KOBE – How It Works: Root Server Operations
Sunday, March 10, 2019 – 15:15 to 16:45 JST
ICANN64 | Kobe, Japan

CATHY PETERSON: Good afternoon. This is How it Works, Root Server Operations, and we'll start in a few minutes. Thank you. Good afternoon, again, everyone, this is the How it Works tutorial on Root Server Operations. This afternoon we have two presenters, Steve Sheng, from ICANN Policy Development and Carlos Reyes will also take over later on. Thank you.

STEVE SHENG: Thank you, Cathy, and good afternoon. My name is Steve. As Cathy mentioned I work in the policy department and our team supports the Root Server System Advisory Committee. Today, my colleague Carlos and I will give a brief tutorial on the root server system and then we have a standard presentation we'll run through, probably in about 20 minutes or so, and then we have the root server operators in the room to answer any of your questions. So, let's get to it. We'll briefly cover an overview of the domain name system, an explanation of Anycast, the root server system today and its features, and my colleague Carlos will cover RSSAC and recent RSSAC activities.

So, just to recap, the fundamental identifier on the Internet is the IP address. All hosts connected to the internet have IP addresses. It's a numeric label. As you see here, we have IP version 4 or IP version 6. There are two examples there.

So, why the domain name system? Well, the original problem is IP addresses are hard to remember. So, I just showed you the IPv4 and the IPv6 addresses. Those are hard to remember, and also sometimes they change. So, you need a more stable identifier to translate the names to the IP addresses. So, that's the original problem.

The modern problem we're having is IP addresses may also be shared, and multiple IP addresses may serve as a single point of entry to a service and which IP address to use. So, very briefly, the domain name system is a local mechanism for translating objects into other objects. So, now here, I didn't just to say just to translate name into IP addresses because that's the most common use, but there are other uses. You can find the mail servers, the IPv6, and also—reversely—have an IP address to look for the domain itself. So, it's a hierarchal database. It's globally distributed, loosely coherent, scalable and dynamic.

In this tutorial, it's very important to distinguish two concepts, the root zone versus the root server system. The root zone is the list of the TLDs and their name servers. So, if you go on to the IANA

website – and you can actually see the latest root zone file – it's managed by ICANN per community policies. So, the policies are developed, for example, the delegation policies, the re-delegations, and those policies, the implementation of those dictates what goes into the root zone. It's compiled and distributed by the root zone maintainer, Verisign, twice a day, and it's distributed to all the root server operators.

So, think of it as the database content. The database content in the root servers, and so the root server system responds with the data from the root zone. Currently, it's distributed with 13 identities over 1,000 instances at physical locations worldwide. The root server system, there is purely technical role to serve the root zone, right, and that's the responsibility of the root server operators. Think of the operator to serve the zone. The zone is the content, the TLDs and their name servers.

Some high-level definitions, the root server system refers to the set of root servers that collectively implement the root service. So, you have a service and you have a set of servers that implements the service. The root zone, as I already mentioned is a DNS zone at the top of the DNS hierarchy. So, if you look at this hierarchy, the DNS is the hierarchal database, and on top of that is the root.

And so, for Anycast instance it is a one-network location responding to DNS queries on the root server operators' IP addresses. So, all the root servers use the anycast technology, where you have many, many servers on the same IP address, responding to queries. So, one instance is that one network location. The root zone administrator is the organization responsible for managing the data contained within the root zone, which involves assigning the operators of top-level domains and maintaining their technical and administrative details.

So, prior to IANA stewardship transition NTIA served as the root zone administrator because these changes need the NTIA approval. After the transition, it's really a two-party – I'm not sure we'll still use the administrator role here. The root zone maintainer, as I just mentioned, is the organization responsible for accepting service data from the root zone administrator and formatting it into zone file format, cryptographically sign it – and so they do that twice a day – and distributing it to the root server operators.

So, with that definition, the root zone administrator today would be ICANN, and the root server operators are the organizations responsible for managing the root service and IP addresses specified in the root zone and the root [inaudible] file. So, you can

find the latest root zones and the root [inaudible] files on the IANA website.

