MONTREAL – How it Works: Root Server Operations
Sunday, November 3, 2019 – 15:15 to 16:45 EDT
ICANN66 | Montréal, Canada

STEVE CONTE: We're going to get started in just a few minutes. We are in a way bigger room than we need to be. If you guys want to come up closer, that's great, we could actually see you then. The way we're going to do this is, the RSSAC is going to make a presentation and then we'll have blocks for questions and we will come out to you for questions. I've already made my steps for the day, so the closer you are the better it'll be for me too, and it's all about me. We'll get started in just a minute.

ANDREW MCCONACHIE: Hello, my name is Andrew McConachie, I work for ICANN supporting the Root Server System Advisory Committee and I'll be talking today on the tutorial for the Root Server System, I guess I can see right there. I like to pace, so I'm not going to sit down, I apologize to anyone remote, the camera has to follow me or something. I promise, the presentation is a lot more interesting than any pictures of me, so we'll just go right ahead.

I'm going to be doing this in coordination with a colleague of mine, Ozan Sahin, who will be doing the second have. I'll first go over the overview of the DNS, which is probably going to be a bit of a review for a lot of the people in this room because I understand you already had

DNS101 but we'll go through that. Then I'm going to go over an explanation of Anycast and how Anycast is different than Unicast, and why it matters for the Root Server System. Then I'll go into the Root Server System today, instances and a bit of its history. After that, Ozan will take over and talk about the RSSAC, the RSSAC Caucus and the ongoing work with the Root Server System Evolution. Let's get started.

Overview of the DNS. I'm guessing most of the people in this room already know about identifiers on the internet and the different kinds of identifiers. This slide is just talking about IP addresses and how important they are for the internet. All hosts connected to the internet require an IP address, there are two different kinds, IPv4 and IPv6. It's a numerical label, it's not a name. It's a fundamental identifier for the internet.

Why DNS? What was the problem with just using IP addresses? Well, IP addresses are hard to remember, they change often as well. There's modern problems as well that came after the introduction of DNS, which is that IP addresses can may be shared. Clients can have multiple IP addresses, servers can have multiple IP addresses, hosts can have more than one IP address, so, which one do you use? You can use the naming system to help with this.

As you probably learned in DNS101 discussion, presentation, DNS System is hieratical, it has a root at the top and then under that, you've got what are called Top Level Domains, here we have some examples of .UK, .ORG, .EDU and underneath that you have the second

level and then third level and then on and on. We can generally think of there being a mapping between names to IP addresses, that would probably be referred to as an A or a Quad A Mapping if you're familiar with those terms. There are other mappings as well, names also map to names of mail servers, there's also reverse DNS Lookups.

Here we have some definitions, there are going to be two slides full of definitions and it's good to go over these before we get deeper into the subject, just so people can be familiar with these terms. These are some terms the RSSAC uses quite frequently, not just in our documents but also in this presentation. The first one that I've already used previously is the Root Server System and this is the set of Root Serves that collectively implements the Root Server. We'll dig down into this a bit later.

Then there's the Root Zone, which is really the data, like on the last slide when I talked about the DNS hierarchy, we're really talking about the data that the Root Server System helps distribute. The Root Zone is the data that the Root Server System helps to distribute. It has no parent and contains all the information necessary to contact the Top-Level Domains underneath it. Finally, on this slide we have Root Server Anycast Instance, all the Root Server Operators are using Anycast, when we refer to individual servers or physical machines, we're really talking about Anycast Instances.

Here we have some Organizational Roles. We have the Root Zone Administrator and this is the organization responsible for managing the data contained with the Root Zone. This is IANA Function

essentially and this involves assigning the operator to Top Level Domains and maintaining their technical administrative details and what that means is that, the ways in which you resolve a Top Level Domain, require that there be a server for that Top Level Domain that can answer queries and the people running that Top Level Domain, will occasionally want to update what server that is and then they have to talk to the Root Zone Administrator.

There's also the Root Zone Maintainer, it's currently Verisign and this is the organization responsible for accepting the data from the Root Zone Administrator and then formatting it into a zone file and then most importantly, cryptographically signing and then after that, distributing it the Root Server Operators. The Root Server Operators, there are 12 of them and these are the organizations responsible for managing the Root Service on the IP addresses specified in the Root Zone and the Root Hines File. These are the individual organizations that are running the physical Root Servers.

I talked a little bit about this already, here we're talking about the difference between the data and how one get's the data. There are servers that serve the data and there's the data itself. The Root Zone is talking about that data, it is Zone File and then the Root Server System is made up of RSOs who have servers who then serve that data. This slide is comparing the two. With regards to the Root Zone, it really is the starting point. It's a list of TLDs and their Name Servers, so how you would resolve. You would go to the Name Servers of the TLDs in order to resolve names underneath those TLDs. The Root Zone is managed by ICANN per Community Policy.

As listed on the last slide, it's complied and distributed by the Root Zone Maintainer, to all the different RSOs and it's the information that's served by the Root Servers. The Root Server System is the system, made up servers, that responds with data from the Root Zone to queries. Currently, there are 26 IP addresses that make up the Root Server System and 13 for IPv4 and 13 IPv6 and there's over 1000 physical instances. The number varies, it's slowly going up over time. Now, we just say over a 1000. We use to have a slide which said exactly the number and then we just kept having to update it and so we said, "Okay, now there's just over a 1000." It's a purely technical role to serve the Root Zone and it's the responsibility of the Root Server Operators.

