ICANN85 Mumbai | CF – GAC Meeting with the RSSAC
Sunday, March 8, 2026 – 09:00 to 10:00 IST

JULIA CHARVOLEN

Welcome to the GAC Meeting with the RSSAC on Sunday, 8 March at 9 a.m. local time. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior, ICANN Community Participant Code of Conduct, and the ICANN Community Anti-Harassment Policy.

Please remember to state your name and the language you will be speaking in case you will be speaking a language other than English. Speak clearly and at a reasonable pace to allow for accurate interpretation, and please make sure to mute all other devices when you were speaking. During the session, questions or comments submitted in the chat will be read aloud if put in the proper form. And with that, I will leave the floor over to Nicolas Caballero, GAC Chair. Thank you, and over to you.

NICOLAS CABALLERO

Thank you so much, Julia. Welcome back, everyone. I hope you enjoyed the beautiful city of Mumbai yesterday, last night, during dinner, or whatever activity you might have had. Before I give the floor and welcome to Jeff and Dave and the whole team, I don't know how many members of your team are with us today, but you're more than welcome, as usual.

Before we start, let me give you just an idea of what the day will look like for us, for the GAC, I mean. After this session with the RSSAC, we'll have a GAC discussion on New gTLDs on the next round, of course, then lunch, and then the meeting with the GNSO. Right after that, we'll be discussing GAC operations, including HLGM agenda, NomCom participation, potential participation or not, and so on and so forth. And finally, we'll have a discussion on the CCG RoR.

So, that's more or less the menu for today. Without further ado, let me introduce once again our good friend Jeff Osborn. He's the chair of the RSSAC. And to my left, another very good friend, David Lawrence, who's also a member of the RSSAC. They're going to be walking us through the latest details and updates and nuances of the Root Server Security Advisory Committee, right?

I sometimes have trouble with the acronyms myself. So, welcome again. The idea for this session is to make it highly interactive, so feel free to raise your hands, interrupt, and ask any kind of question you might have however technical or non-technical it might be. And with that, welcome again, Jeff. The floor is all yours.

JEFF OSBORN Good morning. In 35 years of making presentations to various technical groups, I find it is impossible to ever start with a functional AV system. So, my apologies. I am up here trying to make sure that this shows up on the slide. So, I appreciate your forbearance and your patience. My name is Jeff Osborn. I am the

chair of the RSSAC, which is the Root Server System Advisory Committee here at ICANN. Today I would like to spend some time explaining the place of the RSSAC within the ICANN ecosystem, and then explain about the root server system, how it works, and a little about how the DNS works.

The DNS is a vitally important part of what we all do here at ICANN and an important part of how the internet works. And I find hopefully we all can get a new look at it and get a little more insight if I do this right. So, if you'll bear with me, I'm going to try to be instructional and educational. Two parts, we're going to talk about the Root Server Advisory Committee and how it fits within ICANN, that's the shorter part, and then understanding the DNS root server system. The RSSAC has an extremely narrow scope.

It is really, really not a broad policy publishing organization. We all have a day job which is providing the root server system, and then we come here to advise, as it says, the ICANN community and board on matters relating to the operation, administration, security, and integrity of the Internet's root server system. That's it. This is an extremely narrow remit, and we try very hard to stick to it. I'm going to spend a couple of pages going through the information that is available if you wanted to do a deeper dive.

But it is not necessary and I'm certainly not going to read it out. There is an entire section of the bylaws that created the RSSAC, and there are a number of publications that we've generated that are available and pretty easy to find on the ICANN website. The

composition of the RSSAC is made up of 24 primary members. There are 12 root server operators, and each of them designates a primary member plus an alternate member, and there's one vote per operator, so a total of 12 votes. In addition, we receive liaisons from a number of other places inside ICANN.

The IANA, which is the Internet Assigned Numbers Authority that puts together the root zone, the root zone maintainer that cryptographically ensures that you cannot change anything in the root zone, the Internet Architecture Board, and the Security and Stability Advisory Committee of ICANN, the SSAC. We send liaisons to the ICANN Board, the Customer Standing Committee, The Root Zone Evolution Review Committee, which has the great pronounceable name RZERC, which sounds like a car part that the grease goes in, the ICANN NomCom, or Nominating Committee, and then any ad hoc ICANN organizations that show up, we will try to send a liaison to be on top of it.

There is also a group called the RSSAC Caucus, which is a wider group of experts. It currently numbers about 100 people. These are subject matter experts and the requirement for becoming a member is pretty tough. We are not looking for people who are vaguely interested in learning how DNS works. This is primarily made up of people with either doctorates in DNS or who have implemented a nationwide, corporate-wide, or university-wide implementation of DNS. And that's where a lot of the work gets done. We have those work groups that are performed by the members of the RSSAC and the RSSAC Caucus. Membership is

determined by a membership advisory board and will take applications from anyone.

For somebody who's simply interested, there is a mailing list version of the membership, and we find that people find that a useful way to stay on top of what we're doing without having to be embarrassed by how difficult some of the work actually is. The characteristics of RSSAC as a forum for discussion are really very simple and open.