Very quickly, the domain name, the resolution process, the root server doesn't know everything, it only knows, at the top level, the domains and what is the list of the name servers. So, if you find dot-com, it will give you the list of the dot-com name servers, and vice versa. There is caching of the previous answers, which means you don't have to ask twice. You can cache it, and within the time that the cache is still valid, it no longer goes back to the root for queries.

There are some modern refinements to the DNS. Obviously, we have the DNSSEC, which is the security extension that adds a cryptography signature on the DNS data, so you have the response along with the cryptography signature and you can verify whether you're really receiving the response from the party that it claims to be.

There are privacy enhancements. The DNS queries can leak information. So, for example, if you try to find the IP addresses for www.example.tld, that information goes to the root, goes to the TLD servers, goes to the example, .dot-tld, so the whole query strain is exposed and sometimes it's not necessarily that there are current standards working, going on to address this, what's

called the QNAME [minimization] and there's also encrypting DNS over a TLS, what we call a DOT.

Finally, the Anycast technology is you have multiple servers sharing the same single IP address, so this improves the latency and resiliency that protects against DDoS attacks. I'll share this and go into this in a bit more detail in the later slides.

Again, the domain name resolution process, I already mentioned this before. So, for example, on the right you have an Internet user trying to visit the webserver for www.example.com. [His] computer, there's a resolver library, which first of all goes to the root to ask the root if they know the IP address for www.example.com. The root says, “Well, I only know the name server for dot-com,” and so then they went through the same process, they went to the dot-com name server and the example.com name server to get the IP address for www.example.com. And, after they get the IP address, they declined to establish the connection directly with that server to get the relevant information.

So, here, it's important to note that the root servers are the entry point to the system. Caching is used throughout to avoid repetitive queries. I mean, the example here is kind of a very simple example and it assumes there is no caching here.

The DNS resolution proceeds the actual transactions the users want to do. So, for example, you want to visit a website. You have this transaction, a DNS resolution, before establishing connection with that website directly. And in these days, for any website, they are pulling content from different services, and you end up making lots of DNS queries before loading up the website. So, therefore, the query volume per day for the global DNS system is very high.

Next a quick explanation of Anycast. So, we want to distinguish two terminologies, unicast, uni and any. Unicast is packets from sources all going to the same destination, a single instance serving all the sources. So, in the case of a denial, a distributed denial of service attack, all traffic goes to that single instance.

With Anycast, you have multiple instances, all advertised, that serve the same IP address, and the sources reach the destination based on the routing policies. So, the routing policy selects the instance closest to that, based on the policy to receive the traffic, and the denial of service attack traffic is sent to the closest instance.

So, here is an example. You have a source, you have a destination, there's only one, and then the source, go through a routing policy, select the shortest route and make the connection there.

With Anycast you have multiple destinations in blue that serve the advertisers, same IP address. So, you have a source that, based on the routing policy, so it goes to the nearest destination. So, this is the destination on the left-hand side. So, the benefit of this is in case there is an attacker for the destination, the routing sends that traffic to the nearest instance that serves that IP address, advertises that, but the traffic for others, the other destinations, is unaffected. So, that adds resiliency to the system as a whole.

So, root server system and root server operations. So, the root server system, really, I think the key concept here is the organic growth. The root servers grow with the DNS system.

So, 1983, in IETF, when the DNS system is proposed, so to test the DNS, I think Jon Postel set up the first root server at ISI, where he was at, and then there were four addresses from 1983 to 1986. So, that was the early stages of the DNS, really test out how the DNS worked. In 1987, really, the network role, so it's to accommodate additional root servers that were set up, for example, to accommodate the NSFNET, the research network in the U.S. for root service. And then in 1991 and 1993, this is too for the global growth, the servers are set up, for example, I think in Japan, in Europe and then in 1998 there are 13 addresses.

They all respond to the same information and so you have 13 addresses, but there are a lot more name servers serving the root

zones here. Today, there are thousands of those. The root server operators use the Anycast technology to solve the scaling issues. So, today we have 13 IPv4 and IPv6 address pairs, [serving] a thousand international instances. So, these are the identifiers today and their managers. Oh, it's a map. Okay, so if you, on your computer, you go to root-servers.org, it will show you a map where you can zoom in to see where the instances are in the different geographical locations.

