention because there was also a question about how to get root instances, on the website root-servers.org, you basically have a list of each operator and there you also have their website for the root service. So you can check RIPE NCC's website for the root service, Verisign's website for the root service. And there you can find all the information. For example, our agreements with other RIRs is listed there and how you can get there either directly or indirectly through [your] regional Internet registry is listed there and mentioned.

UNIDENTIFIED MALE: If I may add, a lot of these questions are operational in nature, and you happen to have a number of root operators up here also. But I'd again refer you to the mandate of RSSAC which is we provide advice about the root server system to the board and the community. Where a lot of these questions are not really directed toward RSSAC, and I think the group up here is happy to answer them and we want to be as transparent as possible, but we want to continue to – there is a delineation between root operators and RSSAC. So I just want to draw that out.

STEVE SHENG: Thank you for that. Any other questions? Going once. Oh, Liman.

LARS-JOHAN LIMAN: Just a final remark. If you come up with questions after the session, please come and talk to at least me. I guess it goes for the rest of us. We are here for a reason. We want to talk to you, and I'm happy to give the answers I'm able to give anywhere.

UNIDENTIFIED MALE: And again, I would make a plug for the FAQ that was recently added to the RSSAC web page. Even though, as I stated earlier, a lot of the questions we get are operational in nature, we've gone out of our way to capture all those questions that have happened through multiples of these different presentations as well as just questions that we normally get. That is a constantly evolving FAQ so if there's something that isn't there that you have a question, I'm sure somebody else does too so please share it with us and we'll add it to our FAQ and make that richer. Thank you.

STEVE SHENG: With that, let me show you on the screen on RSSAC website, we have Meetings, Caucus, Publications, and FAQ. When you click on the link, there's a lot more information about meeting minutes, Caucus membership, all the RSSAC publications, and the frequently asked questions.

The root-servers.org website that Kaveh mentioned is a gateway to individual root servers. If you have specific [optional] questions, there's contact information there. This is also where in the presentation we took the Anycast map. You can drill down to more information in that website as well.

With that, if there are no further questions, I thank you for your participation and for the RSSAC members answering your questions. Thank you. This session is adjourned.

CATHY PETERSEN: Thank you, everyone.

[END OF TRANSCRIPTION]