We are very structured, we're very open, we're very transparent. I'm serving in the second year of my second term as RSSAC Chair, and we really very, very rarely have anything that isn't completely open in the public available afterwards for your perusal and available before something you can show up for. We basically develop and express feedback on policies and plans involving, again, just the root server system, very narrow remit. Primarily, we restate the long-established normative framework that's used to operate the root server system.

And yes, that's a sentence that the lawyer spent three hours putting together, so I try to read it out loud. We suggest continued development that grows out of that framework. There are some misunderstandings about what RSSAC is not, and so I would like to state a few of them right here. RSSAC is not an enforcement or supervisory body. It is not a representative body for the root server operators. I don't represent any root server operator except the one that I work for.

And RSSAC is similarly not a representative body for the root server system as a whole. There's 12 organizations that work independently. There's a lot of history, like I said, that's out there, and if you really want to fall asleep on an airplane, try reading the 47 pages of the history of the root server system. I find it interesting because I've actually been involved with the commercial internet since the early 1980s, and it's kind of interesting to read through and remember things that happened and settle arguments on who actually developed what. But again, we fully document all of this stuff where it comes from and the uses where it is.

The service description describes metrics that we use and measurements on how to figure out what your service level is, where you're getting it in something that is a really broad, complicated network. I think a lot of people are unaware of just how broad, complex, and huge the DNS is, and when you've really looked into it, you recognize just what a monstrous task it is trying to keep it all straight and just understandable.

There have been some governance related operations recently, and there is a session I believe, Dave is that the day after tomorrow? Do you know? The day after tomorrow there is an open session of the Governance Working Group. This is the root service system governance working group, which is entirely separate from the RSSAC. I'm not even a member of it. I sent policy people and a lawyer for my company because governance is hard. I prefer the technology part.

And then we give specific advice to the ICANN Board, and that's all available as well. And as things tend to be done at ICANN, the documents are named, the name of the organization, and then a number. So, we're generally RSSAC038, RSSAC001, RSSAC047. Those were available, and I'd recommend if you're interested, you can either present them to somebody who is concerned with learning more about this or if you're interested in learning more about it yourself. Part two, this I think of as the fun part.

This is where if you hang around ICANN for long enough, there are some questions you can't ask. Like there are some acronyms, I've been doing this for 12 years here, there are some acronyms I have no idea what they mean, but it's embarrassing to ask because I've been here so long, it would show that I wasn't paying attention.

So, I think the DNS and the root server system fall under that rubric where it's okay at home to say, like I often do, honey, I do not understand this DNS thing, but then I make my wife swear she won't tell anybody I work with that I didn't understand something. So, we're going to step through how the DNS works and the root servers role in it, and I hope it's instructive, and if something doesn't make sense, feel free to raise your hand because there's a logical flow to this and I believe in about 15 minutes when I'm through this part, you'll have grasped at least one thing that makes you go, oh, I didn't realize how that worked.

So, in introduction, we basically put this package together when we realized that there are a lot of people who need to be able to make

decisions about the DNS and how it works, even writing legislation and regulation, who are brilliant people in their own field but are not technology people. And so, how to explain this was the question. We had an interesting session where I asked a room of about 40 engineers to raise their hand if your mother knew what you did for a living, and no hands went up.

So, this is part of addressing that, where it is complicated, and if you do it every day, it seems very obvious, but explaining it to a broader audience is something challenging, and I'm hoping we're able to do that. So, I want to increase your understanding of the root server system, the root server operators, by explaining the role and the purpose of DNS, the role of an address resolver and an authoritative server, and how, when, and why a resolver consults the root server system, plus some common misunderstandings about the root server system.

The focus, again, the way we put this together, is this is the thing that people in positions of leadership need to know. So, introducing the DNS. At its root, it's really pretty simple. DNS is the domain name system, and it's the way that we find a computer's address, which is a series of numbers, using words. It's just that simple. When I was a kid, you had to know everybody's phone number, a string of numbers.

Now it's just an icon you just push. Similarly, the DNS is a series of words like www.amazon.com, that's all you have to know, and the DNS in the background will find the numbers that are needed to get

there over to the internet and connect to it. So, the computer needs an address like, you've seen numbers like this 40.81.95.116, DNS knows that tata.com, when I looked it up, was at that number. Now the numbers can change. They could take their whole data center and move it from Mumbai to Chennai and as long as they keep the DNS updated, all you need to know is you're looking for tata.com.

So, it's really a clever system and an easy way using human words to figure out where something is in a very complicated network. That's the guts of it. In a perfect world, I'd stop and say any questions, but there's actually a little more complexity to it. Connected devices of any sort, whether it's the computer or the phone or my washing machine or the car, use DNS to find each other on the network. And again, to sum it up, DNS questions ask about domain names, and the answers come back in a numeric IP address.