So, here is an important diagram at ICANN about root zone management and resolution. From the left, you have TLD operators. So, this could be gTLDs or ccTLDs that make a change request, changing the name servers and changing the other information. That request goes through IANA and IANA and the root zone maintainer perform independent tests to check the request, and then once that request is approved, the root zone maintainer modifies the root zone, signs it with the zone-signing key and distributes it twice a day to root servers, to the operators, and each operator, themselves, have their distribution mechanism. Those are very quick, very quickly, throughout all of the thousand instances, worldwide.

It's very important to separate provision from resolution. So, provisioning means provisioning to what to put in the zone and what content is there. Those, that's in ICANN's sphere, governed

by the policy and process set. And then the resolution of that, the serving of that, is by the root server operators.

Over the years, a set of 11 principles has been articulated for the root server system. I think in 1998 when Postel died, the root server operators met in the IETF, where they developed a set of principles. One of them is IANA is a source of DNS root data, and over the years, throughout all personal experiences, they've kind of distilled these 11 principles. I won't go in to detail, but these are kind of the underpinnings for how, for the root server system.

As I mentioned earlier, the root server operators, 12 professional organizations, engineering groups focus on reliability and stability of the service and accessibility for all Internet users. As I mentioned, the root server system is really an organic growth along with the root servers. So, as services grow, more users and more instances are being added, topologically and geographically, to serve the global internet users. The root server operators, themselves, emphasize a technical corporation and professionalism.

One of the important features that Jon Postel set up for the beginning of the system is diversity and that diversity is an important feature for the root server system, and these diversities being expressed from the technical sense, what software, hardware the root server operators use to serve the zone,

organizationally, some are nonprofits, some are educational institutions, research labs, and some are for-profit institution;, geographically, the U.S., Europe, and Asia.

And, finally, funding models, different funding models. I think the diversity really is a very important feature. It really prevents capture by single organizations or entities for the performance of the service.

As I mentioned, they cooperate through meetings, communications, shared data. If you go to the root-servers.org, you can see a lot of these corporations in the news, the reports that they publish on major incidents, changes, periodic activity, support, emergency response capabilities. It's very important to know that the operators themselves are not involved in policymaking. So, the operators are publishers, not authors or editors of the zones. So, provision and serving, provision and distribution – those are very different.

Finally, to correct some myths. So, these are collected from kind of the various past tutorials that have been run, and the first myth is that the root service controls where the Internet traffic goes. The reality is really that the router, the routing system, controls where the Internet traffic goes.

The other myth is that most DNS queries are handled by a root server. The reality is that most DNS queries are not handled by a

root server. So, the TLD servers see a lot more traffic, some of them see a lot more traffic than root servers, and even for root servers, I think about maybe 60% of those queries are for non-existent TLDs or domains.

There's also a myth that the administration and service provision are the same thing. As I've tried to emphasize, again and again, those are different things.

Another myth is the root server identities have a special meaning. So many people say that the a-root-servers on that have special meaning. The reality is that none of the root server identities are special, right? One perception, myth, is that there are only 13 root servers, right? Like I mentioned, there are over a thousand servers globally served with 13 technical identifiers, served by 12 entities.

The root service operators conduct operations independently. The reality is they coordinate and, while they operate, they also collaborate and cooperate.

Last but not least, root service operators only receive the TLD portion of query. Like I said, the root service operators, they receive the entire query. There are standardization efforts ongoing in the IETF, where you will minimize part of the query. So, you don't need to see the whole query to ask for the top level.

We've got some questions from the implementers, from the ISPs, the root system and your network. How do I get faster connections and so forth? You want to have three or four nearby instances. So, I think that's the first thing, but I think with that it's also important to note that it's not the geographic, but there's the peering connection, and it's very important. Sometimes, you may have a root server set up next to you, but because of the peering policy, your traffic still goes abroad to another service. So, it's important to not only be nearby physically to the root service, but also to make sure you're increasing your peering connections.

We have this question often. I have a root server here set up next to my, why is the traffic not going there? So, check your peering arrangement. Secondly, is turn off the DNSSEC validation in your resolvers to ensure that you're getting a modified IANA data.