This slide goes through all the details of a DNS query in response. I'm going to spend a bit of time on this slide and just walk through what happens with a computer when it wants to resolve a name and it's going through a Recursive Name Server.

If we start over on the right, we see we have our user, we have our internet user also know as the client and they're using a computer and they would really like to get to www.example.com. Now, assuming that this Recursive Server was just turned on and it doesn't have anything it's cache, so it's essentially ignorant, it knows nothing except how to get to the Root Name Servers. What that users computer is first going to do, it's going to contact this Recursive Name Server that's located in the middle of the picture and it's going to send a query to that Recursive Name Server saying, "Hey, what is the IP address for www.example.com?"

And that Recursive Name Server, because it's just been turned on and it doesn't have anything in its cache, it will first contact the Root Server System, a Root Name Server. It will most likely send the entire query to that Root Name Server and it'll just ask, "Hey, what's the IP address for www.example.com?" And that Root Name Sever will say, "I don't know the address for www.example.com but I do know how you can talk to .COM and by the way, here's a signature for this data."

And it'll send it back to the Recursive Name Server. The Recursive Name Server then says, "Okay, great. I now know how to get to .COM and oh, I have this signature as well, I will compare that cryptographically to the public part of a Key Signing Key of a the root KSK, which is located within the Recursive Name Server and determine that yes, this answer from the Root Name Server for the location of .COM is correct, I will now contact .COM."

Then the Recursive Name Server goes to the .COM Name Server and says, "Where's www.example.com? " The .COM Name Server says, "I don't know where that is but I do know where example.com is and oh by the way, here's a signature for the data." Then the Recursive Name Server will do basically the same thing, where it says, "Okay, great. I have example.com, I can verify this with cryptographic signatures." And then it will go out to the example.com Name Server and finally get an answer for www.example.com, along with a signature. It goes back to the Recursive Name Server, the Recursive Name Sever says, "Great, I have an IP address for www.example.com now and I have a signature, I'll do the cryptographic comparison. Great, it checks out."

Only now does the Recursive Name Server go back to the user or back to the user's computer and say, "Here's your IP address." What's illustrated here is that through a simple query that initiated by the user, there's recursive, there's an iterative, there's a complex and long-lasting process that goes on in the Cursive Name Server and until it finishes that, it doesn't actually get back to the user.

This is drilling in a little bit more detail about what I just talked about. Root Servers only know what servers need to be asked next. In the last example, they only knew how to get to the .COM servers. They don't know how to resolve the whole name. However, within Recursive Servers, they'll remember those responses, they'll cache those response.

They don't actually have to go to the Root Servers that often and the Time to Live, often called the TTL for that cached information is two days. A Recursive Server will ask a Root Server where say the .COM Name Servers are, it will remember that information for two days probably and then it doesn't have to ask again until that information times out.

There as some modern refinements to DNS. We talked about Cryptographic Signature on DNS data. This is all defined within DNS Sec or the security extensions to DNS. DNS Sec is just used as a short hand term to talk about all the signing and validation that needs to go on in order to cryptographically ensure that the Recursive Servers and the clients are getting the correct information back from the servers. There have been some recent privacy enhancements as well to DNS

because queries can leak information. An original or traditional DNS will report 53 UDP or TCP was just in clear text, there is standards on going to address encrypting DNS transport. We also have DNS over TLS, there's also DNS over HTTPS.

Anycast is another modern refinement. At this point, all of the Root Server Operators that implemented Anycast, we've got a whole section coming up after this about Anycast, we can talk about it then. That leads me right into an explanation of Anycast. In the beginning there was just Unicast and Unicast means that packets from sources all go to the same destination, there's one destination and packets from all sources go to that destination.

Basically, the IP address corresponds with a single server, it corresponds with a single end point. This is fine, unless you want to scale things. Anycast allows scalability, it makes scalability a lot easier and that's important if you want to increase the service, if you're getting more valid queries. It's also really nice if people are launching distributed denial of service attacks against your service and you would like to have some of that traffic offloaded to many different servers, it just becomes a lot easier to handle it.

In Anycast, the mapping between an IP address and end point changes and now you have one IP address mapping to multiple servers, to multiple end points. Different sources will reach different destinations but they're still all communicating or they're still all sending traffic to the same IP address, it's just hitting different physical destinations. Sources get the data faster, there's less intermediate hops because

sources can be closer to their destinations. DDoS attack traffic will be what's called Sink Hold or sent to the closest instance and it won't disrupt other instances.

Here's a pretty simple example of Unicast traffic flow and there's a single shortest route to a single destination. We've got one source and we've got one destination. If you add another source into that, it will go to that same destination. Here's Anycast, you have multiple destinations, that's the purple or blue tear drop looking thing, labeled destination. We still only have one source but you can imagine if another source closer to a different destination going to that other destination. It allows -- it's just a lot easier to distribute traffic.

Here's one of the real benefits, one of the great benefits of Anycast, is what happens under Distributed Denial of Service Attacks is that sources using destinations far away from the DDoS attack don't get effected by it. If it's a local, in some sense of local either geographic or topologically local attack, it will be sink holed at one server and not affect operations and the rest of the world. That's one of the big benefits of Anycast.

