KOBE – How It Works: Root Server Operations
Saturday, March 09, 2019 – 15:15 to 16:45 JST
ICANN64 | Kobe, Japan

CATHY PETERSEN: Hello, everyone. We should be starting in a couple of minutes. Thank you. Welcome, everyone. Good afternoon. Welcome to our third How It Works tutorial for today. This session we'll be talking about root server operations and we have two presenters. The first one is Andrew McConachie from ICANN and he will start now. Take it away.

ANDREW MCCONACHIE: Thanks, Cathy. Okay. My name is Andre McConachie and I'll be presenting a tutorial on the Root server System. There are four sections here I'm going to talk about. I'll be handling the first three, the overview of the DNS, a quick explanation of Anycast and how it works, and then the Root server System today and its features. Then, after that, I'm going to hand it over to my colleague, Carlos Reyes, who will talk about the RSSAC and some recent RSSAC activities.

So, this could be a bit of a review for a lot of folks, but it's good to just kind of get this as a baseline. So, I'm going to begin with an overview of the domain name system.

So, a quick recap on identifiers in the Internet. The main identifier is, of course, IP addresses, the fundamental identifier in the Internet. They're numerical which kind of makes them hard to remember which is one of the reasons why we have DNS. And all [hosts] connected to the Internet need to have IP addresses, and we have two versions – Ipv4 and IPv6.

So, why DNS? Well, there were two original problems. One I already mentioned, that IP addresses are hard to remember. The other one is that IP addresses change a lot. We also have some more modern problems of [why] DNS and that's IP addresses might also be shared – in the case of [NAT], for instance. And there's kind of like a many-to-many or one-to-many or many-to-one type relationships that you can have with IP addresses. So, you can have multiple IP addresses that may serve as entry points to a single service, and which one do you use? The clicker is no longer clicking. Did I turn the clicker off? Ah, there we go. The mouse moved. Is the mouse pointer over … Oh, it looks like we have a frozen computer. I'll just say next slide. Next slide.

So, the domain name system is hierarchal. At the top, we have a root and then, beneath that, we have top-level domains also called TLDs, a very common word here at ICANN meetings. Beneath that, we have second level and then third level and on and on.

Here you have an example of three pretty common top-level domains – UK, dot-org, and edu – and the kind of labels that can be beneath them.

There are many other mappings besides just what we call A or Quad A records. We've got mappings for mail servers. There's also IPv6 and reverse DNS records. Next slide, please.

This slide is pointing out an important difference, kind of when we talk about the root server operations as well as the Root Zone and that there's a key difference between the data and serving of the data and that's what this slide is meant to show.

On the left, we have the Root Zone which can be thought of as the data and then the Root server System which can be thought about the system that serves the data. So, the Root Zone is really the starting point. It's the list of TLDs and their Name Servers and how to go about finding those Name Servers, the records that point to those Name Servers. It's managed by ICANN per community policy. It's compiled and distributed by the Root Zone maintainer. As I previously said, this is the database content for the Root servers.

The Root server System, on the other hand, is the system that responds with the data for the Root Zone. It's a highly distributed system with 13 identities and right now over 1,000 Anycast instances at different physical locations around the world. And

the Root server System is a responsibility to the root serve operators. Next slide.

Okay. So, here are some definitions. The Root server System that I already spoke about is the set of toot servers that collectively implements the root service. That might sound a little bit, like it's repeating, but the Root server System is really the Root servers that implement the Root servers. The Root Zone, as already stated, is the data – it's the DNS zone at the top of the DNS hierarchy. It's also the zone sometimes referred to as having no parent, and it contains all the information necessary to contact the TLDs.

Root server Anycast Instance are – these are considered like one network location which is responding to DNS queries on a Root server's IP address – on a Root server Operator's IP address. Next slide.

These are some roles – these are some definitions of some roles. The Root Zone Administrator is the organization responsible for managing the data contained within the root zone. This involves assigning the operators the top-level domains and maintaining their technical and administrative details, and I'll go into a bit later about how those details get updated and the whole process flow for that. The Root Zone Maintainer is the organization responsible for accepting service data from the Root Zone

Administrator and they format it into a zone file format, cryptographically sign it and distribute it to the root server operators.

And then finally we have the root server operators which are the organizations that actually run the root servers and are responsible for managing the root service on the IP addresses.

So, the root servers only know what servers need to be [asked next], so the root zone only contains the TLDs and the name servers in order to reach those TLDs, and when the root servers are asked where is this query, they only respond with the NS records for the next zone down, so the TLD zone.