You can also participate and contribute to the RSSAC Caucus which is a body of experts where technical advice is developed and created for the root server system, so where that's, really, at ICANN it's what we call a bottom-up process. So, really, this technical advice is generated by these working groups in the caucus. So, we welcome you to participate and contribute in that.

Finally, if you're interested in hosting an Anycast instance, talk to an RSSAC member here after this presentation. You can also send an email to ask-rssac@iccan.org.

So, with that I hand it over to my colleague Carlos to talk about RSSAC and recent RSSAC activities. Thank you.

CARLOS REYES: Thank you, Steve. Hi, everyone, my name is Carlos Reyes. I work with Steve in support of the Root Server System Advisory Committee. So, briefly, talking about RSSAC, this is a direct quote from the ICANN bylaws which you can look up on the ICANN website, but the mission of the Root Server System Advisory Committee is to advise the ICANN board and the community on matters related to the root server system. You'll not that this is a very narrow scope. This is intentional, and the RSSAC is very diligent to ensure that any work that it takes on is within its remit.

So, to expand on that a little bit, what does the RSSAC do, and what does it not do? The RSSAC, of course, is an advisory committee, so its advice is for the board, but occasionally it does respond to requests from ICANN supporting organizations or even other requests. A good example of this is the ongoing policy development process in the Generic Names Supporting Organization on new gLTD subsequent procedures. That PPD has requested input from RSSAC on several occasions RSSAC will assess the question and provide advice. The root server operators, as organizations, are represented within RSSAC, but RSSAC itself as an advisory committee does not get involved in

the operations of the root server operators. So, there's that distinction there. Steve was alluding to it earlier, as well, in that the operators don't develop policies, and they also don't give advice, but that's the role of RSSAC.

So, this is in infographic that gives you a sense of where RSSAC fits in with the other advisor committees and the three supporting organizations in the ICANN multi-stakeholder model. So, let's talk a little bit about how the RSSAC is organized. There are 12 root server operator organizations. Each organization has a representative and an alternate, and there are also liaisons. I'll expand a little bit on the liaisons here in a few slides.

Steve also started talking about the RSSAC Caucus. This is a newer group within RSSAC. It came into being in, or it was founded in 2014, and it's a body of subject matter experts, DNS experts, and the primary purpose of the caucus is to conduct research and produce reports to inform the work of the RSSAC. We'll talk a little bit about work parties later as well, but I think what you can take away at this point is that the caucus is really a vehicle to bring in more technical expertise into the work of RSSAC, beyond the root server operator organizations and the root zone management partners. Every caucus member has a statement of interest that they submit, and all of this is available for everyone to view.

The current co-chairs of RSSAC. The RSSAC operational procedures established by RSSAC will be led by two coaches, and this is also in the ICANN bylaws, and the coaches are currently Brad Verd from Verisign and Fred Baker from ISC.

Liaisons, I mentioned liaisons a few minutes ago. There are two types of liaisons, inward liaisons and outgoing liaisons. The inward liaisons are primarily from organizations that are involved in the management of the root zones, so we have obviously the root zone maintainers. Steve was talking about that earlier; he IANA functions operator, which is currently PTI, the Internet Architecture Board, and the Security and Stability Advisory Committee. That's an ICANN group that has appointed a liaison to RSSAC as well.

In terms of outward liaisons, the RSSAC appoints a non-voting liaison to the ICANN board as well as a non-voting liaison to the ICANN Nominating Committee. One of the outcomes of the IANA stewardship transition, there were two new groups that were set up: the Customer Standing Committee which monitors the performance of PTI and the Root Zone Evaluation Review Committee. So, the RSSAC has liaisons to those groups as well.

A little bit more about the caucus. We have over 100 members now. Like I said, they all have public statements of interest and in RSSAC publications, every caucus member who contributes to

that work does receive credit. So, that is everything from participating in work party calls to drafting and editing the documents. It's a collaborative effort that then goes on to the RSSAC for review and approval.

Again, just expanding on the purpose of the caucus, there are DNS experts, but also that have broader technical experience, so that expands the base of input into RSSAC advice and the statements of the work and the credit that goes to the caucus members is part of transparency efforts to ensure that there's accountability.