Questions asked with words, answer comes back with numbers. That's it in a nutshell. Why is this useful? I'm old enough to remember when you had to know those numbers to add in to get to somewhere else on the Internet 40 years ago, and it was horrible. So, this is all human-friendly. Nobody is going to forget www.example.com, but they would have a hard time remembering 192.168.45.99. The second really useful thing about DNS is, like I described, you could move the entire data center and nobody

would care because you still have the same name for it, it's just you change the addresses in the background.

A remarkable thing about the DNS, like I said before, is when you really get an idea of the scale of the thing, it is a massive distributed database that is global and works very reliably and very, very quickly. And the scale of it is almost unimaginable the closer you look, but we'll describe that. So, devices get their addresses from a thing called a resolver. A resolver is basically a device that resolves this name is this address, this name is this address, this name is this address.

There are millions of these things around the world. And these are what you as an end user would usually interact with. You don't interact with the root server system. That's down in the guts.

The resolvers, you may recognize some of these numbers. If you've ever had to add a DNS device to your computer, you go into the scary part where you're not really comfortable and you put 1.1.1.1 for instance is one of the largest, that's Cloudflare. Google runs a global resolver network that is at 8.8.8.8. There's a fairly newer nonprofit called Quad9 that runs obviously 9.9.9.9. And all these are is the address of a device that takes the name in words that you give it and returns a number.

So, resolvers can find and read the internet's phone books. The resolvers just go out and get information from other places and then answer questions to you. The phone book, the place where this data actually exists is what's called an authoritative server.

And as the name implies, it's authoritative. For instance, the biggest one out there, .com, is operated by VeriSign, which is the authority for where every single thing that ends in .com is. There's an Indian government organization that controls .in, and it can tell you where every single thing that ends in .in is, and is the absolute authority.

But there are many, many, many, many millions of these things all around the world who are authoritative for their own device, and they will always give the reliable answer when you ask their particular question. So, again, the simple way to look at this is what is the number for tata.com? The number is that string of numbers, and this happens in milliseconds and we spent a long time figuring out it happens about 500 trillion times a day. I get lost in the numbers. I think that's 14 zeros, so it's a really, really big number, and it happens pretty reliably and pretty quickly everywhere on earth.

So, the resolvers get their addresses from the authoritative server. The resolver remembers those addresses. This is called caching. And occasionally a resolver will start up from scratch or decide, I've had this data for a whole 24 hours, let's go get new stuff just to make sure nothing changed. So, depending on how much information it needs, a resolver will ask information in one of four simple cases. And then we're going to walk through these four cases. Case number one, the resolver doesn't have to do anything,

because somebody recently asked for that exact name, so it already knows it, it's already there.

You don't have to look up the number for your cousin Joe every time you look it up because once you have it in your phone, it's there. It's very similar to that. The information is cached. I've looked up tata.com, the next person to look for tata.com doesn't need to look it up, it's there in the memory. In the second case, you don't know what the whole string is but you know what the domain name is, okay.

So, if you know tata.com but you're trying to find mail.tata.com or info.tata.com, you'll go to the domain names authoritative server, and Tata controls that authoritative server where those addresses are, where they go, nobody else does, and that's where you're directed by the resolver. In the third case, let's say all you know is it's in India and it's in the Indian TLD .in, that's all I know. This is much rarer, but you go and ask .in, where the heck is Tata? It gives you that number.

That is the authority of that TLD. There are roughly, I'm going to forget the number, let's call it 1,500 TLDs in the world right now, and each of them authoritatively knows the list of everything to the left of its dot. So, now, here's the fun part. This is where we're going to walk you through the process of how the words become numbers by going to the various authoritative places that know that piece of information but not all of it.

Nothing knows all of the information. Everything controls its own authoritative part that's held together by the resolver which then keeps it in memory and gives you a correct answer. So, the big gray box is the resolver, that's this magical device out there that you as an end user will always ask one question and it will always give you one answer.

So, from the user point of view, you are always going to start with the question, what is the IP address for this name that I'm trying to find? And we're going to use the example www.example.com because that's what I used when I made the slide deck. The answer that you're going to want, and the first thing you see out of this big gray box is the answer that is the address. So, to the user, this is really, really simple, but the magic is all going on inside the resolver. So, let's take our first example. What is the IP address for www.example.com?

We ask the resolver, does it know the answer? And it turns out it has it in its cache memory. So, yes we do. There we go. That's how simple the whole process of resolving the domain name was and getting an IP address. And what's interesting is over on the right in the red letters the frequency of this method is showed up there. I think it's closer to 95% but the lawyer said to put 90%. That's how frequently the entire thing you're looking for is literally just sitting in memory, and when you ask the resolver, hey, I need this thing, it immediately knows the answer.

So, when it doesn't know the whole answer, we'll go to the next step. In the next step, we say, what is the IP address for www.example.com, and do I know the answer yet? No, I don't. Uh-oh. So, take a word off from the left and let's see if I know the shorter thing. Do I know example.com? And it turns out, yes, I know where example.com is. That's the authoritative source for everything to the left of it. So, I ask example.com's authoritative server, hey, what's www within your listing of example.com things? And it gives me an IP address and I put it in the memory. Do I know the answer yet?

