MARRAKECH – How It Works: Root Server Operation
Monday, March 07, 2016 – 17:00 to 18:30 WET
ICANN55 | Marrakech, Morocco

[STEVE CONTE]: All right. We're going to go ahead and get going on our final session for today, and our final session for the How It Works in Morocco. I appreciate all of you who have been here the whole day. I know it's a long day, and you guys have weathered through it, so power to you.

This session is on root server operations. This will cover, oddly enough, root server operations, but also the Root Server System Advisory Committee – did I say that correct? – the RSSAC and talk a little bit about what they do within the ICANN structure and what they're working on now.

With us today, we have Jim Martin from F-Root and Wes Hardaker from B-Root, and they're going to go through that. Without further ado, I'm going to pass it onto Jim.

JIM MARTIN: Good afternoon, and thanks very much for showing up. It's nice to see some technical content at ICANN, and I'm glad that you all chose to come join with us.

We've got a couple different sections here. I'll do some. Wes will do some. However, if there are questions along the way, it really does help to understand the earlier portions before we go to the later portions, so please do speak up and ask the questions as we go through.

Also, as we go through, we're going to make an assumption that there's a very basic level of understanding on some of the technical topics. Please, I apologize in advance to those of you in the room that are going to be a little bored, but it's best to have everybody have a baseline understanding.

As Steve just mentioned, both Wes and I are from the Root Server System Advisory Committee (RSSAC). The fact that it was written up there is a good thing because I can never remember what the expansion is. I'm so used to just calling it RSSAC.

To get started, can we have anybody who's on RSSAC – not the Caucus, just RSSAC – raise their hands right now? And that's members of the various root server operators, as well as some liaisons. But RSSAC has a Caucus as well, which is technical members of the community. Who is part of the RSSAC Caucus? You are? Then there's a third piece of this. They're the root server operators. Who in the room is a root server operator?

So it sounded like three very similar – and oddly enough, most of us had the same hands – but different organizations that have

slightly different purposes. We'll go through all that as we go through.

To get things started, what we're going to be doing today is we're going to talk about what the domain name system is to begin with, just an overview – a very, very lightweight overview of the whole domains name system – the history of root server operations, how root servers came about and how they've grown over time, the root server system today and all of the features and things that we're doing with the root server system as it stands right now, and then what RSSAC is doing, in conjunction with the RSSAC Caucus, moving forward.

We've got to start at the very most basic concepts for the DNS system. The most basic identifier on the Internet is an IP address. IP addresses come in a couple of different flavors – IPv4 and IPv6. They tend to be the unique identifier for a host.

However, there are some tricks you can play be assigning the same address to multiple hosts. We'll get into that in a bit.

The way these are allocated, the uniqueness is all handled by the RIRs, the Regional Internet Registries, and they are handed down from IANA, which is part of ICANN.

The original idea behind the DNS was we all had these things called hosts files. They were a file that listed an IP address on

one side and host name on the other. If you wanted to go to host XYZ, it would get looked up in the file, get you an IP address, and that IP address is what your computer would connect to.

That worked great when there 50 devices on the Internet. Unfortunately, that was never going to scale, so a system called the DNS was developed to be able to handle that translation. That was the initial problem.

However, over the years, DNS has changed substantially. It still does that direct translation. However, in the modern environments, you're also doing things like geolocation, getting slightly different answers for different queriers, that sort of thing, to give additional capabilities to the DNS to help expand that we're using the Internet today.

The domain name system itself, as I said, is a lookup mechanism. It's a way to translate from names to numbers and numbers to names, and it's done as a widely distributed database. It's actually a really interesting example of a widely distributed database that has no central point of control. Each of the nodes is managed on their own and yet all hangs together because it's in the node owner's best interest to keep their data current.

The first concept was the mappings from names to numbers, what refer to as A records. You're going to get

www.example.com, or rather just www, translates to 1.2.3.4 and vice versa. You can go back from the name to the number using something called a PTR record.

The actual process for doing a lookup is you've got an end user who asks something called a resolver, a caching DNS server, to lookup www.example.org. We're going to assume for a moment that this resolver was just brand new, just got turned on, and knows nothing. So it's going to go off and it's going to ask the root for where .org is, in this example. The root's going to say, "Go to the .org server at 1.2.3.4." You then ask the org server, for example, and it tells you where the IP address is, for example. Then you go ask example what the record for www is. So you can build the address by following all the way down the tree, starting from the back end of the address, working to the front.

The important part, though, is that that whole setup only has to happen once. The first time you do that, the caching DNS server will remember what root had said about where org is and what example said about where www is. So you've got all that information being kept, so the next time you ask that same question, the server can answer immediately.

Most of this talk is about the root name servers. The truth is that the root name servers, the actual root zone, is a tiny zone. All it needs to know is what the ccTLDs and the gTLDs are and the

other TLDs, what servers are serving them, and then all of the metadata associated with that.

When you ask the root for com, it will give you the list of com servers. That's the end of what the root knows. It's a very, very limited in scope role. However, it is the apex of the distributed tree, so if the root isn't there, then we've got a problem. The whole lookup mechanism fails. That's the reason we go through a lot of effort to ensure that the root is always there and always performing.

As I mentioned, throughout the entire DNS system, there's a consistent process of caching. The caching, which then over time will time out the values – it remembers things for a while and then will eliminate them over time – ensures that you can optimally get responses – reduce the amount of time to get a response – while still ensuring reasonable freshness of the data.

So DNS, as I said, started as a simple names and number mapping. However, over time, we've added some refinements to it. DNSSEC is the one that is often spoken of these days. DNSSEC is a mechanism to cryptographically sign your DNS records to ensure that you can validate all the way from the root down to the leaf record you're talking about. In fact, it was being served by the entity that is allowed to serve that record.

It's very important for things, like if you go to www.bank.com and you really want to know that it's your bank that you've gone to, the first step is to make sure that the address that you're told to go to when you do that lookup is legitimate.

Something that's in process right now is the concept of privacy enhancements in the DNS. Right now, when you make that request for www.example.com, the whole name gets carried to everyone that you make that request of. If that name is something that you may not want everyone to know about – if you've looked up something with bankruptcy in the name, you may not want everybody along the path to know that you're really all that interested in bankruptcy.

So the idea of reducing the amount of private information that is being sent in the DNS is an important part of what's being worked on right now in the IETF.

Then perhaps most fundamentally to what we do with root server operations is the concept of Anycast. The concept of Anycast allows you have to have many servers all answer for the same IP address. You end up using the fundamentals of the routing system to deliver the packet to the name server, rather than things like load balancers and things like that.

[ADEEL SADIQ]: This is the Adeel, a NextGen member from Pakistan. Could you please re-explain that privacy enhancement?

JIM MARTIN: Sure. Privacy enhancement involves reducing the amount of information that is being carried in the DNS request so that when you ask the root, if you're asking for www.example.com, the root doesn't need to know about the www or about example. They just need to know that you need com so it can give you the information there.

There are also additional mechanisms that involve securing the path between your client and the first-hop resolver. There are many additional things to this, but the idea is, for many years, there was no concern about the information in a DNS query being something that could leak information about what you're interested in. Now there's some work actively working in the IETF to address that.

I just want to clarify some important concepts here as we go through this. There's the root zone and there are root servers. The people involved are very different and the way it's managed is very different.

