MONTREAL – How it Works: Root Server Operations
Tuesday, November 5, 2019 – 17:00 to 18:30 EDT
ICANN66 | Montréal, Canada

STEVE CONTE: Alright, we're at the final legs of the 'How it Works Sessions' and I appreciate all people who have joined us throughout the day and on our Sunday sessions as well. I always love this particular session, partially because I, way long ago, used to run a Root Server, L-Root Server, with my colleague, John Crane, but also I just love DNS, so this is one of my favorite sessions and it's evolved in the way that RSSAC handles it with the Root Server Operators is a really engaging way to do it.

So you guys, if you haven't been here before, you're in for a treat; have questions, ask questions, this whole session, this whole series is meant to be a dialogue, so I think Steve, the way that you're going to do it is you're going to do presentations and have question sections, or you want to hold questions until specific times; is that how you...?

STEVE SHENG: Maybe we'll do a quick presentation first and then reserve questions towards the end where the Rooster Operators can answer them.

STEVE CONTE: Okay, great. So with that, this is Steve Sheng, he works for ICANN, he supports the RSSAC and SSAC and many other aspects of ICANN policy

development and other aspects, so I'm going to pass it over to Steve; and please again, engage. This is a great time to speak to the RSSAC and speak to Root Server Operators. Thanks.

STEVE SHENG: Thank you, Steve, and good afternoon everyone, my name is Steve Sheng, with my colleague, Ozan Sahin, we'll be giving the tutorial on Root Server System. As Steve said, we're going to structure today in two parts; the first part is where we'll give the presentation and then the second part, the Q and A session, we'll invite the Root Server Operators here to give the answers where Ozan and I will moderate the queue.

So, here's a brief outline of the tutorial today. I will cover the first three topics and my colleague, Ozan, will cover the last two. We start with an overview of the Domain Name System because, you know, historically speaking, the development of the Root Server System is intricately linked and aligned with the development of the Domain Name System. Then I'll move into a brief -- give a brief explanation of the Anycast; this is one of the key features that is, you know, often misunderstood in some circles, so we want to provide an introductory explanation of that, and then I'll cover the Root Server System today.

So fundamentally, the fundamental identifier on the Internet is an IP address, it's a numeric label; you know, we have IP version 4 or IP version 6, so you have the circle on the lower right, that's kind of an

example of both, and then all hosts connected to the Internet have IP addresses, so that is the fundamental identifier on the Internet.

So why do we need the DNS? The original problems are really the IP addresses often change, where it's hard to remember, you know with IPv4, but definitely more with IPv6, so there's a need to translate that into a more stable identifier, that the underlying IP address can change, but the identifier, the domain name is more stable. The modern problem is the IP addresses may also be shared and multiple IP addresses may serve as an entry point to a single service.

So, the Domain Name System is a lookup mechanism for translating objects into other objects, so I think the key point here is, is not simply, you know, translating IP addresses to domain names, where that's the most common use, but you can also, I'm sorry, translating name to IP addresses, but you could also do other mappings, so for example, mail Servers, IPv6, and you can also do a reverse, you know, it's the IP address to name translation.

On the left, you have essentially the structure of the space, the domain name space so when you think of the domain name system, you can think of it as the namespace, that composes of the names, the infrastructure that then makes the resolvers and recursive, that makes these names available, and also just the protocol itself.

So on the namespace, the DNS is a hierarchical host database; you start, you know, at the top level, which is at the Root and then -- I'm sorry, the very top is the Root and then below, beneath that is the top-

level domains, so these are the country code and the generic name, top-level domains and under that what we call the second level, which is mostly the domains, the domain labels, all of these you construct a domain name by doing the reverse and then by concatenating, that with a dot, so you have the www.icann.Org, the Root is probably kind of the invisible dot at the end of the domain name, which is often omitted. So that's the structure of the database and the namespace.

So before I go further, kind of some definitions high level, the Root Server system is defined as the setup Root Servers that collectively implements the Root service, you know, that serves the Root zone. So the Root zone is a DNS zone at the top of the DNS hierarchy; so if you look at this diagram, the Root zone is at the very top, is above the top level. It has no parent and it contains all the information necessary to contact the TLDs under it.

The Anycast instance is one network location responding to DNS queries on a Root Servers Operators IP address, so the Root Servers have IPv4 and IPv6 addresses and because of the Anycast technology there you could have many Anycast Instances that have the same IP address, so one instance is one network location responding to DNS queries.

When talking about the Root Server system it's also important to this ambiguity, some key actors here. First of all, the Root zone administrator, what we call the RZA here is the Organization responsible for managing the data contained within that Root zone; so

this involves assigning the Operators of top-level domains and maintaining their technical and administrative details.

The current Root zone, prior to the IANA Stewardship transition, you know the administrator is kind of the IANA plus the US NTIA; after the transition, this is really the IANA function fulfilling that with the ICANN community setting the policies for assigning the Operators at the top-level domain names, with the Organization and IANA implementimg those policies.