If you're interested in applying for the caucus, there's a membership committee that reviews statements of interest and then makes recommendations to the RSSAC so feel free to email rssac-membership.iccan.org, and the membership committee can answer any questions or review a statement of interest when you're ready to submit that.

Current work parties. There are three work parties underway in the caucus. One is looking at service coverage of the root service system. Another is studying modern resolver behavior and the most recent work party is looking at metrics for the root server system. Every work party that begins within the caucus is really led by the caucus. So, they propose a work item, they draft a statement of work, the RSSAC approves it, and then the caucus members constitute the party as members, and then they also

elect the work party leader. So, the RSSAC is involved throughout the process. Every RSSAC member is also a caucus member, but the work parties are really where the bulk of the work and the activity happens within RSSAC and that's, again, another testament to the bottom-up model.

Transparency efforts. Steve has touched on some of these, so I won't really go into detail on most of these, but I'll focus on the RSSAC side. In terms of actions or steps that RSSAC has done to promote transparency, obviously, we have an RSSAC website. We have the caucus, the RSSAC publishes minutes from its various meetings, as well as reports from any workshops that it conducts. There is a calendar for both RSSAC meetings and caucus meetings and you're welcome to subscribe to that. It's a Google calendar. RSSAC conducts public meetings. That will happen later this week on Wednesday here at ICANN 64. RSSAC also meets with other ICANN community groups. There are few meetings here. Obviously, it meets with the board at every ICANN meeting and occasionally it's invited by other groups to brief it as well.

Tutorials, that's another point of transparency. Liaison relationships which we discussed earlier and the operational procedures which are also published. And that's actually a good point about publications. Every RSSAC report is published and you can go to the RSSAC webpage and find the publications

library and you can sort through all those publications. We're up to, I think, 41.

The RSOs, Steve covered most of these earlier. I will note that RSSAC 002, that is and advisory that the RSSAC published which sets basically a standard of measurements to establish a baseline of trends for the root server system. In that advisory, RSSAC advises or recommends that the operators implement these measurements and then track them. So, every RSO does publish those statistics.

So, questions? Of course, I mentioned the RSSAC website. There's an FAQ document there if you have some, again, frequently asked questions that the RSSAC members have received via tutorials or elsewhere. They do track that there, and it often answers a lot of the initial questions that people may have about the RSSAC or root server operations, and then there's also a webpage for the caucus and email addresses where you can send queries either to RSSAC or the membership committee.

At this point, I'll stop here. We do have members of RSSAC in the audience to take any questions and thank you very much for attending.

STEVE SHENG: So, we'll open up the queue. Are there any questions in the audience? Go ahead. Do you want to come up here and use this microphone? Thanks. Please identify yourself first before asking the question, thank you..

JOHN: Hi, my name is John from [inaudible]. I'm curious about … I have two questions actually. The first one is about the caucus. Who are the members of the caucus? Is it the personnel, or just another, and is this free volunteers, or at all? Is it free volunteers participating for the caucus?

The second is, I know that this is a question, and the question is about how are the 13 operators maintaining the root server? What is the business model, actually? Because I'm just curious about how you're maintaining the root server and how they can get some revenue stream, as the operators here, I mean, as the 13 operators?

The third question is about we'd like to have an opportunity, if we can put also one of the root servers on or site. Thank you.

STEVE SHENG: Okay, so three questions. First one caucus, second business model for the root server operator, and the last one is how do I get Anycast instance? So, on the caucus model, I can answer. These are volunteer technical experts, and you submit a statement of interest. Every RSSAC member is a caucus member but we have other technical experts who do not work for the root server operators in the caucus. So, and there's an email address.

UNIDENTIFIED MALE: [of mic].

STEVE SHENG: It's volunteers, like any other ICANN SOs and ACs, so the members are really volunteers. So, the second question was regarding funding. the revenue model, sorry, the revenue model for the root server operators. So, any … Liman?

LARS-JOHAN LIMAN: Thank you. My name is Lars-Johan Liman. I work for Netnod and we operate one of the server constellations, and I'll walk up here so that you can all see me.