The root zone is a text file. The root zone is a list of the TLDs and all the metainformation for them and the addresses of the name

servers for them. That's all created and managed by ICANN from inputs from all of the TLD owners, and there's community policy on how that's managed.

That then is handed off to Verisign, who actually goes through a number of checks, puts it into the appropriate format for all the software, and then they hand it off to the root server operators.

The root server operators are organizations that can be unaffiliated with ICANN. One of them happens to be ICANN, but all the rest of them are unaffiliated with ICANN. We all – the root server operators – all provide this service. What we mean by that is we're actually running the servers that are out in different places on the planet, answering the queries. When you send a packet to a root server, it goes to one of the root server operators' root servers.

There are currently 13 identities, all lettered A through M. Those are all a purely technical role. The thing that's most important to keep in mind is that the root server operators and the root servers have no input into what goes into the root zone. The root zone is managed by ICANN.

The root server operators are 12 different organizations. There are 12 root server operators. That's because one of them, Verisign, happens to run A and J.

The things that the root server operators are focused on is making sure that the name service is available everywhere in the world as quickly as possible, at the lowest latency as possible, and as reliably as possible because this is such a key piece of the way the Internet works.

We ensure that we answer queries independent of who you are, whether you are the best person in the world or the worst person in the world or politically correct or incorrect. If you send a query to us, you get a response.

The root server operators work together in technical cooperation, but we are all very separate individual entities. We're made up of diverse organizations and operations, meaning literally we all coordinate to ensure that we don't duplicate things. We ensure that we aren't all running the same software, all running on the same hardware. We ensure that we don't use the same Internet transit. We don't use any commonality so that one single thing failing doesn't ever impact all of the root server operators.

Organizationally, we're spread out all over the map – oh, I'm sorry.

[STEVE CONTE]: Can we use the mic, though? I'm sorry. We only have that one mic.

JIM MARTIN: Well, I will wander over.

UNIDENTIFIED MALE: Who makes the changes in the root server currently? ICANN, right? As a maintainer?

JIM MARTIN: Changes to the root zone come through ICANN. There's a whole process to it. We'll touch on some of that in a minute.

The changes to the actual zone are handled through the community-driven process that goes through ICANN and then through to Verisign, where it is actually signed and then handed off to their RSOs, to the Root Server Operators.

Here's the thing that we want to make as clear as possible. The root server operators, the RSOs, they don't make the decision of what's in that zone. There's no policy happening at the root server operator level, and there's no modification. We all receive exactly the same zone, and we all publish exactly the same zone. So we don't have any input into the inside.

What we do do, though, is ensure that the service itself runs. We do also make recommendations back to ICANN if there are things like technical details, things like the TTLs in the zone or the way things are being signed. Anything that is technical in nature, we can make recommendations back about what's happening in the zone, but we don't control them.

The thing that we spend more time than anything focusing on is robustness and stability. The idea is that this must work 100% of the time for 100% of the planet.

With that, Wes will tell you where we came from.

WES HARDAKER: All right. Thanks, Jim. Again, my name's Wes, and I'm going to be going over a couple slides on the history because, just like the Internet sort of expanded over time and got bigger and bigger and bigger, the number of requests going to the root servers, of course, went up and up and up. So it grew in size and in feature set.

What am I pointing at?

UNIDENTIFIED MALE: [inaudible]

WES HARDAKER: There we go. The very first root servers existed back when the DNS was created back in 1983 through ’86. There was four initial ones. There was only four. You can see the addresses there. They were each running a couple of different pieces of software. JEEVES really isn’t used anymore, but it was back then. There was one, SRI International. There was actually two at ISI in southern California, and then there was Ballistic Research Laboratory, which is run by the U.S. Army. Again, the Internet sort of grew out of an ARPANET project way back in the day in the U.S. government, so a lot of it was based in the U.S.

Then, in 1987, the list of servers grew. One of the things that you’ll notice is that we grew by four. The red ones are the numbers that grew, but there used to be four. One was actually removed. If you notice, there is ISIB and ISIC because ISI had two. In this transition, now the second line is the only one that ISI had. So there was also a transfer that sort of happened there. Now there was seven, with four more being added, one at RPI, one at the University of Maryland, one at the U.S. Airforce Networking Group, and then one at NASA.

You can see, again, more addresses and more software being used. There was more of a diversity of four running BIND and three running JEEVES.

The names you'll notice on the left are going to become important in a minute, and I'll talk about those. Originally, notice the names had no dots. They were just single names. In this change, we also added subdomains. They became names that were internal to the organizations that were running them.

Then in 1991, the root zone was actually expanded outside of the U.S. This was the beginning of the transition to make it more of a global service. One was added in NORDUnet, which is in the Nordic region of the planet.

Then in 1994 and 1995, one of the biggest problems that was being recognized was that the size of the root hints response was approaching 512 bytes. There's a number of problems that existed then with 512 byes. One, some software couldn't handle it. Two, there was firewalls that were dropping anything over 512.

To add more root server instances, we had to figure out how to get that response under 512 to make it work with the software that was out in the world today. So Bill Manning, Mark Kosters, and Paul Vixie devised a plan to rename all of the root servers to all be in a subdomain under rootservers.net, and they each got individual letters. That still exists today. Back then, it was A-I.rootservers.net.

The reason for this – and it was finally approved in 1995 – is that it allowed name compression to take place. Before, when we had all these names with different ending suffixes, we had no compression. By adding rootservers.net to the end, all of that got compressed down into a single instance of that label being propagated with all the letters.

This is the result of the transition. The original name is shown on the left – quite a bit longer. Then you can see the new names, A through I, at the time. Again, the organizations didn't actually change.

Oops. It does like to jump, doesn't it?

Then later, in '96 and '98, Jon Postel, who was the creator of the original infrastructure that later became IANA under ICANN, he used a set of criteria to select new root server operators.

This is the criteria that he had devised. One is need. Where on the planet were new servers needed? What was the connectivity, both internal and external? Was it large enough for capacity purposes, and was there a commitment to respond to those requests without any filtering?

As Jim said just a few minutes ago, one of the things that the root server operators do is that we don't distinguish between a good query and a bad query. We try to answer everything that

we possibly can because it's not our job to do any censorship or anything like that. We serve the zone as-is to anybody that asks.

And then, is there a community consensus? The potential operator should demonstrate the widest possible support for the community being served. Again, the job of the root server operator is to serve the whole world, period. No exceptions.

So four additional servers were added. In Europe, RIPE was chosen to run K-Root. In Asia, WIDE was chosen to run M-Root. J-Root stayed at NSI, and then L-Root was transferred to ICANN as part of the founding of ICANN. L-Root is now run by ICANN still.

The root server operators also met as a formal group and agreed to the principles after Jon Postel's death – five principles. Excuse me. They agreed to operate for the common good of the Internet. Again, this was something that was sort of already decided before, but it was written down.

IANA was the source of the root zone data. That's very important because one organization has to be the source. If each of the root servers were serving different data, it wouldn't work, so all of the root server operators agreed at that time that IANA was the source of the data.

And that there was sufficient investment to operate responsibly. It takes a fair amount of infrastructure and a fair amount of

personnel in order to really responsibly run a root, so they all agreed that they had to have a sufficient investment to operate responsibly, and that if they couldn't, if something was going to happen, they would give proper notice and facilitate the transition if there needed to be one.