The Root Zone Maintainer is the Organization responsible for accepting the service data from the administrator formatting it into the zone file, cryptography sign it and distributing it to the Root Server Operators. So currently the RZM is Verisign, they sign and distribute the zone twice a day to Root Server Operators.

And the Operator is an Organization responsible for managing the Root service, our IP address specified in the Root zone and the Root hints file. So graphically thinking about the Root Zone Administrator is involved in the provisioning of the service where the Root Zone Maintainer takes this data and, you know, cryptographically signs it with a zone signing key and distributes it to the Root Server Operators; and the Operators on the right is about, you know, serving the zone, so this distinction is very, very important, I'll highlight that afterward.

One common question that was asked in these tutorials is, 'What's the difference between a Root Zone and the Root Server system?' So Root zone is a starting point, it contains the list of the TLDs and their name

Servers. It is managed by ICANN per the community policy developed. So at an ICANN meeting you see 3000 people here, those are actively developing the policies for the criteria's and who can get a particular TLD; and when the policies are developed, those are implemented by the ICANN Organization.

The Root Zone is compiled and distributed by the Root Zone Maintainer to all the Root Server Operators and the information is served by the Root Servers. So on the right-hand is the Root Server System is the response with data from the Root zone, so when they see a query, you know, they respond with the data, and it's currently distributed from 26 IP addresses, so 13 of those are IPv4 and 13 of those IPv6. So because of the Anycast nature of this, they're serving the zone from over 1,000 physical Instances, around the globe. This is purely a technical role to serve the Root zone and the Root Server Operators are responsible for the Root Server System.

I think you've probably have seen this slide many times, so probably there's no need to repeat. I think that the key points are the bullets below here and the Root Servers are at the entry points to the system, so that means absent of any other information, a query will begin, you know, the entry point to that DNS query would for the Root Server at the Root Servers.

Please note, the Root Servers only contain the information, where the authority the information for the TLDs, so when you ask a Root Server 'Where is ICANN.Org?', they can only tell you where to find the name Server for the .Org. Then you have to go to the .Org to look for the

name Server information for ICANN.Org, and so I think that's a very important part of the recursive nature.

Also, very important that's often omitted is caching is used throughout the system to avoid repetitive queries, so when you get something from the Root zone, you normally don't have to come back I think within 48 hours -- that's when it's cached -- and the DNS resolution perceives the actual transaction the user wants to do. So if you want to go to CN to see what's going on in the CN, what's the latest news, there are transactions, the DNS resolution precedes the actual transaction, so that when your computer gets the IP address of the www.cn.com, that's where it establishes a connection where it gets all this information.

So that's why the DNS is important because for content and other things that users want, you have the DNS resolution perceives that. I think it's at this point is really preaching to the choir, so there's no need to emphasize.

As I mentioned, the Root Servers only know what Servers need to be asked next, so will give you the list of .com Servers, the list of .net, .Org Servers and the list of TLDs goes on and the caching is for 48 hours.

So modern refinements to the DNS, so when we think about security, we usually think about integrity and confidentiality, so the DNSSEC really protects the integrity, meaning you get a response and that you are fairly confident, those who gave you the response are saying who that person is, so you ask a query, you get a response, you are fairly

sure that the person who claimed they are is actually the one giving you that response.

So that's protecting the integrity and the privacy enhancement is really bringing, protecting the confidentiality, so where you have the DNS over TLS so you are protecting the confidentiality, trying to reduce query leaks with minimizing the amount of queries sent to the various infrastructures in the system.

Anycast is really, as I mentioned, multiple Servers share the same IP address, it relies on the routing to determine where the nearest IP address, nearest location, where to respond to your data and it really protects against distributed denial of service attacks.

So a quick explanation of Anycast. In the DNS terminologies, there's Unicast, there's Anycast, there's also, Multicast, but the Unicast, you know, it's a package from sources all go to the same destination, a single instance serves all the sources. In the case of DDoS attacks, the traffic's all go to the single instance.

Anycast is multiple instances serve the same data to all sources and the sources reach destination based on kind of intermediate routing policies. So this means the sources get the data faster by having a machine closer to you topologically, and also denial of service attacks is sent to the closest instance.

So, here's one example; you have a source, a destination, one IP address serves that destination, so the traffic goes from the source to the destination through the routing. Here you have three

destinations, they all have the same IP address, so they're all advertising the same IP address, and the intermediary routing determines the closest, the shortest path and there the source gets the response from the destination.

So, in the case of denial of service attacks, you have a DDoS attacker launching you know, a flood of traffic and because of the Anycast nature, those traffic's are directed to the nearest, per the routing policy. So, you know, the red that will be -- you know, goes to that on the upper right and so that's where that link will probably go overwhelmed, by all the others traffic's, you see the source, the good traffic's, they keep on working, so I think that's the Anycast nature of it.