Now, I will talk about a bit of the history of Root Server System and the Root Server System today. In 1983 to 86, the Root Server System had four addresses and you can from then it has steadily increased up until 1998, where there are not 13 addresses. Maybe addresses isn't the right term to use there but you can see these changes growing over time and nowadays because IPv6 has been added, we now say the Root Server System has 26 addresses. 13 IPv4 and 13 IPv6. Because of

Anycast, we can say there is 26 IP addresses, 13 IPv4, 13 IPv6, there's actually over 1000 physical instances.

This is a listing of the host names and the IPv4 and IPv6 addresses, as well as the manager, the Root Server Operator for the current Root Server System. All the operators, all the host names have both IPv4 and IPv6 addresses associated with them. You can see they're lettered A through M.

This is a map that's taken from root-servers.org, it's not meant to be accurate, geographically accurate. What you can do is, if you go to root-servers.org you can pull this up, it's right on the front page and then you can actually track down individual cities, you can look at a continent and then a country and then a city and then see which operators are operating servers or operating instances in those cities. This is at the highest level, where you're zoomed all the way out, the mapping software just groups things.

This essentially shows that there are instances all over the world, on every continent, I shouldn't say that because I don't know if there's one on Antarctica, maybe someone can correct but at least every continent other than Antarctica. If you're interested in drilling down more into this, it is fun to go to root-servers.org and drill down into the map and see what city near you have root server instances and who operates them and that kind of thing.

This is showing the process of changes to the Root Zone and how they're provisioned into the actual Root Servers and then how

Resolvers go and get them. On the left we have the TLD operators. A TLD operator let's say needs to make a change request, they would contact IANA and they would say to IANA, "Our TLD, our Name Servers we're using this IP address, we would like to change them to this other IP address."

And they would tell IANA that and then IANA has a bunch of verifications procedures that it goes through to make sure that it really is talking to the TLD operator, that this is a good change, this isn't going to break anything. Once IANA approves that change, they we will then pass that when the IANA tells the Root Zone Maintainer that, "Here is a new Root Zone file." It will just update that information and the next Root Zone file that goes out to the Root Zone Maintainer. The Root Zone Maintainer than complies the Root Zone and cryptographically signs it and then they distribute it to the operators.

Then on the right there, all those little bubbles with RS, those are supposed to represent the individual instances, there's a whole lot of them, for each operator there's a whole lot of them. Then on the far right we can see the Recursive Resolvers sending queries and getting responses. This is the flow of how information, a very high-level flow of how information in the Root Zone changes and the process for that.

A bit more information about the Root Server operators, as I've said, there are 12 of them and they're primarily focused on reliability and stability and accessibility for all internet users. They cooperate with one another through RSSAC and Root Ops and other venues. Focused on professionalism and there are different kinds of organizations,

they're not all non-profits, they're not all governments, they're different kinds of organizations and they're diverse in terms of technically organizationally, geographically, their funding models.

This is some of the ways in which they coordinate. I mentioned some of the various industry meetings and bodies such as ICANN and RSSAC, the ITF, RIR meetings, Network Operator Groups, DNSORC and they also use other different kinds of tools like any organizations that need to coordinate do. They share data amongst one another and they also perform periodic activities, like table top exercises and whatnot to plan for emergency reposes and whatnot.

RSOs are involved with the operation and evolution of the service. Evaluating and deploying suggested technical modifications, maybe to a protocol, like the DNS Protocol. They participate in IETF for that. Making sure that stability, robustness and reachability are maintained. Operators not involved in the policy making and they're certainly not involved in data modifications. The RSOs just publish the data, they're not involved in what the data is.

Piggybacking on that, this slide shows some of the myths that the RSSAC and that I've encountered over the years, as well as the reality. The first myth is that Root Servers control where internet traffic goes and the reality is that packet routers controls where internet traffic goes. Root Servers just answer queries that are given to them from Recursive Servers.

Another myth is that most DNS queries are handled by Root Server and most DNS queries are not handled by Root Server mainly because of caching, because Recursive Servers just remember the answers, they got from the Root Servers so they don't have to ask every time the same question comes in. Another myth is that administration of the Root Zone and Service Provision are the same thing. Again, this goes back to what's the difference between the data and serving the data? Administration of the data is very different from answering queries about the data.

Another myth is that some of the server identities have special meaning. None of them really have any specially meaning. They don't. Another myth is that there are only 13 Root Servers. No, because of Anycast, there are over 1000 but there are only 13 Technical Identities. Root Server Operators do not just operate completely independently, they do cooperate with one another, we talked about on the last slide. This last one, Root Server Operators only receive the TLD portion of a query.

When I walked through the whole how DNS Resolution woks, I had the whole query string going up to the Root Server and that is traditionally how things have worked and that is probably how things work 90 percent of the time now. That's why that usually us bolded, there's a new technology called QNAME Minimization, which is a new privacy technology, I guess we could call it, coming out of the IETF which aims to change that but that's not very widely deployed. It is getting more and more deployed but currently the reality is that Root Server Operators usually receive the entire query.

If you're running a network, something to think about with regards to your DNS and your interactions with the Root Serves. You probably want three or four nearby instances. Neary by meaning, topological and network closeness, not necessarily geographic closeness. You can maybe increase peer in connections if you're experiencing latency. There's many different things you can do.