And then, finally, recognition of each other. All of the root zone operators at that time, they weren't corresponding a huge amount, so they actually started interacting on a much more frequent basis, and they recognized each other as other valid root zone operators.

I'm going to turn it back to Jim, who's going to talk about the root server system today and its features.

JIM MARTIN: All right. Thank you very much for telling me where to point because I was just guessing before, so that helped out.

As Wes had mentioned, he's brought us up to today. This is what the root servers look like today. We've got A through M. You'll note, though, that the IP addresses is that we've got both IPv4 and IPv6. Almost all of the root servers are now publishing both on v4 and v6. It's a set of diverse organizations that are corporate. There are universities. There are governmental agencies. There are non-profits. Where we are today was the

outcome of all of those diversity goals that had been established before.

There are nearly 500 or around 500 instances around the world with the 12 operators, 13 instances, 13 letters. The 500 instances are all using this Anycast mechanism that I described earlier, the idea being that any of these 500 servers will all answer the same queries with the same data. With the diversity of providers and diversity of locations, there's a very high probability that any time you send a question, you will get an answer.

This is meant as flow chart to give you an idea of how requests come in and flow from a – let's say .example wants to change the name servers associated with their TLD. The TLD operator would send a request to IANA to do so. There is a sequence that involves IANA, NTIA, and then back to Verisign to verify it, to approve it. Note obviously this is all right this second. That area right there has been the last several years of discussion right there.

Ultimately, though, the change makes it to Versign. They sign the zone, and then you see that wavy line in the middle. Everything to the left is happening in ICANN in the policy space. Everything to the right of the wavy line are the root server operators, the actual publishing of the zone. It goes from the distribution masters out to the distribution masters for each of

the letters – and we all have slightly different ways of doing our distribution – and then they hand off to all of the individual servers all over the planet to actually get the latest version of the zone file.

The yellow on the right hand side, that's where all the queries come in from the real end users coming in and asking questions. That's how the whole thing lays out.

I keep hitting on this – oh. Yes, sir?

ADEEL SADIQ: Adeel again. What exactly is a distribution master? Why doesn't Verisign distribute it themselves?

JIM MARTIN: I'm sorry?

ADEEL SADIQ: Distribution masters. Can you give an example of distribution masters?

JIM MARTIN: These are systems. These are pieces of hardware, computer systems that receive the zone. There are several distribution masters in the process. There's the distribution master on the

Verisign side that hands it off to the root server operators. The root server operators all have some sort of distribution within their own infrastructure that they get to define however they want, but those distribution masters are going to be the entry point into the RSOs' infrastructure.

UNIDENTIFIED MALE: [inaudible]

JIM MARTIN: Sure.

UNIDENTIFIED MALE: [inaudible]

[STEVE CONTE]: All right. Mic, please.

UNIDENTIFIED MALE: The root servers that the US government [inaudible] queries if we send to the root servers which are on the U.S. department. It's for the public use or it's for the government use?

JIM MARTIN: All of the ones we've just been talking about there are for everyone on the planet. No one is permitted within this system of having their own little private root for their own use, be it the U.S. government or the small shop down the street.

However, people have set up things like that, but this is the public system of root zone distribution.

UNIDENTIFIED MALE: [inaudible]

JIM MARTIN: The question was: some of the RSOs – the addresses on them – some of them were not v6 enabled, and the question was why are they not v6 enabled?

It's because each of the root server operators manages their own infrastructure and makes their own technical decisions on how they operate. That's the set of choices they have made.

UNIDENTIFIED MALE: [inaudible]

[STEVE CONTE]: I'm sorry. We have remote participants, so we really need to have the conversation mic-ed.

UNIDENTIFIED MALE: I'm sorry.

[STEVE CONTE]: So either come to the table if you want to have the dialogue, please. Thank you.

JIM MARTIN: Ah, perfect.

UNIDENTIFIED MALE: Hello? My first question is, why did we not have IP [inaudible]? Second question is, is they are also having a right to not to answer the queries?

JIM MARTIN: No. As we mentioned before, first of all, again, why they don't do v6: each root server operator has to make their own technical decisions on their own. With regards to it not answering queries, part of the commitment that each of the root server operators has made is that they will answer any queries that they receive.

WES HARDAKER: It's worth noting that the transition where the root servers had v6 did not occur overnight where all the rest of them occurred. It

was a slow propagation as each root server had the ability to add their own v6 address.

If one or two of them don't, it actually doesn't prevent you from asking the rest of them because when you're looking up the list of name servers that do v6, you're only going to send questions for those. So you'll only send it to the other 11. You're not going to try and send one to one that doesn't exist. You're only going to pick out of the 11 that do support it. So it actually shouldn't affect you greatly.

UNIDENTIFIED MALE: If some of these main 13 root servers stop working, what happens to the mirror?

JIM MARTIN: I'm sorry. Say that again.

UNIDENTIFIED MALE: If something happens to the main root server, like from A to M, what happens to the mirror that belongs to that particular root server?

JIM MARTIN: All of the root servers are –

UNIDENTIFIED MALE: I'm talking about the main one, like [inaudible] the mirrors. No one is the owner of [anyone], but the root servers, which is the original ones, if something happens to that particular root server –

JIM MARTIN: JEEVES.

UNIDENTIFIED FEMALE: An instance. [inaudible] happens if an instance is [inaudible]

UNIDENTIFIED MALE: Instances. Mirrors, which are located all over the globe.

UNIDENTIFIED MALE: [inaudible]

JIM MARTIN: Sure.

JOHN CRAIN: John Crain with L-Root. Later we're going to get some myths, I think, in the slides. I think you're on one of them. An instance is an instance. All of the L-Root instances are the same.

Now, if something was to happen to our distribution system, that may causes issues with L-Root, but there's plenty of other root servers. So the idea that there is one primary, main L-root and the rest are somehow different is just one of these myths. They're all exactly the same.

So if one of them falls down, we've got 100 more.

UNIDENTIFIED MALE: Actually, in India we did some labs and we tried to query root servers which are located in India. The root servers which are ICANN's root servers were not solving that query in time, actually. So I don't know actually what's happening at [inaudible], but at the [Dahua] Technology, we set up that thing at our office and tried to query. We analyzed, a few of the root servers are not properly responding. I don't know. We've tried to figure out so many ways, but still we are not able to understand. That's why I came here to ask.

JOHN CRAIN: Oh. So that's probably more likely to be routing issues, but if you've got data, and specifically to L-Root, come talk to us and we'll help you figure it out because all the systems are responding to us from a management perspective. So come talk to us with the data and we can work with you.

UNIDENTIFIED MALE: Okay. I just wanted to ask is there's any policy issues or something because these root servers are not responding? Or it's some technical issue?

JIM MARTIN: I believe it's purely a technical issue. All of the root server operators are very interested if there are failure issues, so certainly do reach out.

However, I think Warren was – oh. Go ahead.

WARREN KUMARI: Sorry, Warren Kumari. I help run F-Root. One of the times when a root server might not respond is it thinks it's currently being attacked. So under certain attack scenarios, just to keep the server up and running for a short time, people will block specific IP addresses. But that's not a policy thing. That's a "Help! I'm being attacked right now. Let me try to save myself." It shouldn't be that, but that's one technical possibility.