I'll go more into the DNS resolution process in a bit, but basically caching at the recursive resolvers means that there's very … Most DNS queries do not actually hit the root.

Some modern refinements to DNS include DNSSEC, which you'll hear a lot about here at ICANN, and what DNSSEC does is it provides cryptographic signatures on DNS data and so there is a signing process where the DNS data is signed and then on the recursive resolvers, it can then be validated.

There are some modern privacy enhancements to DNS, because originally queries all happened in plain text and went over the wire visible to anyone. There are some recent enhancements,

specifically DNS over TLS, also sometimes called DOT which provides a secure channel for DNS queries, and then Anycast is very important for scaling the root server system and that allows multiple servers to share a single IP address and it greatly improves latency and resilience and protects against distributed denial of service attacks.

So this is the slide where I'm going to spend a bit of time on and hopefully tie everything together that I just talked about, so I'm going to run through what happens with a DNS lookup by a user on a computer and how that gets resolved, how it involves a root server system, how DNSSEC plays into this, and finally how the user is able to get to a resource on the internet – in this case, the Web.

So, we start with our user on the right-hand side and they would like to go to example.com, to the Web server there, and there is an assumption we're going to make with this slide, and the assumption here is that the recursive name server has never done a query before, or it's just been turned on and it knows absolutely nothing. In reality the Recursive Name Server that's immediately to the left of the user, probably will have already done some queries and probably will have already had its cache filled with a lot of information, so it doesn't have to go through this whole process, but we're going to start from the beginning and show the whole process.

So, the user types in www.example.com in their Web browser, and that triggers a DNS request from the user's computer to the recursive server there in the middle, and the recursive server, not knowing anything except for the addresses of the root servers, will then query a root name server, and it sends the entire query typically. Now the root name server will then respond with, "Well, I don't know exactly where www.example.com is, but I do know where dot-com is." So, it will send back the name server information for dot-com. And then what the recursive server does at that point is it validates the DNSSEC signature of the dot-com name information. It says, "Okay, this is valid, this is signed, this is all valid, this is all good, so now I'll go to dot-com?" nd he says, "Where is www.example.com?"

And once again, the com server doesn't know where www.example.com is but it does know where example.com is, so then it will send the name server information back to the recursive name server, and the recursive name server will do that once again – it will validate the addresses of example.com and then it will go to example.com and it will say where is www.example.com, and this time, it will get a response of where www.example.com is, and it will finally return that answer back to the user and then the user can go to that web page.

So that was a brief, rushed overview of how DNS works. I'm now going to go into an explanation of Anycast.

So, the original problem – problem is maybe not the right way to put it, but originally before there was Anycast, there was just Unicast, and Unicast means packaged from sources all go to the same destination. There was just one physical machine living at one IP address and it receives all the packets destined to that IP address.

This isn't very good for DDoS traffic because it's difficult to scale when you have a lot of hosts DDoSing that single instance. It also doesn't mean that you can't geographically distribute your servers as well.

So, in steps Anycast, and what Anycast allows us to do is you can have multiple instances serve the same data to all sources and the sources reach different destinations based on intermediate routing policies, and the idea is that the sources can get the data faster through lower latencies and that also the DDoS attack traffic is sent to the closest instance, thereby being distributed so it shouldn't overwhelm a single server.

This is a graphical representation of Unicast where the traffic takes the shortest rout to the single destination. The idea with the circles in the middle is that they are making the routing decisions for where the traffic should go and they're just in one destination that they can chose from so they forward the traffic onto that destination.

This is an example, a graphical representation of Anycast and here we see we have multiple destinations, represented by the blue, and we can see that the traffic takes the shortest route to the closest destination. The intermediate nodes can determine hopefully what the closest destination is and go to that via the shortest route.

And here we see one of the main purposes of Anycast is to help distribute DDoS load, so here we have a source in the bottom left, and a DDoS attacker in red, in the top right, and they're going to different destinations. One of them in this graphic is acting as a DDoS sync so it just receives traffic and syncs it so it doesn't affect the traffic, the more legitimate traffic going to a different source, or coming from a different source and going to a different destination.

Now I'm going to talk a bit about the root server system and some root server operators. So, this slide shows a bit of history. This is the growth of the root server system. You can see 1983-1986 there were just four addresses, and then as you progress through time you see the number of addresses increase up until 1998, which is – it hasn't changed since then. And the changes over time were mainly driven by technical demands and scaling issues, although scaling these days is now much better solved using Anycast.