Another thing is to turn on DNS Sec Validation Resolvers, this can ensure that for data is signed in the Root Zone, you're getting the correct data, you're getting unmodified IANA data, between your Recursive Server and the Root Servers and other Authoritative Servers, no one's tampering with the data in transit. DNS Sec Validation for Sign Data, you're validating that data in your local resolver, that just helps.

If you're interested, you can also participate and contribute to the RSSAC Caucus and my colleague Ozan will talk a bit more about what the RSSAC and the RSSAC Caucus are.

If you are a network operator and you are interested in hosting an Anycast Instance, you can talk to an RSSAC member after this presentation, you can certainly bring that up during the question and answer period or you can send a mail to that address ASK-RSSAC@ICANN.ORG.

Now, I'm going to the presentation over to my colleague, Ozan Sahin, who's going to go through the Organizational side of things.

OZAN SAHIN: Thanks, Andrew. Hi everyone, my name is Ozan, I am an ICANN Org member, supporting the work of Root Server System Advisory Committee or RSSAC.

Let's kick off with the role of the RSSAC, it has a narrow one. The role of the Root Server System Advisory Committee is to advise the ICANN Community and the ICANN Board on matters relating the operations, administration, security and integrity of the internet's Root Server System.

One this slide there are two notes on what RSSAC does and what it does not. It's a committee that produces advice primarily to the ICANN Board but also to other ICANN bodies and other organizations involved in overall DNS business. The Root Server Operators are represented inside the RSSAC but the RSSAC does not involve itself in operational matters.

If you look at the organization of the RSSAC, it is composed of appointed members, appointed representatives of the Root Server Operators. There are also alternates to these representatives and there are liaisons. Also, there is this other body called RSSAC Caucus, it's a body of volunteers subject matter experts, DNS subject matter experts. The RSSAC Caucus members are confirmed by RSSAC based on the Statement of Interest. If you want to become a member of the RSSAC Caucus, you basically submit as my colleague Andrew just discussed, your Statement of Interest and the RSSAC confirms your membership.

We have two co-chairs, they're in this room, Brad Verd and Fred Baker. I also want to note that the RSSAC is transitioning into a chair/vice chair model. By the end of the year they're going through the elections for the vice chair. The RSSAC will be -- the leadership will be composed of a chair and a vice chair.

As I just said, there are liaisons at RSSAC, four of them are incoming liaisons and four of them are outgoing, meaning there is a liaison from IANA Functions Operator, one from Root Zone Maintainer, one from the Internet Architecture Board or IAB and one from the Security and Stability Advisory Committee which is another advisory committee within the ICANN system. There are four outgoing liaisons, one is to the ICANN Board, one to ICANN Nominating Committee, one to Customer Standing Committee and one liaison to Root Zone Evolution Review Committee or RZERC.

The RSSAC Caucus has more than 100 members who are DNS technical experts. As I said, anyone interested applies to become a member of the RSSAC Caucus by submitting their Statement of Interest and they get public credit for the work contribution to RSSAC publications. The Caucus indeed adds to the transparency of the RSSAC, so that you can engage in the work of RSSAC by becoming part of the RSSAC Caucus. As I mentioned, these are the DNS experts who bring their expertise to the publications.

There are currently work parties within RSSAC. One is the study of modern resolver behaviors; it studies the behavior of existing deployed software and recursive resolvers through both code basis

and available datasets. The other one is the expectations of the Root Server System and Related Metrics. It is tasked with defining system wide an external verifiable metric which demonstrate the RSS as a whole is online and serving correct and timely responses to end users.

There are some tools and mechanisms that contribute to the transparency of the RSSAC and the Root Server Operators. To name a few, there is RSSAC.ICANN.ORG webpage where you can go and find the names of the RSSAC members, RSSAC Caucus member, you can access the publications. You can also find the minutes from the RSSAC teleconference meetings. Also, the RSSAC has public meetings, if you are interested, you can participate in the meetings.

The RSSAC also has meetings with other groups at ICANN during the ICANN public meetings, for instance during this meeting, RSSAC co-chairs briefing to the Governmental Advisory Committee, also ICANN Board will be meeting with the RSSAC and the RSSAC will have a closed meeting with the Security and Stability Advisory Committee. It's talking to other groups. The RSSAC has a publication 000, which defines its operational procedures, which also adds to the transparency.

RSOs also have some too, they have the root-servers.org webpage. I think the map that my colleague Andrew just showed with instances all over the world can be found on this page. Also, the RSOs Root Server Operators have their individual pages and they publish collaborative reports on major events and also again, Andrew just

talked about it, if you have questions you can submit your questions to ASK-RSSAC@ICANN.ORG and get answers.

In the second phase of the presentation I'll talk about work underway on the Root Server System Evolution. Let's start with looking at the timeline of this work. More than a year ago the RSSAC037 and 38 were published. They basically proposed a new model in the governance of the Root Server System. Then, the ICANN Board directed ICANN Org to look at these documents and work on them. In April 2019 ICANN Org finally published what is called a Concept Paper. In August 2019, the public comment period for the Concept Paper closed. The consultation from different groups were received. By looking at all the comments the next steps that are envisioned are by January 2020, there will be a Governance Working Group in place, which will work on developing the model in 2020 and 2021. By 2022 it is expected the new model will be implemented.