UNIDENTIFIED MALE: One last question, the question [inaudible], I guess. What is the exact way, like the new gTLD process? If we make new entries in

the root servers, what is the exact process to make changes in the root zone or adding a new [inaudible]?

JIM MARTIN: Go for it.

ELISE GERICH: The process is the same for all TLDs. It doesn't matter whether it's a gTLD or a ccTLD. The requests come to the IANA department. The IANA department sees if the request meets the criteria that's already been sent and created the process and it's published, at which point, once a request has been verified and confirmed, then the IANA passes it first to the NTIA, who says yes. They've never, ever turned down a request, but that's what we're all here this week talking about: the potential removal of NTIA from its role.

Then it goes from NTIA to Verisign because then Verisign, on NTIA's authorization and ICANN's confirmation that it met all the criteria, compiles a new root zone. So the gTLDs follow the exact same process as the ccTLDs.

UNIDENTIFIED MALE: At the same time they make changes to the all 13, or one by one or something?

JIM MARTIN: Yeah. Absolutely.

[STEVE CONTE]: I'm sorry. Really quick. Elise, for the room and for the record, can you just –

ELISE GERICH: My name's Elise Gerich from ICANN.

UNIDENTIFIED MALE: My name is [inaudible] from [NIXI], India.

JIM MARTIN: Just jumping back to here, just to be clear, the process that Elise just described is all of the colored blocks to the left of where it says DM –

UNIDENTIFIED MALE: I know the left, but I want to know this part.

JIM MARTIN: Right. Once that's all handed off, the box you see is the single DM. Obviously, that's an abstraction. There are a few of these boxes. They hand off to the distribution infrastructure.

I can use F-Root as an example, to speak for myself. F-Root has distribution masters on many different continents that will then receive copies of the zone from Verisign. We then, from those continental distribution masters, send them out to all of the instances that we operate on the –

UNIDENTIFIED MALE: At the same time to the all 13? Or one by one?

JIM MARTIN: Well, again, I'm speaking for my letter. Think of it as a waterfall process. It comes out of Verisign, is handed off to each of the letters, and they do their own individual distribution.

Part of the concern you're asking is how long does it take from the beginning of the process until it makes it out all the way to the edge? We're going a bit afield here, but there's a document called RSSAC 002, which is RSSAC's technical specifications document, asking for statistics out of the root server operators.

The value you're asking for, which is how long does it take to come out, is being published by many of the root server operators now, so you can go and look at that data on www.rootservers.org.

UNIDENTIFIED MALE: Okay. Thank you very much.

JIM MARTIN: Certainly.

I think I was most of the way through here. The main thing, though, is, though we are fundamentally diverse, we do share some common best practices. It's some simple things, making sure that we have physical system security and that there are no root servers that are living under someone's desk, and the over-provisioning of capacity. We certainly never design around what a nominal traffic load is. We always design around the attack traffic loads.

There is a very close-knit set of staff that actually professionally operate these all.

We do a lot of cooperation, though, though we are diverse. It takes place at industry meetings. A lot of us tend to end up at IETFs and RIPEs and NANOGs and OARCs and all these sorts of things to be able to have interactions. We do have a permanent infrastructure to handle whenever there are crises. So if there's an attack or there's something going wrong, there's an infrastructure to ensure that lots of people's phones will ring and there are mailing lists for communication. We make a point

of exchanging secure credentials so that we can verify each other in time of some event taking place.

We also do periodic activities – what we call tabletop exercises – to ensure that we've run through the exercise of having an attack take place and making sure that all of our mechanisms are operating appropriately.

We also definitely have coordination with other Internet bodies. Within ICANN, the entity there is RSSAC, of which a number of us are members. On the actual DNS standards, that happens at the IETF. Also, through DNS-OARC, we use that as a common platform for distribution of statistics data.

So we definitely also are a dynamic environment. We certainly don't just serve things exactly the way we used to. To your question of IPv6, that has been something that has been added over time. Similarly, IDNs and DNSSEC, these are all things that have been added over time. So as the Internet changes, what we serve and how we serve it changes.

The one thing that is consistent is that we're always working towards increasing the robustness and responsiveness and resilience of this system. There are always root letter operators, RSOs, that are interested in putting out more nodes that want to have more resilience in this system by having more capability out there.

The root servers. There are a huge number of myths associated with it and misperceptions of how we operate. The root servers don't control where the traffic goes. The routers all control the way the traffic flows. Certainly, not every DNS query goes to the root server. It's not like, if the root servers went away, that every DNS query would fail. It's simply that the queries for things in the root would have a problem.

The caching ensures, in fact, that the actual query load on the root servers is relatively minimal because that state is being kept at all levels of the system.

The administration of the root zone is completely separate from the provision of the server. While in general all the root server operators are running the service but not having anything to do with the zone, the one exception is ICANN. If I step on your toes, please yell at me, but there's a very distinct firewall between the side of ICANN that serves the zone and handles the packets coming in, and the policy side and the process associated with generating the zone.

And A is not the best letter and M is the worst letter. They're just letters. There's no preference on any of them. In fact, if you happen to be running – I can't speak for all of them; I can speak for BIND just because I'm familiar with it – the BIND DNS server software goes through and tries to figure out how long it takes to

get an answer from each of the various letters and will keep track of which one is most performant for you and will keep asking that one. So it may start with A or it may start with J or whatever one it does, but it will always bubble up to use the server that is closest to you.

We're not hobbyists. There's a lot of money and a lot infrastructure built to make this all happen in data centers all over the world. This is, again, not sitting under someone's desk.

It's astounding the number of people that are convinced that there are 13 boxes around the world. There's A and there's J and there's K. They're just identifiers that all map back into Anycast, and Anycast takes you to servers that are all over the planet. Often, when you get to that site for the letter, you find a number of servers that are running, when you come down to actual physical hardware. There are huge numbers of devices out there.

The key thing is that, while we are very diverse, we definitely do coordinate our operation to ensure that the ongoing operation of the DNS root is consistently excellent.

To wrap this up …

UNIDENTIFIED SPEAKER: Question.

JIM MARTIN: Actually, give me a minute. Let a new person speak for a little while.

UNIDENTIFIED MALE: Yeah, sure.

UNIDENTIFIED MALE: Yeah. I have seen practices where one IX is having two or three root servers, like other letters. If ten or 15 ISPs are appearing [inaudible] which root server is going to answer? Because they are the same hop count. So is it beneficial? I've seen there's lots of IXs are having two or three root servers.

JIM MARTIN: Sure. Remember that the Anycast is used so that if you send to the F-Root address, you will go to the closest F-Root instance. But it's not that the different letters use the same address. So the address for F-Root, when you send a packet to F-Root, is going to be different than the address that you use for B-Root.

So it's not an issue of hop count. It's a matter of your software making a decision of: is it going to ask A or B or C or D? That was where there is some optimizing mechanisms that are within your software that allow you to use the best one.

To that end, they may all have nearly the same performance. Having multiple in an IX probably doesn't decrease your latency at any given time. However, if one fails, the other one can still continue to operate. So that's how the resiliency advantage takes place.

UNIDENTIFIED MALE: So you recommend multiple root servers [in an IX]?

JIM MARTIN: I always recommend more than one root server. It ensures that you have consistently high availability because, again , the thing that takes out my systems is not what's going to take out Wes' systems.

UNIDENTIFIED MALE: Thank you.