Root Server Systems today; you know, as I mentioned, think of the history of the Root Server system, you know, coincides with the development of the DNS. I think in fact, when the original DNS was developed, people wanted to try out the system, so they needed to develop Root Servers, and therefore, all these Root Servers were setting up in the early stages.

So from '83 to '86, you have these four Root Server addresses, and then the second from '87 is really, you know, you have a larger network, so this is really accommodating to a broader network, so for example, you know, from the early network to the NSFNET, that's where you have your hosting, you need to have more addresses hosting the Root Servers. And then in 1991, we have eight addresses and then up to 1998 we have 13 addresses.

So I think to think of this growing organically, I can hear it's always things [inaudible] very political but I think in terms of the system-wide it's changed over time, responding to the technical demands and scaling issues are now solved with Anycast, because there are over 1,000 Anycast Instances around the world and they're most likely -- you know, you have a few really near your location geographically or topographically.

We have the Root Server system addressing today, we have all the -- in the IPv6 addresses added starting in 2008 and currently, all of them have, you know, IPv4 and IPv6 addresses.

So this is a list of the Root Server identifiers today and the column on the left, the Hostname, the IP Address and the Manager. So let me stop that here. As you can see here, the list of Organizations operate, you know, the service, you have universities, so these are really the early pioneers where they're testing out the Root Server System.

The University of Maryland, for example, was added because of the proximity of connection to the NSFNET, and then you have ISC, that's the Developer of the most popular DNS software, BIND; Department of Defense, you know, Army Research, a lot of these early research implementations were there, and going down, we see those have been expanded internationally, with the RIPE NCC and the Netnod in Europe and with the WIDE project in Japan,

So I want to stress the Organic nature of the growth and the Root Server Operators have been, you know, growing to meet the demands

of a technical change and the demand of the global community, and they have been providing a stellar service for the Internet community for the past several decades.

If you go to the Root-Servers.Org, it's a more animated picture where you can drill down to see where those are. I think here to specify is that, as important to have a Server near you, but sometimes the peering arrangements is also very important. So the Root Server doesn't really control where your network traffic goes, the ISP does; and sometimes if you don't have the right peer arrangements, you could have a Root Server next to you but your queries do not go to that Server, but overall, you can see the map, you can drill down.

This is the graph that I mentioned earlier; Provisioning, Distribution and Resolution. Think of provisioning the IANA function, the Root Zone Maintainer, and then the Distribution in between, from the Root Zone Maintainer distributed to the Root Servers, they in turns had their own distribution instance, distribution mechanism to all their Anycast instances, and then the resolution, the Operators serve the zones where the DNS resolvers query.

I also want to plant a seed here; you know, on the white here, the provisioning aspect, that's really what the IANA transition is about. Bbefore you had the three actors, the IANA function, the US NTIA and the Root Zone Maintainer, so the community through a multi-year policy, changed their model, the aspect. And the governance of that is really what the RSSAC has been working, and that's addressed in the RSSAC 37 and 38.

Root Server Operators; 12 different professional organizations really focusing on providing a stellar service, diverse Organizations and operations.

Very important person, when Jon Postel assembled the group of Operators, he had this diversity principle, I don't know, probably in mind but I think when working on the history document and going through the IETF interviews, he felt that -- when Postel, they want to have a diversity of the Operators that reduces the capture and also the chance of failure, and also the ones that become Root Server Operators must have the community support behind it and make sure it has longevity. So for example, the RIPE and the Y, that's how those two were added.

So I want to specify the diversity principle is really important, and today the Root Servers also are diverse, not only in Organization structure but also in how they approach operations technically, organizationally, geographically and different funding models. So that's a really important feature, how the system is designed and I think it has a lot of wisdom going forward.

Obviously, the Operators co-ordinate through industry meetings so through RSSAC at ICANN, you know, they have communication tools, they share data, and they have emergency response capabilities. So, I want to point you to the Root-Servers.Org; here you come to ICANN and see this presentation and you may participate in the RSSAC meetings, but I also want to point you to the Root-Servers.Org website

directly, and they too also coordinate and publish some documents., so please give a read there.

One thing, you know, already I wouldn't repeat what the Operators are doing, what the Operators are not doing is policymaking. So for example, determine who has the right to operate which TLD and which one has the right to go into the Root Zone. The Operators are committed to serve the IANA zone, which is kind of the global interoperable Internet, the Root zone, but they are not involved in policymaking. And finally, they are definitely not involved in data modification, so they serve the data that is given by IANA, distributed by the Root Zone Maintainer. Let me see

Some myths corrected. I already alluded, in the past Root Servers control where the Internet traffic goes; the reality is the router's control where the traffic goes. A similar myth is most DNS queries are handled by a Root Server; actually no, the most DNS queries are not handled by a Root Server, so some TLD Servers like the dot com, they handle much more traffic.