Like I said earlier, there are over 1000 instances spread around the world right now using Anycast. There are still 13 addresses, although when IPV6 was added there were 13 IPv6 addresses added as well, so now you can see that there are 26.

And these are the different operators and their IPV4 and IPV6 addresses. What graphic are we missing there? Okay, we will just skip that.

So, this slide shows how the root zone is provisioned and how distribution resolution work. On the far, on the far left side, we have TLD operators. So, these are operators for TLDs that are in the root zone and they maintain their own records in the root zone. For instance, the records for their NS servers are possibly glue A records. And when they need to change them they contact IANA function, and it's the IANA's job to make sure that they're communicating with the correct TLD operator and everything is above board there.

And then they pass that information on to the Root Zone Maintainer. The Root Zone Maintainer actually compiles that into what's called the root zone and then distributes it to the different operators, which are marked here as RS, RS, RS. All the way down the right side, you see the very small bubbles with RS in them. Those are the individual Anycast instances. And then from there

queries come from recursive resolvers located around the Internet to the root servers.

So, recently the RSSAC published RSSAC 37, which is a very important and long document for the RSSAC, and in it there are 11 principles of the root server system, and I encourage you to go read that document if you're interested more in learning about the RSSAC and the evolution of the RSSAC and learning more about these principles. I don't want to read them because I'm losing my voice and you can read the slide.

But a bit more about root server operators, there are 12 different root server operators, and these are professional engineering groups focused on reliability and stability of the service and accessibility for all Internet users, cooperating both technically and professionally, and they're a diverse group of organizations and operations. They're diverse and there's four different ways technically, organizationally and geographically and different funding models.

With regards to cooperation and coordination, these are some of the meetings that they cooperate through. They've also got communication tools that they can use in case of emergency situations, so back-up methods for communication in case something – there's some kind of emergency that needs to be

dealt with. They do share data, and they also have periodic activities for emergency response.

Some of the things that operators are involved in is careful operational evolution of service, so as the service evolves the operators involved in working with that, evaluating, deploying different types of technical modifications that need to come to the service, like as new things come about such as DNSSEC. And in general, to make sure every effort, to ensure stability robustness and reachability in the root server system. Operators are not involved in policy making, and they're certainly not involved in data modification. The root server operators, they're strictly publishers and they're not authors or editors. They simply publish what they get from IANA.

These are some of the myths of both I guess the root server operators and the root server system, and we have on the left-hand side we have the myth, and then on the right-hand side we have the reality. An old myth is that the root servers control where Internet traffic goes, and the reality is that the routers do that. Root servers might control what the … They serve an address. They answer a query, but it's ultimately the router's control where that traffic ends up.

Another myth that I talked about previously was that most DNS queries are handled by a root server, and the reality is that simply

because of caching and recursive resolvers, most DNS queries are not handled by a root server and instead they're handled in the recursive resolver.

Administration of the root zone and service provision are not the same thing, so in the graph where I showed how a TLD operator can update their information and how that works its way through the system in order to get all the way to the recursives, the administration of the root zone is something that is handled by the IANA function, and it's separate from the service provisioning and the root server system.

Another old myth is that roots or identities have special meaning, and none of the root server identities are really any different than the others – none of them are, so to say, special.

Another myth is that there are only 13 root servers at one point, but now there are so many Anycast instances. And I have to say it's more than a thousand, because I don't … It's changing right. It's always growing, so it changed – I think when I first started doing this it was like more than 700, and then at some point we changed it and it said more than 800, and now it's up to more than a thousand. I guess at some point I'll have to say more than 1100 but right now it's still more than a thousand.

The root server operators do cooperate. They cooperate with one another to provide better service.

A final myth is that the root server operators only receive the TLD portion of the query, and there's a caveat here because this is maybe an older slide, the old reality use to just say root server operators received the entire query. However, there is a new standard called [queue name anonymization] which isn't that much deployed, and this is something that came out of the IETF and is meant to address some issues with privacy in the DNS. And with [queue name anonymization], actually root server operators will only receive the TLD portion, but it's just not all that well deployed yet so that's why the word 'usually' is there bolded. That is changing.

So, if you're a network operator and especially if you are on a recursive resolver, these are some things that you can keep in mind for the root server system in your networks. Generally, you'll want to have three to four nearby instances, and by nearby it might mean direct peers, like BGP peers. It might mean … Basically you want them to be close and you maybe wanting to have low latency connections to them.