[ADEEL SADIQ]: So if I send a query to, for example, Google.com, will I be returned an IPv4 address or IPv6 address? Is there any preference, and if there is, it's based on what? [inaudible]

JIM MARTIN: If I understood your question correctly, you're asking: is there a preference for making a query over v4 or v6?

[ADEEL SADIQ]: Yes.

JIM MARTIN: Wow. You're getting definitely down into questions of how the software on your client is written – not just your client, but also the operating system you're running on and that sort of thing. From the perspective of the root server operators, they're treated equally. There's no difference between IPv4 and IPv6 from a service perspective. Nothing is preferred.

As far as I know, that is the case. I can tell you that on F there's no preference between any of them.

As you might note here, I keep saying that. While we do certainly coordinate, we all are very careful about not speaking for each other. Each entity speaks for itself, so while many of us have a pretty detailed understanding of each other's infrastructure, we'll never speak for the others.

So if you have a specific question, you always would need to want to speak to each of the RSOs.

UNIDENTIFIED MALE: Maybe I can answer as a software developer. For the resolvers, there are different ways to deal with that, but there's something called in the IETF Happy Eyeballs. If there's an IPv6 address available, they randomly select IPv4 or IPv6, or they give IPv6 a small head start of a couple of seconds. If they don't receive an answer, they ask IPv4. But that's from the open source resolvers I know of.

JIM MARTIN: I'm not leaving. There was somebody back here who had a question.

UNIDENTIFIED FEMALE: Thank you. In the past, concerning the [corrected] myths, you said that not every DNS query is handled by the root servers. I'm wondering what types of queries are handled.

JIM MARTIN: Again, just to be as clear as I can, using the example of www.example.com, what will happen is that, if you're a resolver, if that first hop that you make a request to has no information, it will ask the root server for where .com is. The root server tells that resolver, "Go talk to these servers to answer all your questions about .com." That's all that goes to the root server. Or .ca, if you're asking about something a Canadian site, or .cn if

you're asking about a Chinese site, or .gTLD for one of the new gTLDs. It's just that very last label, the thing all the way to the right in the name, that is answered by the root server.

Now, that is what the theory says. I can tell you that the root servers get a lot of queries that have nothing to do with just knowing what's that last label. But the design of the root server is just to answer: what is that last rightmost label?

UNIDENTIFIED MALE: What we are doing in the lab is we've tried to restrict our software to query only three root servers, which are located in India. And we achieved it, somehow, there. So my question is, I'd really to know, if we are having only three instances, like A, B, or C, is it wise to go for the other instances? Because if my software is working on the probability, like if it's [granting] all the 13 root servers at the same time, if in my country we are having only three root servers in my country, my probability of getting a response reduced to the 25%. Are you getting my point?

Even I'm not clear. I'm not making any statement. Even I'm not clear. It's working on the probability or RTT, like round trip time, it calculates on what basis it decides to whom to query which.

JIM MARTIN: Again, what you're asking there are implementation questions on specific name server software. I can tell you, again, some of the operations on BIND –

UNIDENTIFIED MALE: No, no. My question is, is it wise to install all the 13 root servers, or only –

JIM MARTIN: Yes. Most certainly.

WARREN KUMARI: Yeah. Warren Kumari, F again. Yeah, it makes sense to list all of them because most software implementations will automatically try all of the servers and will figure out which one is fastest. So it will query A, B, C, D – all of them – and it will learn which one is faster and it will use the fastest one. Every now and then, it will try one of the other servers, in case one of the other servers [inaudible] –

UNIDENTIFIED MALE: It is not possible for a developing country to install all the 13 root servers in the country.

WARREN KUMARI: Well, does somebody else have a – yeah.

JIM MARTIN: Liman would like to jump in. I'd like to jump in – oh, I'm sorry, Benno, [you did say]. Sorry.

BENNO OVEREINDER: If you limit yourself to only root instances in India and people know that, you make yourself vulnerable also for an effective DDOS.

JIM MARTIN: Exactly.

UNIDENTIFIED MALE: So there's no [inaudible] then because normally the routing system would take over [inaudible] –

UNIDENTIFIED MALE: No, no. I'm not in favor to restrict for the three. I am just asking a simple question. Is it wise to install all 13? If I contacted to ICANN and ask them to install 13 L-Root servers in India, or I should ask all the 13 root servers operators to install there? What will be the wise choice? Which will help me do –

JIM MARTIN: Diversity of name server operators will certainly aid you in the case. Again, you want to envision the scenario – we all have the same data. We all publish the same data, so you're better off ensuring that you've got different operators that have different software sets, different administrative controls, different approaches to solving things. So envision the scenario of there's a major catastrophe on the Internet. You want to ensure that you have the greatest likelihood of finally continuing to operate than if were to have yourself locked down to one root. If that root goes bad, you're done.

UNIDENTIFIED MALE: Thank you.

JIM MARTIN: Liman, you wanted to say something?

LARS-JOHAN LIMAN: Yes. I'm Lars Liman from Netnod. We operate I-Root. Why are you so focused on inside India? Because if the service inside India breaks, the ones outside India still work, so the traffic will go outside and then go back –

UNIDENTIFIED MALE: I know the traffic at the root server is very low. When it comes through the [inaudible]. I'm just asking a normal question. I'm not worried about the traffic or anything. I'm just asking, is it wise to install 13 root servers or the same root server? I know about the traffic. It's not much traffic on the root servers because once it comes through the caching, it remains for the 24 hours as per the settings for the longer time. Right?

LARS-JOHAN LIMAN: I would say it makes sense to have a variety. Maybe not all 13, but a few spread over a few sites so that you have a selection to choose from and so that there is redundancy within your region. India is a huge country with almost an infinite number of Internet users, so it definitely makes sense to have several inside India to support the community. Thanks.

JIM MARTIN: John, and then Warren.

JOHN CRAIN: You were also conflating various issues here. You were talking about your resolver configuration. You were talking about having servers in the country versus having access to letters that may be outside the country. So what you should probably do is

what Liman does. But apart from that, don't limit your resolver to just the letters inside the country.

UNIDENTIFIED MALE: Actually, I just [inaudible] limit resolver because if I want to do that, I have to develop a policy and I have to ask each ISP to query particulars. That I can do. That anyone can do as well because –

JOHN CRAIN: Okay. What you said was that you're resolver [inaudible] –

UNIDENTIFIED MALE: I tried that in a lab.

JOHN CRAIN: Okay. So do not do that. In the lab, it's cool. Outside the lab, bad idea.

JIM MARTIN: Warren withdraws, and I will hand off to Wes.

WES HARDAKER: It's worth noting that we keep diving down into technical details of the DNS, but the root zone is just one instance of a zone. So

we may talk about, "Does the root zone exist in political boundaries?" which doesn't make a whole lot of sense because networks don't go along map lines, and, two, it doesn't really help you if the root zone happens to be you have all 13 letters. I don't think any country has all 13 letters, actually.

But if you're trying to go to example.org, well, example servers aren't going to be in your political boundary, anyway.

This discussion today is really about how the root works, which is sort of independent of how DNS resolvers work and how they most effectively ask org servers and example servers. So you need to get distribution across everything.

I'm going to go onto RSSAC and what that is within ICANN – that's a special body within ICANN; it's an advisory committee – and the recent activities of RSSAC within ICANN.