Administration of the Root Zone and service provisioning are the same thing, I think I have made that point clear in the previous diagram. The Root Server identifiers have special meaning; so the reality is none of the Root Servers identities are special. There was a point that, there's a myth that maybe A Root is special because that's the first letter; no, they are not, because when the recursive resolver -- when they do a start, they see which one they have gives the closest,

quickest response, and they use it there and when that fails it kind of sometimes randomly selects from the list of remaining Root Servers.

It's also not true there are only 13 Root Servers, there are more than 1,000 Servers globally, but there are 13 technical identities, 12 Organizations managing the Servers. You want to keep it in a fairly small number because, you know, for obviously planning, coordination purposes, you don't want this group to be too large.

The Root Serer Operators conduct operations independently; they do so because of their -- you know, they make independent decisions on how they conduct operations, but they also coordinate the operation as a whole.

And finally, the myth was that the Root Server Operators only receive the TLD portion of a query; so the reality here is they usually receive the entire query, although now we have the latest technical standards that minimize the amount of exposure query to different parts of the DNS infrastructure.

So with that, I hear a beep so I think my time is up, so I'm going to hand over to my colleague. Okay, let me -- this one last slide, so what does this mean that if you are an ISP or if you are from a developing country, thinking about the Root Server System and your networks so you probably want to have 3 to 4 nearby Instances; as I mentioned. it is very important to increase the peering connections.

I think we have one Instance from India where you have servers very close but it's going elsewhere, where in turn is the peering connection issue.

Turn on DNSSEC validation in resolvers to ensure that you're getting a modified IANA data, so that you can verify the integrity. We always make a pitch here for you to participate and contribute to the RSSAC Caucus, where the technical advice is created. My colleague, Ozan, will go over this in a bit of detail.

And if you're interested in hosting in a Anycast instance, talk to the RSSAC members after this presentation. Can I have the RSSAC members, raise your hand? Okay, you can see we have some RSSAC members here; and also you can send an email to askrssac@icann.Org. I believe the Root Servers Operatives themselves are also setting up a mailing address, so in the future, we can probably add those here.

With that, let me hand this over to my colleague, Ozan, to take you through RSSAC and RSSAC caucus. Thanks.

OZAN SAHIN: Thank you, Steve. My name is Ozan, I am an ICANN Org member in support of the Root Server System Advisory Committee and working out of the ICANN Middle East Africa Regional Office located in Istanbul. So in this presentation, I will talk about the Root Services System Advisory Committee or RSSAC and they are RSSAC caucus, and in the

second part, I will talk about the work underway on the evolution of the Root Server system.

So let's start with looking at the role of the RSSAC. It has a narrow scope; the role of the RSSAC is to advise the ICANN board and the ICANN community on matters relating to the operation, administration, security and integrity of the Internet's Root Server System.

So this slide is to give you a better understanding of its role; RSSAC is a committee that produces advice primarily to the board but also to other ICANN bodies and organizations involved in the overall DNS business. The Root Servers Operators are represented inside the RSSAC but the RSSAC does not involve itself in the operational matters.

So if you look at the organization of the RSSAC, the RSSAC is composed of appointed representatives of the Root Server Operators, and they're also alternates to these representatives, and there are Liaisons. The RSSAC caucus is a body and of volunteer subject matter experts, and the RSSAC Caucus members are confirmed by the RSSAC based on statement of interests.

The RSSAC has currently two co-chairs, Brad Verd and Steve Baker, who are in the room with us today; Fred Baker and Brad Verd, I'm sorry, Fred. And also I'd like to note that RSSAC is transitioning into Vice-Chair/Chair leadership model and RSSAC will have its elections for the Vice-Chair in less than a month.

I talked about the Liaisons, there are four Liaisons of the RSSAC and there are incoming to the RSSAC from other structures and there are four of them from the RSSAC to the other structures, so if you look at them, the incoming ones, there's a liaison from the IANA Functions Operator, a liaison from the Root Zone Maintainer, a liaison from the Internet Architecture Board, and a liaison from the Security and Stability Advisory Committee, which is another Advisory Committee of the four advisory committees that ICANN has.

And looking at the outgoing Liaisons from RSSAC to other structures, there's a Liaison to the ICANN board, a Liaison to the ICANN Nominating Committee, a Liaison to the Customer Standing Committee or CSC, and a Liaison to the Root Zone Evolution Review Committee or RZERC.

So looking into the RSSAC Caucus in detail, there are over 100 technical experts in the RSSAC Caucus currently and these members basically apply to be a member of the RSSAC Caucus by submitting their statement of interest and they get public credit for their individual work performed at the RSSAC Caucus so they can contribute to some of the publications of the RSSAC and they get public credit there.

The purpose of the RSSAC Caucus is basically to bring expertise to publications because they're DNS experts, also this provides transparency to the RSSAC because with the RSSAC Caucus, you can basically engage in the publications and the RSSAC work.

There are currently two work parties within RSSAC, one is the Modern Resolver Behaviors Work Party and the other one is the Root Server System and Metrics Work Party. So the first one is to study the behavior of existing deployed software and recursive resolvers through both code bases and available data sets, and the RSS Metrics Work Party is to define system-wide external verifiable metrics which demonstrate that the RSS Root Service System as a whole is online, serving correct and timely responses to end-users.