One of the aspects of the root server operators is that we are all different, and that is one of the strengths of the systems. One of the models that we operate by is actually diversity is good, and

one of the things that we are very diverse about is how we generate the money that we need to operate the service. To begin with, this is nothing that generates income. There's no root server operator that gets money for operating the root from central source somehow, so we have to find a way to make sure that we have the money that we need to operate the service.

I can only respond for Netnod. I can tell you what we do, but you have to ask all the others individually how they do it. In the case of Netnod, we provide DNS service, not only for the root, but we also provide DNS service for, for instance, top-level domains. We have something like 50 top-level domains that we also run on servers in the same compounds. We have the root server and then we have other servers where we sell DNS service, and there we can charge. So, we charge money for the top-level domain registries and from other people who want DNS service as well, and from that money we take a little of that and we put that on the operation of the root servers, and that creates, in our case, we think, a win-win situation because the top-level domains, they wouldn't be able to operate unless there was the root service, so they see that, okay, if you pay for this, they can also play for the root servers and we need that, and so that is good, and for us it's good because we can use a little of that money to operate the root server.

So, that is our business model, but then there are 11 more root server operators and they all have different models, and that's good because if our model fails, I'm quite sure that the model for the Internet Systems Consortium, or the model for Verisign, which is different, it continues to work. So, this is actually part of the stability package.

STEVE SHENG: Thank you, Liman. And to your third question, I guess you can come and talk to the RSSAC members up front later about Anycast instances in your area. Any other questions? I will take the microphone, so less walking for you, no? Yes? Yes.

[NATALIA]: Hi, my name is [Natalia] [inaudible], ICANN Org. Just to follow up on the question on setting up an instance of any of the root servers, is there is an organization that is an operator who is interested to set up an instance? Are there any kind of criteria or differences on which root server from the 13 that the operator wants to set up, like K or L, or are they all equal? How do you pick from which server do you want to get an instance, and can the root server operator refuse a specific organization to set up an instance? Were there any cases when you just refused an organization to set up an instance of your root server? Thank you.

BRAD VERD: Brad Verd, Verisign, co-chair of RSSAC. I'll let Fred talk after me. So, the instances. There are a number of root server organizations. Oh, am I in the way? I'm sorry. So, there are a number of root server organizations that offer – is the mic not working? There, is that better? So, a number of root server organizations.

UNIDENTIFIED MALE: Let's trade.

BRAD VERD: How is that, better? Alright. So, there are a number of root server organizations that have their own process for providing instances to people who are requesting them and would like to host one, and RSSAC is happy to take your request and forward it on to the root server operators, or you could come talk to them. There are a number of them here. You could talk to them individually. Each of the RSOs, the Root ServerOperators, have their own set of criteria for where an instance would go or could go. As far as refusing service, I'll speak for Verisign, yes, if you come to Verisign and ask for a root server operator and it's going to go in the basement of your house, you're probably not going to get it. It's a large investment on Verisign. We're spending numbers on

servers, when you basically have to provide space, power and ping, and we want the root service to serve the Internet, so it needs to go in a place that is justifiable for that service, so it's not just anywhere.

So, I can't speak for the other root server operators, but maybe somebody else wants to say something. But if you are looking for an instance, there are a number of different ways to get one and I'm happy to have that conversation with you.

FRED BAKER: I'll pick a microphone that works. Our situation is really very similar. Coming back to your question about how we're funded, ISC is an open source software developer and also operates one of the root services. So, now open source software is an interesting problem in its own right. I'm giving something away for free. How do I derive money from that? And what that usually means is that we sell support for that software. And, well, yeah, so some of the money that comes in gets used to operate a root service, and nobody is paying us for the root service, but some of the money that we get from various sources is used to operate the root service.

Now, if you go to ISC.org, one of the things that you'll find is, “I’d like to have my own instance. Could you please send me one?” And they'll ask a number of questions. Electricity needs to come

from somewhere. We need to have connectivity peering, and so on and so forth, so there's a dialogue. There's actually a set of questions right up front, and then there's a dialogue that happens, and then if we wind up feeling that you're a viable candidate, then fine, we'll mail you the relevant [information].