The role of the Root Server System Advisory Committee – and I forget what that acronym spells out all the time, too, gentlemen; I'm glad it's not just me – is to advise the ICANN community and Board on matters relating to the operation, administration, security, and integrity of the Internet's root server system. It's a very narrow scope. Our only goal is to provide advice on a very narrow list of things to a small set of bodies.

What does it do? Well, RSSAC is a committee that produces advice, as I said, primarily to the Board but also to other bodies within ICANN and other organizations involved in the overall DNS business.

The Root Server Operators, or RSOs, are represented inside of RSSAC. RSSAC contains membership of which the root server operators have membership in RSSAC. However, RSSAC itself does not involve itself in operational matters. RSSAC has nothing to do with how the boxes run. RSSAC is only an ICANN entity involved in ICANN policies and things like that, and only giving advice. RSSAC actually doesn't make decisions on its own.

Where is RSSAC in ICANN's organizational structure? It is here in this little box. This is some version of the ICANN organization, which is changing all the time.

It's composed of representatives from each of the root server operators. Each of the letters has a representative and an alternate as well. Then there's a large list of liaisons, which I'll go over again in a different slide in a minute.

More importantly is actually the RSSAC Caucus because that's actually where the technical work gets done, not within RSSAC itself, but actually within the RSSAC Caucus. It's composed of a body of volunteers that have said, "Hey, I will help. I'm a subject matter expert. I know a whole lot about DNS. When you start

getting into nitty-gritty numbers and details and how things work, I can help." That's where all of the work is generally done.

The RSSAC Caucus members are appointed by RSSAC, and you can apply to be one. If you believe you have expertise in the DNS world and you want to be on the Caucus, please do so because we need the more the better.

The current RSSAC Chairs are Brad and Tripti, from Verisign and the University of Maryland. The Chairs change over time. Brad is actually the newest of them. Lars Liman at the back seat was the previous Chair that just ended his term in December.

The liaisons that I mentioned before: there's RSSAC liaisons to IANA, to the root zone maintainer, i.e., Verisign, to the IANA functions administrator, to NTIA, the IAB/SSAC, which is another advisory committee within ICANN, the ICANN Board, and to the NomCom committee.

The RSSAC Caucus, which is where I said all the technical work happens in the Caucus, where the meat of stuff happens, there's 71 technical experts within the RSSAC Caucus. 45% are actually not even root zone operators in the first place. They have to submit public statements of interest. You write down your history with the DNS, why you want to be a member, what your background is, and things like that.

Most importantly, the members get the public credit. If you look at the RSSAC publications, you'll find that it doesn't say, "Published by RSSAC." It says, "Published by the RSSAC Caucus" or "Written by the RSSAC Caucus."

The purpose of it is to be a pool of experts who produce documents that are expert-level documents. There's a critical mass of people that have read it and understood it, so there's wide technical agreement. And there's transparency for who actually does the work. Who actually helped author this document? What expertise did they have, and what were their hats that they were wearing? Are they root zone operators? Are they software developers? Any of the above.

There's a framework for getting things done. The RSSAC Caucus, as I said, is where the work actually gets done. So it's a framework for how that happens. The rules of operation and how long a document has to be are proposed and then worked on and then things like that.

If you want to apply, there's the e-mail address. It should be on the slides that you can download from the ICANN website, but it's rssac-membership@ICANN.org.

Recent RSSAC publications include RSSAC 001, which is the service expectations of root zone operators. 002, which Jim has mentioned earlier, is the advisory on measurements of the root

zone system. This is actually a fairly important one because it's actually about to be revised a little bit, as I mentioned, based on lessons learned. This is where the root zone operators actually publish their statistics for how long it takes to load a zone once they get notified by Verisign, all these different types of parameters, and how many requests they get on a rate kind of system. It's all kept historically, so you can go back and analyze the data for as long as that root zone operator has been publishing it.

RSSAC 003 is the report on the root zone TTLs, which is actually the most recent one. It was published in September of last year, and it's advice on how long the TTLs should be for stuff that's published in the root zone.

There's also been a number of statements. There was an RSSAC workshop 2015 report that reports on some conclusions by RSSAC from their workshop. There's an RSSAC commitment on the ICG proposal. There's an RSSAC comment on the CCWG Work Stream 1 report, and there's an IAB liaison to RSSAC statement.

The current work. One of the things that the Caucus is working on right now is documenting the history of the root zone system. You saw our slides earlier that document that. We're trying to publish a much longer and more detailed document for exactly

what happened with the root zone over time and things like that, where the transition points in time occurred.

It's going to contain a chronological history, as I mentioned, of the root zone system, from the very beginning to its current structure, and divided into historical periods. It's going to contain a description of the current operators and their histories involving the root zone. That's provided each root zone operator is providing that history.

That was it, right? Yeah. RSSAC 002. I talked about the metrics a minute ago, where all of the root zone operators have been working on publishing metrics. Like any technical thing, once you do something, you learn a little bit, and then you go back and revise it.

Well, we're undertaking some changes in those metrics and how they define because some of the numbers, when you actually start measuring it, are either nearly impossible to measure or every difficult or not as meaningful as we like.

So we're going back and here's a list of seven things that we're going to change. It's open to more interpretation. If you go look in the document and you need something else or it doesn't quite make sense, the RSSAC Caucus certainly wants to hear about it.

They contain things like whether the zone size metrics should actually continue. Is that actually an important element that people are actually going to care about? There's a whole other list that I won't go into detail today.

The other things. There's the root server system naming scheme. One of the interesting things: we talked earlier about that there's rootservers.net, which is where all of the letters current exist under. It's not entirely clear whether that should be maintained as a separate zone as it is. Right now, we have a.rootservers.net and b.rootservers.net. It's a question of whether we should keep that as a separate zone under .net, or whether it should sort of be rolled up into top-level domain names within the root zone itself.

There's some tradeoffs there. There's some security things to consider with respect to DNSSEC, where the information is stored and everything like that.

So there's going to be some analysis done on that subject itself as to whether those changes should be rolled up or whether it should be kept like it is, and if so, whether that sub-zone should be signed or not. Right now, it's actually not signed. There's some analysis that says that it shouldn't be signed, and some people think that it should, but it's never formally documented anywhere, and we're trying to fix that.

All right. That's it for RSSAC. RSSAC meets at ICANN meetings. And the Caucus, we meet once a year, sometimes more if more work needs to be done. There's a mailing list that is actually quite active.

If there is any questions about the entire presentation – we're at the conclusion of this presentation – about how the root zone works, we'd love to hear about more.

[STEVE CONTE]: [inaudible] both of you at separate mics, and then I'm going to run the mic around. So go ahead and speak.

UNIDENTIFIED MALE: Three questions. First, does [any] root server contain the data of only itself or all the 13 root servers? Second, for example, if somehow a root server is offline, what happens to all of its instances? They go down immediately, or there is a synchronization time?

Lastly, if a root server is offline permanently, how do you recover it? For example, it has been destroyed. All the data has been lost. How do you recover it?

WES HARDAKER: A couple of things. First, unfortunately, I missed the first question. I was shuffling things around, but I can answer the second and the third, and I'll let Jim answer the first.

If a root zone goes offline entirely, that's very, very rare because with many, many Anycast instances, they don't all go down at once. They're in physically different regions. They're designed so that if an earthquake takes out one machine, the rest of them still keep operating.