Let's look at what the RSSAC037 was about it. It defined 11 principles for the operation and evolution of the Root Server System. Basically, it proposes an initial Governance Model for the Root Server System and its operators. It also demonstrates how the RSSAC037 model works through a set of scenarios on designation and removal of Root Server Operators.

On this slide you see three recommendations that compliment RSSAC037. The first one is to initiate a process to produce a final version of the model based on 37. Also, the second, estimate costs of the Root Server System developing the model. Initial efforts also

should focus on developing a timeline and also implement the final version of the model based upon principles of accountability, transparency, sustainable service and integrity.

On this slide you will see a graph depicting what the proposal was a graph. You see three different areas, one is governance, the other DNS root operations and then there is onboarding and offboarding of Root Server Operators. In the governance area you see three stakeholders there, the ICANN Community, IFT and Internet Architecture Board and the Root Server Operators. The slide also gives the five functions proposed by the model.

Namely they are, the performance monitoring and measurements function, designation and removal function, financial function, there strategy architecture policy function and there is secretariat function. Five functions were all proposed by the model. Then you note at the bottom of the slide there is some performance metrics, which will be used on the onboarding, offboarding area of the Root Server Operators. This relates to designation of Root Server Operators and the removal of the operators.

We just talked about the functions and going back to the concept model, it envisions the following structures based on the 37 model that corresponds to the five functions. One is the Root Server System Governance Board, the other one Root Server System Standing Committee, Root Server Operator Review Panel, the last two functions which are the finance function and secretariat functions, ICANN Org.

The Concept Paper outlines the community driven process, to finalize a new cooperation and governance model for the RSS based on recommendation one in RSSAC 38, which is another publication that relates to the evolution of Root Server System In phase one the ICANN Org is reviewing and evaluating RSSAC037 at the direction of the ICANN Board, which is done.

In the second phase, RSSAC037 Concept Paper Governance Working Group, document where available for public comment and we talked about it. In the third phase, it related to developing a new cooperation and governance model for the RSS and this has two tracks, one is the structural track, the other one is administrative track. Structural track governance work group develops a model, in the administrative track it's to plan for implementation of the governance working group model lead by ICANN Org.

What is Governance Working Group and its composition? It's composed of representatives from the RSSAC, ccTLD Name Supporting Organization, Register Stakeholder Group and the Security and Stability Advisory Committee. There will also be liaisons from the ICANN Board, the IANA and the Root Zone Maintainer.

The Governance Working Group is tasked with working out the details of the model. The Concept Paper also outlines some of the guidelines for the Governance Working Group, which are committing to a timeline with clear milestones, working opening and transparently, seeking informed contributions when necessary, also embracing the

principles outlined in RSSAC037 and basically referring RSSAC037 Concept Paper and the public comment feedback that were received.

We are now at the questions and answers part of the session. In the room we have Root Server Operator representatives who are members of the RSSAC. If I may please suggest them to come on the stage and take their seats to take questions from audience. There will be some roving mics, if you raise your hand, we will give you the mic so that you can ask your question.

STEVE CONTE: While they get it setup, I do have a question over here but we'll wait till all the Root Servers are seated. Our first question was over on this side over here.

UNKNOWN SPEAKER: I was just wondering, coming from a nontechnical background, I was wondering how do you determine when and where you need another instance of the Root Server? How is it different with the Name Server and what is a Recursive Name Server, and what's the difference between that and your Regular Name Server?

FRED BAKER: Well, we decide to place a new Root Server by throwing a dart a board and seeing where it lands. Not so. We actually have a process that we go through that starts out with need. Why do we need a new Root Server? Why would we need a new RSO? Then in the event that there

actually is a valid need, then there's a set of principles by which we would identify a company or an organization that can do that, that wants to do that. I would frankly suggest that you read RSSAC037 to cover that because it's -- there's a fair amount there.

WES HARDAKER: I actually believe, Fred, that he asked about a Root Server Instance, not a Root Server Operator.

FRED BAKER: Oh, okay. A Root Server Instance, you would come to one of us, one of the Root Server Operators and for example, my company ISC, if you go to the webpage at ISC.ORG, you'll find something right on the page that says, "Click here if you want a new Root Server or Root Server Instance."

And I imagine that's true of all of us. Then we'll talk with you about your requirements. We have some expectations; we're going to want both IPv4 and IPv6 connectivity to the system. There needs to be adequate bandwidth. There needs to electricity, that kind of thing. We ultimately exchange an MOU. Then start operating and we would operate the server, it would sit in your rack but we would operate it remotely. Basically, if you want one you ask and we start that dialog.

WES HARDAKER: If I can add to that a little bit. I think what you're hearing is that we get requirements from multiple places. We get external requirements

where people are sending up new internet exchange points and things like that, which is the last one we placed in was -- new internet exchange point that came to us and said, “We’re going to be brand new and we have nothing there, are you willing to help? ” We were willing to do that and we were the first one in the door because of that but we have internal requirements of how we analyze, how the world is being served right now.

I’m currently in the process of deploying more and I’m trying hit geographically disperse regions to get as good coverage as we can. It’s a combination of both internal and external metrics and we value input from anybody that believes that you don’t have adequate service in your area.