Yes, I do. Excellent. That turns out to be about another 5% of the ways these happen where you're simply going to like I said tata.com, example.com, amazon.com and asking basically the organization, hey, I'm looking for stuff to your left. So, by the time we've used things that are already in the memory or things that are just a further explanation, we're up to well over 95% of how often these things happen.

Now let's suppose you come across something very unusual or your resolver just started up and it doesn't know things, we're going to add another step recursively as it goes additionally to other authoritative sources. What is the IP address for www.example.com? Do I know the answer? No. Uh-oh. Do I know the authoritative place to look up example.com? Nope. Oh, no. Do I know .com? Yes, it turns out I know where .com is.

I'm going to go ask those guys, hey, do you know where example.com is? It looks it up in the authoritative server and says, I've got it. Here it is. Let's stick it in memory. And now I know where example.com is. Do I know the answer yet? No, because the question is where's www.example.com? So, now we go and ask again. Now I know where example.com is, where's www? It looks it up, it gives me the address, it puts it in memory. Now do I know where the answer is? Yes, I do. This, it turns out, is about 2% of the cases.

So, we're getting really down to the edge cases, and if you stick with me for the next step, you'll see why this is important in describing the root server system. Now let's say you don't know anything. This resolver came out of the box, you plugged it in, you turned it on, it's got no memory, it's got no nothing, it doesn't know anything at all. So, when we ask it, what is the address for www.example.com? It goes, I don't know. We ask it, do you know the address for example.com? No. Do you know the address for .com? No. So, it knows nothing.

Basically, what this is a piece of resolver software. My company, by the way, develops an open-source piece of software we give away called Bind9, which is probably still the most widely used of these products. We give it away. And one thing it has in it is a little short list, and it says, here are the 12 guys who have access to the root server and their address, ask one of them. So, this is where the root

server comes in because you don't know anything and you have no memory of anything.

So, with that, you can say, hey, you've got the address for .com. Excellent, .com Where is .com? And it gives you .com, it puts it in the memory. I still obviously have some recursing to do. Now I've found .com.

Next, we're going to ask where example.com is. I'm going to get that answer. Then I'm going to ask where www.example.com is and get that answer. And finally, put that in memory so the next time nobody has to ask again. Do I have the answer yet? Yes, I do. Now all of that happened to you as the end user invisibly. You just asked the resolver and sometimes it knew right away and sometimes it had to go back to the source to figure out the first part and then back to the source of the next part and back to the source of the next part. That's what allows this to be completely distributed and controlled by each of the individuals responsible for one authoritative thing.

There's no overwhelming authority. But here's the part where people miss the root server system. We estimate you have that bottom case of having to ask the root server about one in five to ten thousand queries, 0.02% of the time. And the problem is if you read a computer textbook and it describes DNS lookups, it does this upside down because you're starting with a new resolver out of the box that knows nothing. And so, every student of computer science knows the first thing you have to do is go to the root server.

Well, what they don't mention is, and then you won't have to go back to it for 10,000 times. So, it looks very much like it's a controlling function, which it absolutely isn't, because everything is probably in memory. And if that makes sense, I'm very happy about it, because this is a hard concept, but that's pretty much DNS in a nutshell. So, there's some numbers out there. A zone file is a list of addresses. So, there are something like 350 million domain names out there on earth. And that's a here's the name and here's the top-level domain.

They can have infinite numbers of words to the left of that and it's all controlled by the domain name registrant. That's the folks who got the domain name. So, if it's ibm.com, they control it all. If it's tata.com, they control it all. There are 350 million of them out there, and each of them can have dozens or hundreds or millions of additional addresses. Like I said, the complexity of this and the size of it is really amazing. But everybody who is authoritative gets to control their part of this. The TLD zones are the top-level domains where you know what all those are, they're all to the right of the dot. So .in is responsible for every address within .in. .com is responsible for every domain name within .com.

They are authoritative and there's a total only of 1,450 or so of those top-level domains. And then we get to the root zone, and the root zone is remarkably small. It's one file with the addresses for the 1450 top-level domains that are out there. It tends to be two of them because there's at least one of each of the types of IP addresses. There's an IPv4 and an IPv6, so there's at least two.

There can be multiples. But basically, it's a really simple file, and all it does is tell you the address of the guys who know where everything in the top-level domain is.

So, in review, the root servers each hold a copy of the root zone. All the root zone is a list of 1,450-ish top-level domains and the IP address to go to if you want to look for a domain name that ends in that top level domain. So, there's a .com server, there's an .in server, there's an .ai server. The TLD's authoritative server knows the next word. We know the thing to the right of the dot, they know the thing to the left of the first dot. They know where amazon.com or flipkart.in or claude.ai is. And then when you get down to those guys, their domain name server, they can add all the stuff to the left.