If you look at the some of the tools and mechanisms that contribute to the transparency of the RSSAC and RSOs, the RSSAC has a web page which is rssac.icann.org, so if you go there you can find the names of the members, also the minutes of the RSSAC meetings and you can also see the RSSAC publications there.

Also, I just talked about the existence of Caucus which provides transparency to the RSSAC; also, the RSSAC meets during ICANN public meetings, as we are doing now, and these meetings are open to the public. The RSSAC also meets with other groups; for instance, for this meeting at ICANN66, the RSSAC met with the SSAC, Security Stability Advisory Committee, and it will also meet the ICANN Board. Also, the RSSAC coaches provided briefing to the Governmental Advisory Committee earlier in the week.

We are having a tutorial now which also adds to the transparency of the RSSAC and the co-chairs and the RSO representatives are here to answer your questions, and RSSAC has a document which is called RSSAC 000, which defines its operational procedures, which also is

another bullet point here contributing to the transparency of the Advisory Committee.

Similarly, RSOs also have some transparency mechanisms and tools. Steve just mentioned that you can ask your questions to Root Server Operators by sending by email your question to ask-rssac@icann.Org; they do have a web page, as Steve mentioned, Root-Servers.Org. And then RSOs also participate in RSSAC and they have individual web pages too and they also collaborate, they also work together on collaborative reports on major events.

So that was the part of the presentation with respect to RSSAC and the RSSAC Caucus. The presentation will continue with the work on Root Server System Evolution, which is underway. If you look at the timeline, this started more than a year ago with the publication of RSSAC037 and RSSAC038, which proposed a new model on the governance of the Root Server System, and then the ICANN Board and directed ICANN Org to look at these documents and come up with what is called a concept paper.

This concept paper was published back in April 20 19 and then it went up for public comments. The public comments period was closed back in August 2019, so moving forward, the expectation is that there will be a governance working group formed by the end of the year, and this governance working group will be working on first developing this model and then the implementation by 2022.

So here's an overview of RSSAC037. It defines 11 principles for the operation and evolution of the Root Server System. It proposes an initial governance model for the Root Server System and its Operators, and it also demonstrates how the RSSAC037 model works through a set of scenarios on designation and removal of Operators.

If you look at RSSAC038, which complements RSSAC037 with recommendations, on this document the RSSAC recommends the ICANN Board to initiate a process to produce a final version of the model based on RSSAC037, and also estimate costs of the RSS and developing the model, and the initial efforts should focus on developing a timeline. The RSSAC also recommends the ICANN Board to implement the final version of the model, based upon the principles of accountability, transparency, sustainability and service integrity.

This slide illustrates what is proposed on RSSAC037, so you will note three areas: the governance, DNS Root Operations and Onboarding and Offboarding of the RSOs. So, if you look at the top on the governance area, you will see the stakeholders which is the ICANN Community, the Internet Engineering Task Force and the Internet Architecture Board, and also other stakeholders, the Root Server Operators.

RSSAC037 proposed five functions, and on this document, the proposed functions were: Performance Monitoring Measurements Function, Designation and Removal Function, Financial Function, Strategic Architect Policy Function, and a Secretariat Function. And you note on the bottom half of the slide that there are some

performance metrics and I just talked about a working party within the RSSAC which is the Root Server System Metrics Work Party. So this work party is now trying to define some of the metrics that will be used in the system, and with those performance metrics, you also note the designation and removal function so, operators can be removed or new operators can be designated to the system.

And the concept paper drafted by ICANN Org it envisions some structures that correspond to those functions. So what does it envision: the Root Server System Governance Board, the Root Server System Standing Committee, the Root Server Operator Review Panel; and for the other two functions, the Finance and Secretariat Functions, it envisions ICANN Org.

Also the concept paper outlines a community-driven process to finalize a new cooperation and governance model for the Root Server System based on recommendation 1 in RSSAC038. So in the first phase, ICANN Org is reviewing and emulating RSSAC037 at the direction of the ICANN Board which is done in the second phase RSSAC037, the Concept Paper and Governance Working Group documents are made available for public comment which is also done. And in the third phase, which is the final phase, this is developing a new Cooperation and Governance Model for the Root Server System.

And the implementation phase has two tracks; one is the structural track where the Governance Working Group and develops a model, and on the Administrator Track it's to plan for implementation of the Governance Working Model which is led by ICANN Org.

So let's look at the Governance Working Group in detail. With respect to its composition, there will be representatives from the RSSAC, ccTLD, the Name Supporting Organization, the Internet Architecture Board, Registry Stakeholder Group and the Security and Stability Advisory Committee. Also, there will be liaisons from the ICAN Board, the IANA and the Root Zone Maintainer.