Turn on DNSSEC validation in resolvers. This ensures that you're getting unmodified IANA data and you're not being lied to either via some kind of on-path attack or some other type of attack. DNSSEC validation will prevent DNS lies pretty much.

You can participate in and contribute to the RSSAC Caucus if you're interested. My colleague, Carlos, is going to talk a bit more about the RSSAC and the RSSAC Caucus. But basically, it's where technical advice is created, and if you're interested in the root server system and you're interested in the evolution of the root servers then I encourage you to join the RSSAC Caucus and contribute there.

Finally, if you're interested in hosting an Anycast instance of a root server, you can talk to an RSSAC member after this presentation, or send a mail to ask-RSSAC@icann.org and after Carlos and I are done presenting, there will be a question and answer period and the answers will be handled by root server operators, and after this meeting if you are interested in hosting an Anycast instance you can always talk to one of them. So now I'm going to hand it over to my colleague, Carlos Reyes, who will talk about RSSAC.

CARLOS REYES: Thanks Andrew. Hi, everyone. My name is Carlos Reyes. I work for the Policy Development Support Department here at ICANN and I'm assigned to support the Root Server System Advisory Committee and I'll be providing an overview of the RSSAC today.

So, what is RSSAC? RSSAC is the Root Server System Advisory Committee as I mentioned, and its mission is to advice the ICANN

Board and Community on matters relating to the operation, administration, security, and integrity of the root server system. This is directly pulled from the ICANN bylaws. The ICANN bylaws spell out the Charter of RSSAC.

You'll note that this is a very narrow scope and that's by design and we'll cover that here. So the RSSAC, as the name implies, is a Committee that produces advice. That advice generally goes to the board but it can also produce advice to other ICANN supporting organizations. Fr example, the Generic Names Supporting Organization which is the supporting organization that develops policies for generic top-level domains, has requested input from our second in the past so if there is an ongoing policy development process there can be requests for RSSAC to provide advice.

There was a previous slide where Andrew was talking about how the operators do not develop policy, so the converse is that RSSAC does not involve itself in the operations of and the operational work of the RSOs, so I think that's an important distinction here is that RSSAC is the group within the ICANN multi-stakeholder model that develops this advice and the operators continue to perform their service.

So, this is just a graphic representation of where RSSAC falls within the ICANN multi-stakeholder model. We have the three

supporting organizations that develop policies and the four Advisory Committees, and RSSAC is one of those Advisory Committees.

So, let's talk a little bit about the organization of RSSAC. RSSAC is composed of representatives from the Root Server Operators. So, there are 12 Root Server Operators where there are 12 representatives. Each Root Server Operator also has an alternate, so we have 24 representatives from Root Server Operators, and then there are liaisons. We'll talk a little bit about the liaisons in an upcoming slide.

The RSSAC, much like every other group within ICANN, undergoes periodic reviews. Those have been roughly every five years. The first review in the 2007-2008 timeframe produced a recommendation that the RSSAC should be expanded, so in 2013 the bylaws were amended and the new RSSAC also created and established the RSSAC Caucus. The RSSAC Caucus came into effect in 2014 and it's a body of volunteer experts – DNS experts – that broaden the technical base available for RSSAC work. I'll talk a little bit about the RSSAC Caucus in another slide here, but the members of the RSSAC Caucus are confirmed by the RSSAC after a recommendation from the Membership Committee, so statements of interest and applications are reviewed by the Membership Committee and then they're approved by RSSAC.

Current co-chairs. The ICANN bylaws established two co-chairs for the RSSAC. The current co-chairs are Brad Verd from Verisign and Fred Baker from ISC.

So, earlier I mentioned liaisons. There are two types of liaisons. There are liaisons that are appointed from other organisations involved in the management of the root zone to RSSAC, and then there are liaisons from RSSAC to other groups. So, the Inward liaisons – there's the IANA Functions Operator, the Root Zone Maintainer and the Internet Architecture Board. These are of course outside of the ICANN structure.

And then within ICANN, the Security and Stability Adviser Committee (the SSAC) has a liaison to RSSAC. Then the RSSAC appoints Liaisons to the ICANN Board. That liaison is a non-voting member of the ICANN Board. The ICANN Nominating Committee, also a non-voting liaison, and then the Customer Standing Committee and the Root Zone Evolution Review Committee. These are two new groups that resulted from the IANA stewardship transition.