And sometimes there are just stunning numbers of information out to the left. I've given this presentation twice when somebody came up to me afterwards and said, we add our customer names to the left. So, we have 30 million of these things two dots down in the name. And it kind of boggled my mind that the system is so adaptable, you could simply put every human being on Earth three fields down, and there is a structure for organizing it, there is a way to find them, and it's all orderly, and it just works. It's kind of amazing.

So, again, the primary thing going on here is a resolver finds and returns an answer to the question, where is this string of words, here is the number where it resides. So, I was going to put this on

a t-shirt so they made me put it on the slide. This is really an important thing to remember about the root server system. In the millisecond world of resolvers, the queries to the root server system are remarkably rare. And I would argue that is not obvious to most people who haven't spent a lot of time around them. So, just a couple of other things that people get wrong on occasion and I just want to make clear. All the root server system provides is an address. It is literally a string of numbers that can tell you where the authoritative server for the top-level domain is.

There is no information, there is no content, we don't see email, we don't see web addresses, we have no idea where you're trying to look, and the end user doesn't even ever get here. The end user asks a resolver, the resolver says, if I don't know that already, I'll look up if I know absolutely nothing how to get to the top-level domain address. The root server system isn't a gatekeeper. This is one where somebody wrote a computer science textbook 40 years ago and really did a disservice because they said you get the resolver out of the box and the first thing it does is talk to the root server system, which makes it sound like everything that ever happens in DNS requires starting with the root server system, and it's not the case.

You start this thing once, it gets the root zone, and then it's very, very infrequent that it's ever going to need to look there again, and it's almost never the end user. I'd say never, but there are corner cases because engineers are that way. You can't say never around an engineer. They'll fight you all the way to lunch. So, traffic is

transmitted without the need to wait on the root server. The idea of the root server system ever causing latency that a normal user can recognize is probably not true, and I would say always not true, but there are engineers listening and they'll prove to me that in the third phase of the Saturn moon it sometimes happens.

So, primarily, this is the case. The system is very stable and secure and resilient. This is a system that was built years ago and was trying to keep up with the growth of the internet, and about 20 years ago, a concept called Anycast showed up. And once Anycast was applied to this, there really isn't much need to grow much because it is a massively redundant system where we've figured out one of these 1900 server machines could handle the root server needs of the entire earth. It would need to be a fairly hefty machine, but I think you could do it for about $10,000.

So, the redundancy of this is remarkable. Banks often use a thing called failover, and there are two machines. One is a hot standby and one is the actual machine. And the billions of dollars in the bank count on if this one goes down you have another one and that's good enough. You could knock down 1800 of these 1900 and nobody outside the geeks who watch the monitoring would notice an effect. It could literally have zero effect on the end users on the earth. So, it's stunning how um redundant the system is.

Each server contains an entire hundred percent of the root zone, you don't need multiple ones, they are run on very diverse versions of hardware. Everybody pretty much picks their brand and makes

sure it's different. We use diverse operating systems, although most of them are flavors of Unix. We happen to like FreeBSD. We use diverse DNS applications. We try to have no more than a couple of people using the same one. We obviously use our own DNS because we wrote it and it's easier.

And then there's diverse routing where there was a famous incident where Digital Equipment Corporation in the 1980s built a fiber optic ring around its New England campus and one pickup truck hit the wrong telephone pole and took the whole network down. So, we are very, very careful about making sure that we have diverse routing, and it would take more than a few bad incidents to have any reaction, and again, it's difficult to even hypothetically create a situation where an end-user would notice. There is no single point of technological failure.

Again, we're 12 very diverse organizations, government, education, I run a non-profit, very separate, very different types of organizations. And so, while we communicate a lot, we're very independent of each other, and an unfortunate force majeure could affect one, I would imagine, but as I've said before, the disappearance of any one of these would have no effect on the rest of the network. It is simply a failover where every single part of this could replace any other part of it. The system has operated since the 1980s and has never had a blackout yet. Knock on wood, but I'm pretty confident.

If you look at the architecture, it is a remarkable example of something that is just incredibly resilient. So, 30 years of operating anything without a major failure in the computer world I'd argue is something you have to call reliable. So, here we go once again to say we don't decide what's in the root zone. IANA decides what's in the root zone under very careful controls, and if you're not familiar with IANA, it'd be interesting to get that information.

They take that information under incredibly controlled circumstances, cryptographically lock it through the group called the RZM, the Root Zone Maintainer, and this is all information determined by each step. The registrant who owns the domain name gets to decide what the addresses are to get to it.

The TLD register gets to decide its authoritative server and how that information goes and they give that to IANA. IANA authenticates that all that information is correct and provides it to the root zone maintainer. The root zone maintainer cryptographically signs it and hands it over to us. So, the root server operators have absolutely nothing to do with the content, we merely take a signed thing we can't possibly change and make it available to the world.

So, in summary, the root server system provides the addresses to get to the TLD's authoritative servers. We don't decide what information appears in the root zone, we are not a gatekeeper, we don't hold up anything. I would argue it is almost impossible to

determine a circumstance where you could tell an end user latency had anything to do with the root server, and the system is resilient.