The Governance Working Group is tasked with working out the details of the model and also the Concept Paper tasks the Governance Working Group with some of the guidelines; for instance, committing to a timeline with clear milestones, working openly and transparently, seeking informed contributions when necessary, also embracing the principles outlined in RSSAC037, and referring to the RSSAC037 Concept Paper and public comment feedback as references.

So this concludes our presentation here, we are transitioning to the Q&A session, and as Steve mentioned, we have the Root Server Operator representatives here in the room with us, so, please raise your hand if you have a question and they will be answered.

STEVE SHENG: Yes please, gentlemen on the right.

YAZID AKANHO: Thank you very much, my name is Yazid Akanho, I'm an ICANN66 fellow. I have two questions; my first question is to understand, just understand, do all the Instances of a same Roots Server have the same

fully qualified domain name? They have the same IP address, I understand, but do they have also the same DNS name or not?

My second question is also to better understand, does DNSSEC [inaudible] the answer or only it provides security in who is giving me this answer? So if the answer itself is compromised from the Root itself, does DNSSEC [inaudible] prevent that? Thank you very much.

STEVE SHENG: So who would like to answer those two questions?

BRAD VERD: So I'll take the first one, Wes will take the second one. I think your first one, you are asking if all the Instances of a given identifier had the same name, and the answer is 'yes', they all announce the same IP space and provide the same service. I think what's most important here is that all the identifiers, all 13 of them, and all thousand instances share the same exact data, you will get no different answer from any one of the thousand servers. So they're all identical.

YAZID AKANHO: No, sorry, my question was not about the data they have. Let's say a Root Server, okay, we have x Instances of a Root Server, but my question is to understand if all those Instances they have the same IP address, I understand that, but do they have the same name? That is my question.

STEVE SHENG: The Anycasts instance of a particular Root, do they have the same domain name?

FRED BAKER: Well, sort of; the Anycast address is used by all of them and there is a name that returns that Anycast address, so in that sense, yes, they all have the same name. On the other hand, we as operators need to be able to access them, and so they will also have a maintenance address and a maintenance name to go with that. So, sort of.

WES HADAKER: So one additional bit of information, which is that if you query any of the Root Servers and you the NSID flag to your request, which is a way of saying 'who am I talking to?' and it will respond with a name, and you will find that the same multiple Instances of a given identifier will return at least some segment of that internal name, so that you can actually figure out which sort of sub piece that you're talking to. And we use that all the time for debugging and if you use RIPE ATLAS, which is a well-known DNS measurement system, it will actually list that for you.

YAZID AKANHO: That's where I wanted to reach, how do we identify each piece of the same instance?

WES HARDAKER: Yeah, so you can't direct your query to a particular Instance. If a Root Identifier has a whole bunch of Instances, all you can do is say, "I want to send out my request but I want to add DNS ID Flag," which not every Operator I don't think supports, most of them do, you'll get an answer back that will say, for example from me, you might get something back saying, 'I'm B3-LAX, which is an AB Identifier within our node near LA." So different things.

BRAD VERD: But its key to point out that the fully qualified name is identical to all the Instances.

WES HARDAKER: Yes, yeah, so when you request something from, you know, b.RootServers.net, you'll get back -- you can't force it to get to one instance or another.

For the DNSSEC part of your question, that's a great question. IANA and the Root Zone Maintainer sign the data long before any of the Operators get to see it, so if any Operator actually returned invalid data to you, that would be highly weird and probably something in the middle since we are all committed to serving the IANA Root exactly as it's given to us.

DNSSEC signs each blob of data and it doesn't matter how you get it, it doesn't matter whether you got it over from one Operator or another,

or whether you got it written down on scratch paper and handed to you and you typed it back in. It guarantees that it has not been modified no matter how you get it. So it's very different than transactional-based where I know that I'm talking to this server, but I have no idea how they got it and whether it was okay. DNSSEC provides that chain of authentication, the whole way through so it's very secure if you're just checking the contents of the data. Does that makes sense?

STEVE SHENG: Thank you. Any other questions?

WES HARDAKER: If you're in in the back or the side, you can raise your hands and somebody will bring you a mic.

STEVE SHENG: To the gentleman on the left, please go ahead.

NIKESH SIMMANDREE: My name is Nikesh Simmandree, from Mauritius. I am an ICANN66 fellow and I am speaking on my own capacity. So actually, I want to know how do you ensure that a Root Zone is properly replicated because nowadays we have lots of attackers and malwares, so in a secure way, how do you make sure that it is fully replicated.

STEVE SHENG: Question is: how to ensure the Root Zone is properly replicated. You mean properly securely distributed or…?

NIKESH SIMMANDREE: What I don't understand, that the Root Zone needs to be replicated, but how do you ensure that it is properly replicated?