Again, the RSSAC Caucus. Currently there are 106 members. That's as of last month. Every member has a public statement of work – statement of interest, excuse me – that documents their qualifications or background and their motivation for joining the caucus. And every caucus member that contributes to an RSSAC

publication or an RSSAC work party receives credit for that contribution.

So, the purpose of the caucus as I mentioned earlier, it's to bring diverse technical expertise to RSSAC work and one of the principles in the operational procedures of RSSAC is to ensure that there is transparency of who is doing the work, so that refers to the credit earlier. And also the caucus provides a process and a framework for how work parties are created within RSSAC, how they operate, and then how they deliver their products to the RSSAC for review.

If you're interested in joining the RSSAC caucus, feel free to approach me, or any of my colleagues, or any of the RSSAC members, and you can also submit an email to RSSAC-membership@icann.org.

So, current work parties. There are three work parties underway within the caucus. This is how the RSSAC and the caucus organize their work efforts. The first work party here is service covered to the root server system, and the second work party here is studying modern resolver behavior, and the third work party is the newest, which is looking at metrics for the root server system.

Every work party has a shepherd, which means that someone from the RSSAC is monitoring the work and reporting back to the RSSAC. Every work party also has a leader, so the caucus work

party will elect a leader to, obviously, lead the work party, and guide their work and agenda. I think all the shepherds for the various work parties are here in the room, so again, feel free to each out to me and I can direct you if you'd like to talk to them about any particular work effort.

Transparency is very important to the RSSAC. It's also a principle that is taken directly from the ICANN bylaws. So, on this slide we capture some of the transparency efforts both by RSSAC and the RSOs. So, I'll go over the RSSAC list first.

RSSAC has a web page that's part of the ICANN.org website. Obviously, we have the caucus which I mentioned earlier, which broadens the base of technical expertise available for RSSAC work. RSSAC publishes minutes from all of its monthly meetings, and also publishes reports from its workshops. The RSSAC conducts workshops – one or two a year, depending on its workload. Public RSSAC and Caucus calendar, that is maintained on a Google calendar, so if you'd like to subscribe and know when work party meetings are happening, that's one way to do that.

The RSSAC conducts public meetings. Here at ICANN 64 that will happen on Wednesday, so you'll see the RSSAC conduct one of its monthly meetings.

The RSSAC also meets with various community groups. Those are typically open unless the other group requests a closed meeting.

You'll see the RSSAC meet with the ICANN Board and the Board Technical Committee here. There will also be meetings with the Office of the Chief Technology Officer of ICANN Org and the RSSAC will also be briefing the At-Large Advisory Committee, so there's a lot of interaction. These tutorials, that's another example of transparency. Obviously, the liaison relationships ensure that there's communication between RSSAC and other groups.

And then the operational procedures. This is essentially how RSSAC conducts its work, how it operates day to day, and the RSSAC reviews that and revises that document annually as well.

The RSOs. Similarly, there's an RSO website (root.servers.org). Agendas from root ops meetings are published.

RSSAC 002 statistics. A few years ago, RSSAC produced a publication, which is RSSAC 002, and those statistics, every RSO publishes the statistics on their website. Obviously, RSOs participate in RSSAC. Every RSO also has individual web pages and occasionally – I think Andrew eluded to this – the RSOs produce reports or collaborative reports on major events related to the root server system.

A few years ago, RSSAC also agreed to respond, or at least convey, questions that are communicated to RSSAC back to the RSOs, so that's another mechanism there where the RSOs can respond to questions that the RSSAC receives within ICANN, and then of

course Andrew mentioned this e-mail address earlier. We have ask-rssac@icann.org if you have any questions. That is monitored.

So, this is my last slide. Again, we have information about the RSSAC here. We have an FAQ document that the RSSAC regularly maintains. If you have any general questions ask-rssac@icann.org and then we have information about the Caucus as well.

So, I'll pause here. I'll invite members of the RSSAC to come to the table here, and if you have questions about RSO operations or the RSSAC feel free to line up and we'll proceed through those questions. Thank you very much.

CATHY PETERSEN: Please feel free to use the microphone in the center of the room, and just provide your name and affiliation please.

JOAN KATAMBI: Thank you. Joan Katambi is my name. A fellow, from the Fellowship Program. I want to ask. Andrew presented and he talked about the Anycast. I'm wondering because I'm from the academia, so if we are talking about communication models, we have the Unicast, the Multicast and also the Broadcast, so I want to know is [End] cast also coming up as a new communication

model and to also understand how it's actually getting different from the three. Thank you.