There is no single point of technological failure and no single point of institutional failure. Thank you very much for your time. I hope that was instructive and we have some time for questions.

NICOLAS CABALLERO: Thank you so much, Jeff. Indeed, very detailed and nuanced presentation. Thank you so much for the energy and the detailed presentation as usual. A pleasure to have you here. There are some questions already and there's a hand up. Let me read the question first, and then I'll give the floor to Netherlands and then I have India.

The first question being, "The Root Server System Governance Working Group submitted the report to the ICANN Board in the last couple of weeks. That group was formed as a result of an RSSAC advice, and it says RSSAC037. What do you have to say about that? Any updates?"

JEFF OSBORN: I believe when I opened, I let it be known, there is an organization made up of some of the root server operators and other organizations as well. I'm not a member of it myself and I have been warned that RSSAC and the Governance Working Group are two separate organizations. So, I am not trying to avoid the

question, and in fact brought a Board member here who is much more knowledge about it than I am, Dave Lawrence. So, Dave.

DAVE LAWRENCE

Hi, all. I'm the IETF liaison to the ICANN Board and also have been one of the Board representatives to the RSS GWG. And as Jeff says, they have been separate, while many root server operators have participated in the governance working group. As the question noted, the initial recommendations now have been finalized by the governance working group and the comment period has ended. The GWG has increased its capacity to provide incorporated feedback from the comment period, and the report will now be considered by the ICANN Board to determine what the next steps are.

NICOLAS CABALLERO

Thank you very much, Dave. I have the Netherlands and then India.

MARCO HOGEWONING

Thank you and good morning colleagues from the Netherlands. This is Marco, one of the vice chairs, but putting on my hat as a country representative here, and thank you, Jeff, for this, once again, very useful introduction to the DNS system. As I mentioned in chat, hopefully we can carve this out into a clip for every newcomer, not only to the GAC, but to all of ICANN. This is I think the basics.

And then to pick up from where you just left, and indeed we did ask the GWG to join us, they politely declined pointing to their recently published report. I saw RSSAC also responded positively to the first conclusions of that group. But to my question and I've been looking at several other RSSAC publications more recently. For instance, what is it, RSSAC60, the one that talked about incident reporting, and I noticed that a lot of the recommendations developed by your group seem to hinge on the Root Server System Governance Working Group completing their work.

In other words, I see a lot of more recent RSSAC publications that sort of go like, oh, if the root server governance body exists, then this or this. I wonder if you have any views on, as RSSAC, views on the timeline for completion of this work. And would you recommend to the root server operators to implement certain features without that governance body existing, just to shorten some of the timelines, for instance, on getting clarity about the incident reporting? Thank you.

JEFF OSBORN

Marco, thanks. These are some questions we get a lot, and I appreciate in our going back and forth over this that you recognize I have a little bit of a political complexity that I am not encouraged to speak for what the GWG does, and so, thank you for splitting it out and making it the RSSAC part. This is going to sound a little self-serving, but I really don't know what else to say about it.

The root server system is kind of remarkable in a couple of ways where you normally think about a brief outage of something being meaningful, where we're talking about the address of a top-level domain. There are top-level domains that haven't changed their IP address in 30 years. So, if, God forbid, the system didn't work for 10 minutes, it is possible and likely that nothing changed. In fact, I've seen these things go day after day after day, where nothing changed other than in a big TLD that has 35 machines, they added another one. So, the idea of what is a failure is really very different.

So, every time we try to come up with metrics that say things like, there was a 20-second outage, who cares? There is a time to live, which is the measurement of how long is it before I throw away the data and get a new one. On most of this stuff, it's about 24 hours or 48 hours. So, what we've had a problem with is when we measure the metrics, we can come up with a bunch of failures, but they don't mean anything.

So, once you have a failure in the computer world… like, okay, my son playing a game and having a latency increase from 3 milliseconds to 12 will come running down the stairs and yell at me that I've ruined his life because it shows up and it has an effect. But the root server system is so far down the line of things that it really takes a long time before any of this becomes noticeable. So, again, I'm not trying to say our mistakes are not important, but I'm saying this is really, really a glacial paste set of numbers.

My company, wrong company, ISC, has a product called Bind9. And one of the things they did years ago was realize that if, God forbid, you couldn't get to something, if you have an island nation with two fiber optic lines and both of them get cut, this is an actual real situation, it defaults to a thing called surf stale, which means it goes, well, I can't get any information, I'm going to act like nothing changed. And the incidence of anybody noticing these is tiny, because nothing changed.

Furthermore, if you can't get to us, you can't get to anything else anyway. So, Marco, you and I could have a long discussion about this, I'd be really interested. I've actually given a similar presentation to this to RIPE that I think you were at. Further discussion is all going to be good, but the issue of the governance and the issue of the operation of the network are two separate things, and I don't want to step on people's toes by talking about the political. Did that help at all or am I just avoiding your question?