WES HARDAKER: So that actually falls back to the DNSSEC question so that you know that the data has gotten, you know, to the Root Server properly within the distribution network, which was on one of the slides, so IANA gives data to the Root Zone Maintainer and the Root Zone Maintainer signs it and gives it to all of the Root Servers; that's done over a secure transport as well, so that we know that we're talking to the right entity so that we know it hasn't been modified anywhere from IANA to the Root Zone Maintainer to each Instance, so all of that communication is authenticated, in addition to the fact it is authenticated with DNSSEC, so it's sort of like double protected in that regard.

NIKESH SIMMANDREE: Thank you. Is there a way like -- is there someone who verifies the interaction between this?

WES HARDAKER: Each Operator's responsible for doing that themselves. I can tell you that our monitoring system has a whole bunch of checks to make sure that things don't go wrong along that path. Every Operator has a different mechanism for how they double-check that.

NIKESH SIMMANDREE: Thank you, sir.

STEVE CONTE: We have a question in the back.

GANGESH VARMA: Hi, my name is Gangesh and I think during the presentation, you mentioned something about the distance between Instances and having better peering. I wanted to understand that little better, I have a background in policy and not in tech.

STEVE SHENG: Question about peering arrangements and increasing instance, but also, you know, peering so any willing to take it? Right so we had -- I think it was a couple of ICANN meetings ago, it was someone from India saying, "We have these Root Servers really in our data center, but our traffic is not reaching that, we're literally so close to that and it turns out it is the peering arrangements," right? That dictates where the traffic goes, so because of the peering arrangement, the traffic is

not going to the location next to you, it's going across the sea, so there's nothing about increasing peering. Go ahead.

BRAD VERD: Also in the in the presentation, they say that BGP routing is how traffic is sent over the Internet, so just because I'm physically close to my Server doesn't mean I'm going to talk to it if my routes take me somewhere else. And I think the slide that you're referring to was talking about ways to improve kind of, you know, your experience with the Root Servers; one of the ways is peering with them. Peering is a usually a bilateral set up between two parties to share traffic free of charge, usually, and so obviously, that would get a higher preference than you sending a packet over a circuit that you're paying for.

So some of the challenges that we've seen, as you referred it to India, we've documented a lot of different instances where -- instances, that's the wrong word -- examples where traffic was actually leaving India and coming back in to India to talk to the Root Server that was in India, and that was a result of their peering relationships and how they were set up and the routing was.

So you really need to look at the topology of your network, if you're having challenges reaching any of the Root Servers because you got to find out which route you're actually taking to get to it, if that helps.

GANGESH VARMA: And that's for the Instance Operators to figure it out, or…?

BRAD VERD: No, no, we have no control over that, we have no control over the BGP routing table, so that's up to you working --

GANGESH VARMA: [CROSSTALK] that deals with this?

BRAD VERD: I don't know if there is one.

WES HARDAKER: Well, I mean, routing is a very tricky game and the reality is, is that there are many parties that can affect that answer, and the Operators have some control and some leverage to be able to manipulate things near them, but in the end, it's up to stuff beyond them to figure out who sends them traffic and there are things that we can do to try and trick, you know, other things to send it closer, you know, to send it down the path that you want, but in reality, there's no way to guarantee it and it requires human to human negotiation, it's a lot of our time to do that.

STEVE SHENG: Jeff?

JEFF OSBORN: I never get to answer any here, but you came from policy and I'm the lone business guy in the RSSAC thing, so let me just address from a slightly different point of view. There are roughly 1,000 Instances from 12 organizations; if you have a place you feel is underserved, it's really simple to get a Root Server instance.

My organization, for instance, has several hundred, we're a nonprofit and if somebody comes and says, "I'd like a Root Server instance," we literally get all excited. It's either, "If you have $4,000 in a rack to put it in, we'll do it tomorrow. If we need to buy it for you, it might take a couple of weeks, but literally on-demand, come on down and we'll open one up for you." I think with L.ROOT, very similarly, with ICANN I know the right folks are in the same position, so this isn't a matter of, this is all being controlled and we're trying to keep it somewhere, it's literally free and available.

What people find challenging is, once you've gotten it in and up and operating, the fabric of Internet routing is crazy complicated; if you've ever watched anybody try to model the weather, that's what it feels like, it is so complex and driven by humans running their own businesses for their own purposes, and then you pile up millions of them and it's -- that's where a lot of the inequity comes from, it's not that people aren't trying to put things in the right places, it's the routing itself is unbelievably complex.

GANGESH VARMA: What does it mean to be an underserved region? Is that the traffic is moving too slowly or is there congestion? What does it mean by if some areas are underserved?

STEVE SHENG: Fred?

FRED BAKER: It means that you don't like it for some reason. The reason might vary, it might be that you're just taking a long time to get an answer back, or that things are getting lost or something. So part of the question when you come and say, "I have an underserved area," I think it's underserved, I want to know what problem you're experiencing, you know, and then I'm going to look around and say, "So what's available and why are you not using it?"

I'm thinking of a particular case, a guy in Zimbabwe wanted a server and the server that he was, in fact, using was in Malawi, which is not very far away; and, you know, so what is it that he's trying to solve and what's getting in the way of doing that? And we talked to him through a number of, "Okay, tell me what your problem is," and were ultimately able to make him happy. But yeah, underserved fundamentally means you're not happy with what's happening.