So as a root zone operator, I need to make sure that my system is diverse. I actually speak from B, where we actually don't have any Anycast instances. I'm the one letter that doesn't.

If I went down, which is possible if mine's in L.A. and if there was a huge earthquake in L.A. that took out B, you actually would not care. You would get a slight delay when you were asking questions to B, and we wouldn't answer you. Your software, assuming you use one of the ones that we've been talking about today, would stop asking B. It would realize that there was an issue, and you'd actually stop asking. So at the most, you'd probably get a delay of one second because you would start asking the other ones instead. Software fortunately is designed to handle that kind of thing.

In terms of an entire organization going away, if ICANN went belly-up tomorrow – I shouldn't pick on ICANN – but if ICANN

went belly-up tomorrow and they just stopped answering questions and their entire infrastructure went away, again, you probably wouldn't care because you'd start asking the other 12. That's the benefit of having so many, of having 13.

JIM MARTIN: Liman, did you want to add to this specific question?

LARS-JOHAN LIMAN: Yes. Lars Liman for Netnod again. I think that there is another myth, a misconception, that we need to kill here, and that is these Anycast instances are mirrors of something. They're not. They're self-sustained machines. They can operate. As long as they have power, they will respond to queries. As long as they get updated within a couple of days, they're just fine. So they're self-sustained machines that can operate. They don't have to talk to the [inaudible] every second.

UNIDENTIFIED MALE: I have a question. When we make a new entry –

JIM MARTIN: I'm sorry. Before you ask a question, we're queuing, so I'll put you in queue. Warren?

WARREN KUMARI: Yeah. Warren, F. Yeah, so as Lars was saying, they're all separate entities. Individual machines die every now and then. That just happens. The system as a whole doesn't notice it. First BGP Anycast takes of the routing side, so it flows to other ones at the same letter. Systems will automatically fail over to a different letter as well. There are multiple levels of redundancies, so breaking the whole system doesn't really happen.

JIM MARTIN: Elise?

ELISE GERICH: It's on the same point, but I think – and I don't know if you said it in the presentation today, this afternoon, but it was certainly in an earlier one – there's something like 500 instances. So it's pretty unlikely that all 500 of them across the globe would suffer catastrophic failure all at the same time. It would be odd.

JIM MARTIN: Indeed. Just to go back to your first question, you asked if there were different versions of the zone file for each of the different letters. Everybody serves exactly the same zone. Or was that not your question? What was the first question? I thought it was…

UNIDENTIFIED MALE: Data of a root server. Is it distributed in all other servers as well, or it contains only the data of itself? The A-Root server will contain data of the A-Root server only, or B, C, D, E, and F – every data?

JIM MARTIN: Every root server serves the same root zone file. There's no A-Root data and B-Root data and C-Root. It's all identical. The idea is that any one of the letters could serve for the entire global community. The reason there are multiple letters is not to split up the queries being asked, but to ensure that there is resiliency.

[STEVEN CONTE]: Wes, and then we can go to a new question.

WES HARDAKER: Okay. One important thing that you ought to do – I put this slide back up, which is RSSAC 003. If you care about TTLs and you care about what happens when things fall off and how long data is good for, RSSAC 003 is actually one of the published documents that goes into that. You really ought to read it because it talks about how long TTLs should be and what happens when data doesn't change frequently enough.

It's not a complete answer to what you want, but it's a good start in the direction of what you're thinking.

JIM MARTIN: Warren has an asterisk.

WARREN KUMARI: Yeah. Warren Kumari, F-Root. Following up from what Jim said, because of DNSSEC, all of the root servers have to serve exactly the same set of data. Even if they wanted to change stuff or add stuff or remove stuff, the root server operators cannot do that. DNSSEC locks the entire system, so that's the data that has to get served. The data that gets served, ICANN/IANA determine exactly what that is. All the root servers do is distribute that.

JIM MARTIN: Okay. This gentleman, and then the gentleman over there.

UNIDENTIFIED MALE: When IANA receives a new request for the new entry in the zone file, at the same time, simultaneously, it updates all the root servers – I'm not going to use the word "mirror" now – at the same time, or it takes a different time to update different root servers?

UNIDENTIFIED FEMALE: Wes [inaudible]

JIM MARTIN: Wes, and then Liman.

WES HARDAKER: Going back to the other architecture, if you remember, once it has been approved by the NTIA and ICANN, Verisign gets it because they're the ones that maintain the root zone in terms of signing it with the zone signing key and sending it out.

Functionally what happens is their servers send all the rest of the letters a notification saying, "I've changed." All the rest of the letters –

UNIDENTIFIED MALE: So there is no concept of primary or secondary, right?

WES HARDAKER: Right. All the rest of the letters will immediately go, "Oh. I need to go grab a new copy." That's what should happen. So from these –

UNIDENTIFIED MALE: So all the 500 root servers across the globe update it at the same time?

WES HARDAKER: I can only speak for my own in this case because that's the only data I know, to be honest. So I've measured it –

UNIDENTIFIED MALE: Okay. We will talk about yours. So you get all at the same time where you make the entry. How do you do that? I just want to know.

WES HARDAKER: I actually measure it. Because, again, of RSSAC 002, I've measured ours. When it comes in, we take less than two seconds to actually update all of our instances at B.

UNIDENTIFIED MALE: Oh, okay. At the same time.

WES HARDAKER: Two seconds. That's it.

UNIDENTIFIED MALE: Two seconds.

JIM MARTIN: Liman, did you want to add to this? Okay.

LARS-JOHAN LIMAN: Two short points. Lars Liman from Netnod again. The first point is, we serve the same zone file and it is public. You can download it from an FTP site. The things that we serve right now is there without delay. So just download it. If you want the URL, come and see me afterwards.

Second thing is, I pulled up the distribution numbers for I-Root. I speak for I-Root, not for B, not for F. For I-Root, the latest zone transfer from the point where Verisign said, "I have the new zone. Come and get it," until it hit the last of our 50 servers: 11 seconds. Okay. The average was 4.5. We have a few that are very far away that take a little extra to transfer zones. So that's out order of magnitude.

UNIDENTIFIED MALE: Still I'm not able to understand what exactly you do to do that.

LARS-JOHAN LIMAN: Come and talk to me afterwards. I'll happily discuss with you.

JIM MARTIN: Maybe we can takes this offline, then, and continue the details of that [inaudible] because we've had a very patient gentleman in the corner.

GREGORY MOUNIER: Thank you very much. Gregory Mounier from the European Cybercrime Centre. Are you aware of any malware that might have been written specifically to attack root end servers?

JIM MARTIN: There's no doubt that over time there have been attacks on the root server. We tend to focus more on the operational action of remediation, rather than trying to root out the source.

We do have collaboration in our community with various law enforcement and that sort of thing, and we will pass information associated with that to law enforcement. However, our focus is much more on an operational role of ensuring that the zone is available, and less about trying to track down the root sources.

UNIDENTIFIED MALE: John?

JOHN CRAIN: Yeah. There have been possibly malware – because we've not gone and actually got the source because that's not our expertise – that have hit the root servers. But my understanding, having a look at what they're doing – I know a little bit about malware, not a lot – is that typically it's when somebody makes

a typo. They're not actually aiming at the root servers, but if you make a typo at the end – you forget to put a dot on the end of something or you meant to say com but you say something that doesn't exist – it ends up coming to us.