MARCO HOGEWONING No, no, thank you for that. I know that it is a political minefield, and let's pick this up when we happen to be in the same room. But I think this is helpful, although I might have a slightly different view on it. The usefulness of let's put it reporting and then providing clarity on their misses in terms of future learnings, but thank you for your answer. I'll leave the floor to other people. Thank you.

JEFF OSBORN: And Marco, I would encourage the ongoing discussion. I got busy before coming here and wasn't able to be as responsive as I wanted tom but please let's continue it.

NICOLAS CABALLERO: Thank you, Netherlands, and thank you, Jeff, for the answer. I have India next.

SUSHIL PAL: Thank you, Nico. Thank you, Jeff, for a remarkable presentation. I mean, maybe I would start from the last point, which is the resilience part. I think it's a wonderful tone that the system is working for the past 30 years, although with some ifs and buts. I hope it would be worthwhile also to point out and correct me if I'm wrong regarding the cache part. I think cache part is also not like, it's not like a permanent thing, right? So, this works for what, seven days? TTL in the cache is what? TTL in the cache is for 48 hours or seven days?

DAVE LAWRENCE: A typical TTL in the cache for a top-level delegation is anywhere from a day to two days.

SUSHIL PAL: A day to two days.

DAVE LAWRENCE It does vary by registry. Two days maximum though.

SUSHIL PAL So, in case nobody kind of searches for the same domain name, it will have to go to the root zone, right?

DAVE LAWRENCE Yes. So, after two days, you have to refresh the idea of where the .india, yes.

SUSHIL PAL I get a point. So, the point I'm highlighting is that the dependent of the root zone still lies, right? So, that number may not actually the 0.02%, which is 1 in 5 or 10,000. They do not actually reveal the true story because the domain name is being searched again and again by various citizens, so by various people, right?

JEFF OSBORN No, once you've called it up, it lives in the resolver cache.

SUSHIL PAL For 48 hours, right?

DAVE LAWRENCE Right. And I will say I think I understand your question to the point of view of, yes, like somebody that's asking Google's cache, can you tell me about the delegations for .in? Google will only have to ask.

But then somebody asks Cloudflare's cache, and that will get a separate. So, yes, it is by the number of caches that are out in the world, we'll have to refresh that data.

SUSHIL PAL

So, I mean, this is important. So, I mean, there's a story behind that number, and that aspect also needs to be clarified when we throw that number up for a normal because that doesn't reflect the true resilience part, I think. If I'm not searching, if any citizen is not searching from the Google resolver, it will have to go to the root zone, right?

The number is low because somebody is searching, there are too many Internet users, and as the more and more Internet users grow, that number might still come down. But the resilience part still is not addressed because it still depends on the root zone. So, that's one part.

And also given in view of, again, in view of the changing geopolitics, so as to say a concentration of root servers in a few geographies or larger maybe say over a concentration, we've been talking of a supply chain resilience in all aspects of economy, and this is also one of the fundamental aspects of our economy because almost all nations' economy depends on the Internet as of now, does the RSSAC think of giving any advice to the ICANN for diversification of more root servers? To say it's not resilient, it is still, but that doesn't take away the option of making it more resilient.

JEFF OSBORN

Obviously, we've had this question before, and I am not a politician. I am somebody who has run a series of technology businesses. My father used to always say there's no excuse but there might be an explanation. And I think the explanation is important where the internet began in four locations in the western United States. And I was around soon enough after that, I was fortunate to be involved in putting some of the first connections across oceans and it was really thrilling.

It was so fun to do those first connections and everything. What happened was the footprint of the internet started in the United States and then spread to Japan and Western Europe, and about that time, we realized, oh it doesn't have to get any bigger. And so, being a bunch of technology people running the root server system, we said great, let's stop adding them. The political reality of that is nuts, but if you've met engineers, they don't think that way.

And so, again, this is not an excuse, but it's an explanation. I think the reason is very much not we want to shut anybody out. The additional growth was always technical, and when we got to Anycast, Anycast is magical. It just means everything just worked. So, I don't know how this is going to end up getting solved politically, but I think you're right, it has to.

SUSHIL PAL

I think Anycast, I guess, that answers the latency part, it doesn't answer the resilience part, right?

DAVE LAWRENCE

It does, if I can add. And so, like, one of the interesting things, root servers provide a map of all of the instances around the world, and pretty much everybody sitting in this room has root server instances in their country. And in part, one of the most widely distributed, so again, there are 13 letters for 12 organizations which represent 26 different addresses over IPv4 and IPv6.

Jeff's root server instance actually will happily take anybody that requests in order to host the servers in their country, including then it is no longer directly under Jeff's organization's control. It is under the control of the ICANA root system and the distribution thereof. It's running in an Indian data center, but it's running with Indian administrators. And I see you shaking your head that this is not a satisfactory answer, but the resiliency part is well addressed by widely diverse geographical and network topological distribution.