STEVE SHENG: Wes.

WES HARDAKER: So let me end with one little bit of guidance that we typically give people. A lot of people have this notion that you need all 13 identifiers as close as possible to you, hopefully in your living room. That is not true, the general guidance is that if you look at how -- and there's been a lot of studies; in fact, I watched a presentation last week about this -- that you really only need three or four near you to ensure that you have fairly rapid responses, because most resolvers that are trying to reach you, they will pick the top couple that are really, really close and they'll just cycle between them and they ignore all the faraway ones, so you really only need three to four, sort of the general advice that we give a lot of people.

BRAD VERD: Just to add to that; it's not necessarily near you, it's near your resolver. So whoever your resolver or is, if you're using your ISP or somebody else, then it's near them, not necessarily near you. And I say that because you get Instances where somebody travels, we go to a foreign country, we connect back home via VPN and now that latency is, I'm asking the question from my home office, not from here. So these are all things you got to look into when it comes to latency and whatnot.

STEVE SHENG: Thank you, thank you. Very good question. Gentlemen on the right.

YAZID AKANHO: Thank you, Yazid again. I have two questions; the first one, have we ever faced a situation where the Root Zone Maintainer published wrong data into the Root Zone file?

BRAD VERD: The short answer is no, since VeriSign has only contractually been the Root Zone Maintainer for five years now, four years. If you go back through history, I believe there might have been an issue -- this predates most everybody here, so 1998, something like that there was something that happened that corrupted the database, and a corrupted file got published, but since then, as Wes alluded to earlier, wwe have a secure channel for distribution, the amount of validation that goes on before, after, and once it's out there is, you know, kind of off the charts, it's next level.

YAZID AKANHO: Fine. My second question and last, have we made a study to see what is the impact of not doing Qname minimization on the Root Zone System?

STEVE SHENG: What is the impact of doing minimization or not doing minimization?

YAZID AKANHO: No, doing minimization, will we just ask a short request, a short question to the Root Zone? I understand that and I understand that actually, not all resolvers are doing minimization, so I'm asking the question, what is the impact?

STEVE SHENG: So what is the impact of the Qname minimization on the Root?

WES HARDAKER: I work for a university and so we're constantly thinking of new things to go measure. That is an area of ongoing study by most of the world and there's a couple of problems; one, only the newest software is turning on Qname minimization by default; and in fact, I think one software is actually going to turn it off for a little bit and then turn it back on later, because they discovered a bug in their implementation of it. There hasn't been any studies yet that have measured it sort of at the what's the deployment level like. There have been studies that have measured how much extra time does it take for you to go resolve information?

I did a study in 2018 that was trying to measure how much extra privacy you've got by looking at Qname minimization versus other, you know, technology, so I compared it against TLS, I compared it against using the hyper-local Root System that ICANN talks about sometimes, which is actually the most private. So there's varying levels and you get different benefits depending on what you pick, but right now, there's a lot of people that are hoping to study really the

deployment model and how far that's gotten, but nobody, as far as I know, has done that yet.

STEVE SHENG: So the answer is not known yet, but some data. Gentleman on the right.

BENJAMIN AKINMOYEJE: I just want to go back about the question about the corrupt data about the Root Zone. I heard a story about somebody trying to hack a ccTLD, by hacking the Server and trying to do a delegation but during the process, ICANN validation stepped; basically, they replied and said, “Okay, I'm the one that asked for the delegation,” but since the process was manually, they stop it, so they have not been changed but I heard another story that ICANN is going to do like automation, but if they go to automation, maybe it's something that may happen in the future.

BRAD VERD: There's a lot of 'what if's' in that story, and I'd rather not comment on speculation; it sounds like what you're talking about is a conversation between IANA and some customer regarding a ccTLD, so that's before we ever get to the Root, so there's no data in the Root. I mean, any changes that happen there have to go through the IANA process, which is pretty intensive and lots of validation, before it even gets sent to the Root Zone Maintainer to then get published in the Root Zone.

So there's a lot of gates that that would have to make it through in order to get there, but I don't know anything about that story, and again, I wouldn't want to comment on a speculative story.

I will add that I believe IANA has implemented, they have two factor authentication on their system and the challenge, I think -- I wish Naela was here -- the challenge that they have is that a lot of the ccTLDs are not under contract so they can't force them to do certain things to get into their system, so they're working through that. I know IANA wants that system to be as secure as possible, but they have to talk to users which are the ccTLD owners.

STEVE SHENG: Thanks, Brad. Any other questions? Very good questions, thank you for the fellows, especially for the fellas. With that, I want to conclude by saying thank you for coming to this presentation; more importantly, for the Root Server Operators here answering the questions and providing a stellar service to the Internet communities for the past number of decades, thank you. The session's adjourned.

[END OF TRANSCRIPTION]