[FRED BAKER]: I'll take a crack at that. So, Multicast and Broadcast and Unicast actually have different kinds of addresses, so in the network if I send a broadcast on an ethernet for example, all of the systems on the ethernet receive it. If I send a multicast, then all of the systems that are interested in that particular address receive it. There may be other systems, but [inaudible] of the systems on the LAN will receive it. If I send a Unicast it goes to the computer that I'm talking with.

Anycast actually uses Unicast addresses and operates very much like Unicast. What's different is that I have multiple systems – might be two, might be 2000 systems – that are using the same address and the packet that I send will go to whichever one routing thinks is nearest to me. Now, routing, there are different routing technologies. They may come up with different answers. It's not like it'll go to some specific one but it'll go to whichever one routing thinks is closest to me.

BRAD VERD: I'll add to that. As Anycast has been around for a long, long time, it was first implemented on the root system in 2002 I think, or

2000, and then widely used in the root server system from 2004 on.

ANDREW MCCONACHIE: Any other questions?

ROB MONSTER: Hi, Rob Monster, I'm the CEO of Epik, epik.com, registrar. We also operate the Anonymised dot-com privacy proxy VPN service. I have a question about censorship, a rising pattern of censorship where ISPs or nations would in theory [know] route or misroute certain domains, and to what extent there is an endorsement or advocacy for being able to essentially pair VPNs with the authoritative root server, so that parties who would use a particular VPN would always know that they're being routed to the authoritative root. Did that make sense?

[FRED BAKER]: Well, I think your question makes sense. What I don't understand is what is an authoritative root is. The way the root works is that the Internet Assigned Number Authority collects a database of associations between character strings and addresses. That is handed off to the Root Zone Maintainer who collects it up in a file, and then the root zone system – the RSS that Carlos was talking about – has a thousand different root zone servers that distribute

that information out. No one of them is authoritative and they all have the same information. So, I'm not sure where that VPN would go.

UNIDENTIFIED MALE: Yeah. I guess what he meant, that question is more about how to basically address perhaps some of the challenges arise due to patterns of censorship but in fact, it may be ultimately that we're really solving different sets of issues, right? There was an issue that is basically dealing with the individual registries, and then there's the issue of dealing with the root.

[FRED BAKER]: Okay. So, there's at least two sides of that problem that we address. That is if you Google a document called RSSAC 001, that's a set of expectations that we as root operators sign up to. This is what we will do. And one of them says whatever IANA tells us, we will distribute. We will not change it. So, that would not be something that any of the RSOs would do.

Now, a resolver might do that. they might very possibly do that. Open DNS, for example, has a service in which they don't distribute names that would get you nowhere, and in your [parlance] that's censorship. I'm not sure I disagree with you. It's

whether it's good or bad. So, we're not going to edit the zone file. We deliver the zone file as it is.

Now, the other side of that is that when we deliver the zone file, we deliver it signed. It comes to us from the IANA signed. We deliver it signed. And so the resolver or whoever receives that has the capability of checking the signature and saying, "Aha, that was in fact what IANA sent," or, "No, it wasn't." And I think that would be an important thing for you to rely on.

Your next question is, I suppose, how did you get to a resolver that gives you assurances that you're looking for? And that could be done with a VPN. But I don't think there would actually be a lot of value in doing that for a root zone server.

ROB MONSTER: Okay, great, thanks.

BRAD VERD: Just to add that the root servers that are represented in the root zone, the 13 identities, they all serve the IANA root. The IANA root is signed, so any censorship that you refer to from a root server is not come in from one of the official root server identities. Any further questions?

JOAN KATAMBI: My question to you, as the RSSAC, currently in our country we face a lot of Internet shutdowns and we have an issue with the OTT tax so Internet is free. You must pay a tax even after paying for the [GBs] to be able to use. So, our people have resorted to using VPN, so I'm wondering is VPN good or bad. Thank you.

BRAD VERD: To be honest with you, that question is outside the scope of RSSAC and we're not in a position to answer that.

CATHY PETERSEN: Are there any other questions from the room?

Alright, I guess we will conclude this tutorial on root server operations. We will have another one tomorrow, here in the same room, at 3:15 if you want to come back. Meanwhile, thank you again for your attendance.

[END OF TRANSCRIPTION]