NICOLAS CABALLERO

Thank you, India, for the question. We still have five minutes. Any other questions? Because I do have a question, but I would like to give you the floor first. I don't see any hand online. Anybody in the room? There is one online. Oh, the European Commission. Please go ahead, Gemma.

GEMMA CAROLILLO

Thank you very much, Chair. Thank you so much for this presentation. This is Gemma Carolillo from the European Commission. And of course, I'm not going to repeat what the colleagues said, it is very reassuring always to hear how the system is incredibly resilient and redundant in the way this is organized.

But of course, you have a job in advising the Board and the community about the security of the system. Any trend that you have seen, we are not looking for something alarming, we are looking for something informative. For example, linked to the use of emerging technologies, AI, in the types of attacks that you can see over the system. Thank you.

JEFF OSBORN

We were just talking about this the other night, and one of the funny things in ICANN is one of the questions they always like to ask is what keeps you up at night. And Kurtis asked me that the first time I met him in an official capacity, and I kind of stopped and thought about it and I said paying for my kids college bills. And he said, no, no about the root server system.

And I said it doesn't keep me up at night, paying my kids college tuition bills keep me up at night. So, we are a bunch of engineers running a system and we're pretty proud of how we've done it and we recognize there are political realities and we are not politicians. So, a group was spun off that is attempting to address the

governance thing, and I'm really glad to be able to point it to them because what you're saying is so obvious from a political point of view.

It's like talking colors to a colorblind person. I do engineering. I can tell you how to pay for open-source software that we don't charge for and still maintain payroll. I can do some things really well, but politics is a complete mystery. I try to vote every year and that's about it. So, that's why I get to throw it at Dave. I'm really happy about that.

DAVE LAWRENCE: Well, and again, the Governance Working Group did just complete its original mandate and has made a series of recommendations to the Board as to what the next phase of attempting to set up a governance system would look like. And then that Governance Working Group's remit would be to consider the additional political issues.

NICOLAS CABALLERO: Thank you so much again, Dave and Jeff. I have the U.S. first and then India, right? And we have three more minutes. So, U.S. first and then India. Sorry to keep you waiting, Sushil. U.S., please, go ahead.

SUSAN CHALMERS: Thank you, Chair. Susan Chalmers for the United States. Just three brief points. First, to iterate my colleague's observation from the

European Commission that it is reassuring to hear about the stability and the resiliency of the root server system that continues to endure after so many decades of serving the internet well. I would like to point out that it would be useful at some point to follow up on the cyber incident reporting work.

I believe the GAC would appreciate an update on that. And also, just to note that we did lodge a submission on discussions during the public, well, we did lodge a submission on the proposed functional model, but that we are still closely following the unfolding dialogue on this governance structure to ensure that the outcomes continue to provide for the root server systems long-term stability and security. Thank you.

NICOLAS CABALLERO Thank you so much. Jeff, would you like to say anything in that regard?

JEFF OSBORN Thank you. I'm glad to hear that this is being followed at that level. One thing I feel like I should have made the statement of is this is all volunteer work. My company is a nonprofit that is a small company that gives away software for the good of the world, and then spends millions of dollars operating a root server system. We don't get advertising from it, we don't get funding from it, we don't get subsidies from it, we don't get anything else from it.

And so, part of the concern I feel like is a bunch of people standing around it feels like saying why don't you do all of these things you're doing out of the goodness of your heart better. And what keeps me up at night? Money. Like I have to pay for all of this stuff.

This is a free good that is given by all of these organizations for the good of the Internet. And if that sounds naïve, it's not, it's genuinely what we're trying to do. So, if you're ever trying to get some government direction somewhere, funding would go a long way.

NICOLAS CABALLERO Thank you, Jeff. US, would you like to reply?

SUSAN CHALMERS No, just thank you. We appreciate that. The point is that it would be good to have some updates on the cyber incident reporting work for the GAC at some point, and that we look forward to monitor the discussions on the governance structure, which I appreciate that we did not have an update on that today and that you did not address that in your presentation, and that was not planned to be part of this discussion. Thank you.

NICOLAS CABALLERO Thank you so much again, USA. We're absolutely over time. I'll give you the last one, India.

SUSHIL PAL: Thank you, Jeff. I admire your frank and honest confessions, and also the clinical way you separate the political and the technical part. As a technical income, if you wear your technical hat, would you be open for capacity building in such a internet infrastructure to the other academic institutions?

JEFF OSBORN: I can't speak for everyone by definition. I can speak for my own organization. Literally all we care about is that we don't break this. This has worked for a really long time, I don't want to break it. That's literally our entire concern. I have the board of directors behind me, I have my organization behind me, all we want to do is not break it.

NICOLAS CABALLERO: Thank you so very much, Jeff and Dave, and thank you for your questions and for your meaningful engagement. We're going to have a 30-minute coffee break now, and the next session is going to be the GAC discussion on New gTLDs at 10:30. Please be back in the room at 10:30. Let's give a big round of applause to our colleagues from the RSSAC. Thank you so much.

[END OF TRANSCRIPTION]