BRAD VERD: What you’re hearing here is different approaches because each of the Root Server Operators chose it in their own way. It could be based upon a need, an actual traffic need. It could be based on geopolitical need. There’s any number of different needs that would justify an instance.

I think the second part of your question, what was the difference between a Recursive Resolver and an Authoritative Server, we run Authoritative Servers, we’re Authoritative for the Root. Your Recursive Resolver is what you talk to within your ISP. Your Recursive Resolver is your intermediary between all the authoritative.

For example .COM is an Authoritative Server, the Root is an Authoritative Server, .US and so forth. You, for the most part, probably don't talk directly to us, you talk to your Recursive and your Recursive then talks to us when it's needed and if the answer isn't already cached. I hope that answered your question.

OZAN SAHIN: We have another question here.

UNKNOWN SPEAKER: Thank you. I don't know if this is for the panel, the gentleman before or if I'm in the wrong room. It's reading DNS over HTTPS, so my understanding is that Firefox and Chrome are going to be moving that way in very short order, like a matter of weeks or months, which means 95 percent of DNS traffic could be encrypted, I don't know. I see a shaking head up there.

That's my whole point, is that I don't know anything about, how it's going to effect it, how it's going to effect that entire presentation at the start of this meeting on how DNS works, if it's going to change or if it's just going to be the same but hidden? I didn't see any panels on it in the entire schedule of ICANN and it seems like something really important. If can't find out about here, maybe you guys can tell where I can learn about it?

WES HARDAKER: I also sit on the Internet Architecture Board, which is in participation with IAFT where that work is going forward in terms of a standardization process. A little bit of facts about the deployment that Firefox and Chrome are doing. They're doing it very differently so don't mistake how that's going to happen. The other thing that they're doing, that is only between A, the web browser and a resolver.

In the case of Firefox they are picking a Resolver, the default Resolver is run my Cloudeflare, they will have a drop down menu in your configuration when you can pick the resolver that you want to use and they are turning it on by default in the United States this month, the rest of the world they're looking for other partners to do that. That resolver is not doing encrypted DNS to the rest of the system and that includes the Root Servers, it includes the TLD Servers like .COM and the ccTLD servers, example.com and things like that.

Chrome on the other hand is doing something slightly different. They are trying to see if your local ISP's resolver supports DoH and if it does and if it's on their approved list, they will communicate with DoH to the ISP. They are not doing what Firefox is doing and sending it all to one location. There is a very different deployment issue and unfortunately there is a lot of confusion because that information is changing so quickly, even Firefox desire to only do the US with Cloudeflare was a very recent decision and what happens next week is subject to debate.

In two weeks, it will be much more heavily discussed at the IETF which the Internet Engineering Task Force and that will be in Singapore.

That's where the technical conversation is going on. In a couple of months, it's going to be different, it's changing very rapidly right now.

BRAD VERD: If I can add to that. Again, reiterate that that is between the client and the resolver where the encryption is happening and not between the resolver and the authoritative right now. I will say, this has been talked about at ICANN, there was a high interest topic in Marrakesh regarding this, so I would go back to the Marrakesh agenda.

I don't remember what day it was on but there was a big presentation put together by SSAC and I think it was the CCNSO, it was a deck in there that talks about it and they continue to talk about. I know SSAC is working on it and there's other work being done. This has lots of eyes.

FRED BAKER: So Wes, let me ask a question about DNSSEC; browsers generally don't implement DNSSEC, they depend on somebody else doing that, so when the browser goes to DoH, is DNS verified?

WES HARDAKER: That's an excellent leading question, Fred. Very good question. There are two aspects to security in general, that most people think about. There is encryption, in other words is your data being protected and then there's the question of whether your data is authentic, is it

actually the correct data? It could be encrypted and still be wrong. You could get the wrong encrypted file for example.

In DNSSEC, it protects from the origin, from the place where the data was create, in this case it's IANA and through the Root Zone Maintainers signing of that data, the rest of the Root Zone data and actually for that matter most TLDs and anything else that is signed below, it doesn't matter, you could actually hand that to me on a piece paper, I could read and scan it in and I'm able to verify the signature against where it was originally created from IANA and all the way down the tree.

Fred's leading question was, DoH also does integrity but it's only between two points, so if that entity up above, that resolver that you were talking to insecurely, it wouldn't know and it would hand it to you in an unverified way. Some DoH resolvers will be doing DNSSEC validation. If you know that you're talking to a DoH resolver over a secured integrity protected channel and you know that they're doing verification, you're probably secure end to end. Cloudflare is one resolver that I believe is doing validation by default, I don't know about the rest.

BRAD VERD: But that's something worth checking if you're looking into that.

STEVE CONTE: Is that a quick followup that you have? Okay.

UNKNOWN SPEAKER: The data's not visible to the ISP in HTTPS environment, they can no longer see errors? Everything that's flying through their network but you guys can still see it? Who will still be able to see -- who will have visibility into the DNS requests and who will not in a DNS or a HTTPS world?

WES HARDAKER: That's very hard to answer today because as I said, next month it'll probably be different. That's true for people using Firefox that's communicating to a DoH provider somewhere else but that DoH provider, like Cloudeflare will be able to see it. From there, it gets distributed to the entire -- at some point you have to ask somebody a question, at some point you have to go to somebody and say, "I have to ask a question."