One fairly interesting problem that we ran into recently, we were talking with an organization that wanted to host an instance, and in the course of it they were really being fairly demanding about certain things, and eventually they phrased it to us as, if we're going to offer you the service of hosting your instance, this needs to be true. Understand, they're not offering us a service. We're offering them a service. And when we got that clear with them, a whole lot of the other questions just kind of disappeared. So, we're doing this for the good of the Internet. It's a public good. And the dialogue is basically with the view to making sure that it, in fact, does that, and that it's a public good.

STEVE SHEG: Thank you, Brad and Fred. Any other questions? Oh, I'm sorry, the gentleman in the back.

UNIDENTIFIED MALE: Thank you. My name is [inaudible] from dot-ng. Since it assumed or said that all the root servers, about 13 of them, have the

content the same data … Now, how are the clients, how do you load the traffic from the clients, distributed so that the traffic will not go through a specific, maybe half or [inaudible]. So, how does the traffic, the queries from the client, distributed fairly over the 13 root servers?

STEVE CHANG: So, he's asking how do the different root servers receive the traffic, I think?

FRED BAKER: Well, the fact is that the root servers don't distribute the traffic. The computer that is asking the question has 26 different addresses to pick from, and 13 of them are IPv6 and 13 of them are IPv4, and he really doesn't know the difference, but he might know something about his history with them, that this one seems to respond more quickly because he's closer, this one seems to respond more slowly, and that kind of thing. And so that computer is going to decide on its own which addresses, or which address, that it wants to use on any particular access. And the way the protocol is set up, we ask the guy to try all of them over a period of time because we want to make sure that he has that information. But, at the end of the day, he's going to pick, he's going to optimize for himself, and so he's going to pick the ones that make sense.

Now, what happens with those requests, when we talk about Anycast, what we are saying is that we have a cluster of computers – or cluster is the wrong word. We have a set of computers that are all using the same address, and usually they're in different locations. And what that means is that if the packet gets sent to that address and goes along this route, it's going to a particular server and that server will respond. If that server goes down, bad things happen, then routing will now take that packet to another server somewhere else that's using the same address. So, the routing changes, and it goes to another server that happens to be working at the moment.

And so, the distribution of traffic is somewhat random. It has to with routing and how routing distributes traffic, but we have routing set up in such a way that we achieve the distribution we're looking for. Does that answer your question? Okay.

STEVE SHENG: Thank you, Fred. Any other questions?

UNIDENTIFIED MALE: Okay, I have a question about we have seen that their mechanism to ensure [inaudible] will be correctly distributed with the operator, but for one operator, for one of the 13 operator was, in

theory, the mechanism [inaudible] to modify the data it, themselves, yeah?

STEVE SHENG: So, how so you ensure the operator doesn't modify the data? Is that the question?

UNIDENTIFIED MALE: Okay, alright, Liman is raising his hand.

LARS-JOHAN LIMAN: I'll walk up to the front. To be quite frank, you cannot ensure that the root server doesn't modify the data, but you will immediately notice. And for me, that would be, if I was to do that, it would be an immediate disaster because a lot of people will notice that, and I would no longer be trustworthy, and I will be tossed out of the business, immediately. So, for business reasons, I cannot do that.

And the way to notice if I have have made any changes to the file, is DNSSEC. Security DNS will tell you immediately if I have made modifications to the data that you receive from me. It's properly signed with keys that I don't have. Only the proper root zone maintainer and administrator and maintainer channel has the keys that are needed to generate the signatures. I don't have

access to those keys. They are very carefully locked into safes that I don't have access to. So, I cannot modify the signatures. I can possibly modify the data, it would kill my business immediately, and you would notice immediately. So, there is a security thing in there that is very easy to use.

STEVE SHENG: Thank you, Liman. Any other questions? No. Any questions from this side? No. This side? Well, thank you very much. Could I ask the RSSAC members and the root server operators to stand up so that, later on, there may be conversations and discussions. So, thank them for providing the Internet, good service. Thank you.

CATHY PETERSEN: Thank you everyone. The slide materials are already uploaded in the public schedule, and the link to the recording will also be added to the public schedule within the week. Thank you.

**[END OF TRANSCRIPTION]**