That person you have to ask a question to, like where is that website, they're always going to be able to see it. There's always somebody that has to know your question. For Chrome on the other hand, since they are going to be DoH to the ISP, to your resolver at your ISP, that's not going to change the ISP's visibility.

It greatly depends on the deployment situation and Chrome and Firefox are doing different things. You mail reader on the other hand, there is no mail reader plans to do DoH for example. If you're doing mail inside of a web browser, then there will be. It's a very -- it's not a yes or no questions. Does that make sense?

FRED BAKER: So, now it seems like a good point to mention QNAME Minimization. This is a project that's being worked on in the IETF and will be coming to some software near you some time. The idea is to betray as little information as it can and still get the question answered. If I was looking up www.example.com for example, I might ask my Recursive Server and the Recursive Server would say, "I don't know where .COM is, I haven't figured that out yet."

So, it would now, instead of sending the entire name up to a Root Server, which is what it does now, it would send COM to the Root Server, the Root Server would know that it asked .COM and it would give it that name and then it might ask example.com. Really, only that Recursive Resolver would have the access to that information. You want to be watching QNAME Minimization.

UNKNOWN SPEAKER: The other piece of this is aggregation. If you've got thousands and thousands of people going to the same Recursive Resolver, yes that Recursive Resolver knows that you made that request but now if you're getting thousands of requests from that same Recursive all over, it can't necessarily be attributed to the individual, just somebody, one of the many, many people that actually that use that recursive.

STEVE QUESTIUON: Thank you. Next question over on this side.

UNKNOWN SPEAKER: I want to know if the DSOs retain logs and if so, are there any rules about privacy for example? Just out of interest, how are the DSOs funded for the domain service operations?

BRAD VERD: Did you mean RSOs? We keep hearing DSOs, I just wanted to make sure you meant RSO, correct?

UNKNOWN SPEAKER: Yes, sorry.

BRAD VERD: Okay, I'm going to go backwards. I didn't hear the first part of the question but the last part is very simple on how they're funded. Right now, it's a nonfunded mandated, it's all volunteer basis. As the internet grew organically, there were volunteers to run these Root Servers and they were added over time, basically between 82 and 98 and there hasn't been a new Root Server since 1998. I think it was 2001 Anycast introduced and started being used by the Root Servers, so we went from 13 identifiers, 13 servers to now we're at thousand servers but the same 13 identities. By using Anycast, we're able to spread that far and wide and that's all done, that's all self-funded by

each of the organizations. So do you remember the other parts of the questions, or could you repeat them?

UNKNOWN SPEAKER: I'm interested if you retain logs of requests and if there are any rules about privacy of those?

BRAD VERD: The only thing I could refer you to would be the yearly there is what's referred to as Diddle collection which is a Day in the Life and it's referred to as Diddle and it's a collection of a 48 hour window that allows researchers a whole bunch of data to see what the internet does in a certain day. All the Root Server Operators can contribute to it, along with a lot of other TLDs and ccTLDs and large organization, big DNS operators, it's community effort and that data is stored in the DNS Database and in order to access the database you have to become a member and sign confidentiality stuff and whatnot there.

WES HARDAKER: One further point, which is that a lot of operators anonymize the data before giving it to OARC and what that really means is they anonymize the IPS, the requesting IP address is often a resolver in the first place so it doesn't really tie back to one particular end users machine, it ties back to a resolver that's serving lots of stuff. I think especially with the more recent ones, especially since GDPR I think most operators are

anonymizing but you'd have to go talk to each one, I don't remember the current status of who's anonymizing and to what level.

OZAN SAHIN: We have another question.

UNKNOWN SPEAKER: First, I must apologize because I think this question would be kind of hard to answer. Earlier in the slides I notice that most of the organizations that owns the IP address for the Root Servers are all American Organizations and I was wondering in a world for example the American Government is not really going to commit to neutrality; what can ensure or is there any mechanism that ensures that Root Servers for the internet backbone stay neutral?

FRED BAKER: In answer to that, and there's probably more than one answer but one of the things that the RSSAC is working on right now is a question of how do you measure the system? One of the measurements, one of the things that we're concerned about is whether the RSO that's being measured and we're measuring all of them but it talks about one of them at any given time, is actually serving the system that came from IANA.

If the Root Zone that you're getting from a particular server is different, than that's a violation of some things that we consider fairly important. That would be considered a bad thing. What we do, is we

literally download the information, every few minutes something like that, download the information from IANA, we serve it for a while and we'll download so more. We're always passing what IANA gave us.

Now, what did IANA give us? The TLDs, the ccTLDs, the gTLDs, turn around and inform the IANA, I had these names and they had these records associated with them and the IANA is the neutral party and manages that. I think what you're really depending on in terms of neutrality is the ethic of the IANA, of the ccTLDs, which are doing this as a business and of the RSOs. Does that answer your question?

WES HARDAKER: Let me add a little bit more, Fred, before you go on. I strongly recommend you read RSSAC023, which is the 0 to 23 document, it's the history of how the Root Server System got established the way it is today. In fact, we were talking about it earlier today in our RSSAC meeting, which you're welcome to come and listen to as well.

It explains how we got to the organizations that are currently serving it and it's entirely due to history, it hasn't changed in 20 years and one of the goals of RSSAC037, which is an architecture for how do we make changes to this process in the future since the last person that made changes passed away 20 years ago, is on the table.

More importantly, the same conversations that was I talking about earlier with respect to DNSSEC, if you are doing a DNSSEC validation and the domain that you are doing is validated from the top to the bottom, then you know it hasn't been modified and it really doesn't

matter what countries it went through, it's totally politically independent because technically it would be [inaudible] impossible to fake that data. That's the safest thing I can tell you to do, is make sure that you're using a DNSSEC validating resolver.

BRAD VERD: And I'll just add a little bit more, which is RSSAC037, which has been published and is out there for everybody to read and in there states the guiding principles as defined by the Root Server Operators. One of those guiding principles is to remain neutral, it's an Apolitical view, there is no politics involved, we serve the Root Zone that is given to us from IANA. Regarding the comment around US based, that is purely a result of organic growth. The internet started in the US, it grew in the US, there was a need for Root Server Operators and that's how it happened, there was no other reason other than organic growth.

FRED BAKER: And we do of course have Root Server Operators that are outside the US, we have one in Sweden, one in the Netherlands and one in Japan.

STEVE CONTE: I'm going to take the next one over here in the shadows, over to your right.

UNKNOWN SPEAKER: I know that denial of service and distributed denial of service attacks, like many things security, it's cat and mouse, the attackers get strong, security gets stronger and vise versa and it continues. What exactly is being done to keep the Root Servers protected continuously? Does RSSAC meet up with SSAC and have meetings on keeping the Root Servers protected or how does that work?

BRAD VERD: The first place I can send you now is just recently the Root Server Operators published a document that kind of addresses your question directly. It's threats to the Root Server System and what they do and what they've done to mitigate some of those overarching threats.

Regarding RSSAC, RSSAC is contact conversation with the Board, with SSAC regarding any type of threat to the system, which is why I think -- there was an early question about DoH and DoT, we've talked about that and the implications to the infrastructure when something like that happens. Those conversations are happening all the time. Those are not really operational though, the operational questions happen within the RSOs and between the RSOs and in the event of a DDoS or something else, the information is shared and mitigations happen.

WES HARDAKER: Where can he find that document Brad?

BRAD VERD: I'm sorry. That document is not an RSSAC document, it is located on the Root Server webpage, which is at www.root-servers.org; I think it's on the top of the page.

UNKNOWN SPEAKER: Perfect, thank you very much. Am I able to ask one last question? I know with DNSSEC misconfigured servers, there is also the potential to launch amplification attacks and I know that there's a lot of push for rolling out DNSSEC, I'm curious if there is any possible way to also push hardened DNSSEC configured servers? For example, adding rate limiting and other methods? I know there was a large discussion back in 2013 on behalf of ICANN on hardening DNS security.

WES HARDAKER: Can you refine your question about where DNSSEC misconfigured servers caused problems because that's either too broad or you're referring to something very specific?

UNKNOWN SPEAKER: So, if someone sets up a misconfigured DNSSEC supporting server and someone launches UDP based attack and they spoof an IP address, that you can basically just gather a list of misconfigured DNSSEC servers, and then it just allows the amplification attack to be stronger than a normal DNS amplification.

WES HARDAKER: You're worried about a reflection attack. DNSSEC adds a whole lot of data to the signatures become quite large, the keys are quite large and you're absolutely right, especially in the past DNSSEC was used as a reflection attack because you send a very small packet faking it from the address that you want to attack and a very large amount of traffic goes that direction.

Today, most servers and I'm not talking just the Root Servers but most servers have technologies using something that's called a Responsive Rate Limiting and what that does is it won't allow one server to ask too many questions, it just cuts them off and says, "I'm not talking to you anymore. You can still talk to me but you have to come back over TCP, which is much harder to spoof."

What the settings for each Authoritative Server does is very, very different. I can tell you that before most people turn that, you would see just daily spikes of traffic because everybody was using it and then years following, most people's traffic graphs just went fairly flat because people realized that that's no longer viable way. You don't see reflection attacks in DNSSEC much anymore. I'm sure they still exist somewhere.

STEVE CONTE: Alright, thank you very much. Any other questions in the room? I have checked online, there are no questions online. Last chance for questions in the room.

WES HARDAKER: These are all great technical questions, thank you.

STEVE CONTE: I would like to thank Andrew; I would like to thanK Ozan for a wonderful presentation. I'd like to thank our Root Server Operators for answering these questions, and I'm going to do a shameless plug, I know the day's almost over but for the How It Works Series, we have a brand-new session starting at five in 512G. We have Aaron coming to talk about Regional Internet Registries, what they do and a foundational knowledge on that. I ask you to please join us for one more session, five o'clock, 512G for a session on Regional Internet Registries. With that, thank you, gentlemen. Thank you, Andrew, thank you, Ozan, and thank you guys for your time.

WES HARDAKER: Can I make one additional plug?

STEVE CONTE: Yes, please.

WES HARDAKER: With a conflict unfortunately, Steve, I apologize. There is also a DNSSEC for Beginner's Talk that's going on also at five o'clock.

STEVE CONTE: Alright, well you've got to choose your poison then. Thank you, all.

[END OF TRANSCRIPTION]