GREGORY MOUNIER: But would it make sense according to you to target the root server specifically? From a criminal mind perspective, would it make sense or not if you really wanted to make a big mess?

JIM MARTIN: Warren?

WARREN KUMARI: There have been a bunch of discussions on that and there have been times when people have tried to attack the root server system. It manages to adapt.

But one of the issues with that is cybercriminals want to be able to use the Internet to actually commit crime. If you do something where you slash and burn and everything goes away, suddenly you're not able to use the Internet for all of your other crime things. So it doesn't really work out in your favor. So there is definitely something to that.

John looks like he's going to explode if I don't stop talking.

JOHN CRAIN: Boom! I'm [inaudible] we're going to stop answering. No, that wasn't it.

JIM MARTIN: All the remote users appreciate the explosion that just blew their eardrums out.

JOHN CRAIN: Good. Excellent. There is a published case of threats against a root server, which is probably what you're alluding to, and that was an anonymous operation that said they were going to attack the root servers. We actually all worked together with international law enforcement to prepare for that because that's what you do.

But if you follow the discussions amongst those groups, they started with, "We should go get this," and they went, "Oh. Maybe not. Maybe cutting off our own head to spite somebody is not the right idea." So even the ones that actually threatened to take us down never even attempted it because they realized it's not in their benefit.

JIM MARTIN: Warren's bouncing, so go ahead.

WARREN KUMARI: Yes. Some root operators actually kind of like that because it was easier to go along to their managers and get justification to add more nodes and more servers.

JIM MARTIN: Any other questions from anyone? We got one in the back. Hang on one second.

UNIDENTIFIED MALE: Thank you very much. With the introduction of new gTLDs and ccTLDs by ICANN, what measures have been taken by the IANA towards the availability, the liability, and stability of the root server operations?

My second question is that the policymakers of ICANN are facing some issues in the new TLDs regarding the allocation of new TLDs, some TLDs like .wine, etc. Is there any technical viable solution? You are experts in the root zones. Is there something, architectural changes, that can be made for the availability of a single gTLD to different registries? Thank you.

JIM MARTIN: Do you want to take the first part, Elise?

ELISE GERICH: I'll repeat the question to make sure I have it correct. I think what you asked was, was there anything that the IANA department did in order to prepare for the New gTLD Program and the additional entry of TLDs into the root zone?

It wasn't necessarily the IANA department that had to take steps to see if the root zone was going to be capable, but ICANN did reach out to RSSAC, the Advisory Committee, which is comprised of representatives of the root zone operators, and said, "Well, if we add 500" – whatever the number is; I think we only thought we might get 200 new TLDs – "can the root zone support that many more entries?"

RSSAC came back and said yes, and I think they said that they didn't anticipate any problems and they thought, even if it wasn't 200, if it was 1000, we'd be okay. They indicated what kind of headroom they have in their systems based on their own measurements. So it wasn't necessarily the IANA department, but it was ICANN, depending on the Advisory Committee, that said, "Yes. The root zone should be okay if you add more TLDs."

WES HARDAKER: One of the reasons that RSSAC 002 was written, which is to measure all of the root operator machines, was specifically so that we could track that over time, so that we can look at how fast all of the various letters are responding, how fast each

system is responding. If it suddenly looks like, as we are adding TLDs – you notice that they didn't add all 1000 at once. It was a long rollout of a couple a week kind of thing.

ELISE GERICH: Actually, a year and a half.

WES HARDAKER: Year and a half. Excuse me. One of those reasons is to make sure that we didn't take out the whole root zone at once. But the statistics that we're measuring should give us a better clue over time of how, if we added 10,000, what would happen, hopefully.

ELISE GERICH: Elise, just one second. Elise and then Liman.

UNIDENTIFIED MALE: Just to add something about the introduction of new TLDs, I think most of the concerns and the problems were from the client side, the resolver, so, I think, more than the root operators.

JIM MARTIN: Elise?

ELISE GERICH: I said Benno could go first, and then I [inaudible]

JIM MARTIN: Okay. And then Liman.

[BENNO OVEREINDER]: Actually, as part of your question you asked about the impact of new gTLDs on the root and the root servers, ICANN commissioned a project to some impact analysis. Tomorrow morning at 8:00, very early, there's a presentation of the preliminary results and approach of this study. So if you're interested, tomorrow morning, 8:00, in [inaudible], if I pronounced it correctly. It's nearby.

In this session there will be a presentation of approach, first results, preliminary results, on the analysis of the impact of the introduction of new gTLDs in the past two years.

JIM MARTIN: Liman?

LARS-JOHAN LIMAN: Lars Liman from Netnod again. As reading material for that presentation, I can recommend the report that was actually done in preparation for the new gTLDs. It's called "Scaling the Root." It's a 50-page report about what we thought would

happen and how we [projected] everything. So it was a substantial amount of work put into this beforehand.

ELISE GERICH: Just to conclude on my first response, I think if you look back at the data on the root servers, you'll find out that, when we introduced DNSSEC, that grew the size of the root zone far more than the entry of any individual TLDs or even any large number of TLDs. So it's the change in technology that seems to make a difference, and I think, when each TLD has IPv6/Ipv4, that also increases the size the root zone. So it's more the technology, not necessarily the number of TLDs.

JIM MARTIN: Any other questions from anybody? Go ahead.

UNIDENTIFIED MALE: Where can I find the document which is the root server operators are accountable to IANA? In what way?

ELISE GERICH: I had this question in the Newcomers' Group, too. The root server operators are not accountable to IANA. They're not accountable to ICANN. They're accountable to the community. All the root server operators were established – actually, yeah,

all of them – before ICANN ever existed. They've made a commitment. Several of them have letters of commitment stating that they support the global Internet community and that they do good by just doing it. They don't receive any remuneration from anyone – well, maybe if you put instances out. I don't know if you do or not. Individuals might.

Basically, it's a service. It's a public service, and the organizations that do that make a public service. As you saw earlier, what they've said as a collective group, they believe that it's the IANA functions operator that receives requests to change information in the root zone, and they'll only trust that if it's gone through the process of meeting the criteria. Then that gets sent to whoever and distributes it. Does that answer your question?

UNIDENTIFIED MALE: So there's no other option we can force them, if they fail to perform a certain set of criteria or performance?

ELISE GERICH: There's no SLA. No. And I don't think there's been any need for it. All of the 12 organizations have served the community for many, many years for the good of the community and the Internet.

UNIDENTIFIED MALE: So there should be some stress test now.

ELISE GERICH: Well, I don't know. We just said that DDOS attacks, a B could drop out and you wouldn't notice. Several instances could drop out and you wouldn't notice. It's actually the beauty of the design, the diversity of the operations of each operator, and the hierarchy and the distribution. Maybe people who are closer to it can speak to it than I, but we don't own and we don't regulate.

JIM MARTIN: Liman has volunteered to speak to that.

LARS-JOHAN LIMAN: Since I happen to be author of the document, there is a pair of documents that go together that define the expectations that the community should have on the root server operators. These are, on the quality side, so to speak, what type of protocols should we provide and so on, and that document was created inside the IETF [side]. So that's an RFC document.

The other one is more amounts to volumes and stuff, and that's an RSSAC document. So RSSAC 001 and RFC 7220, I believe. At the top. It's at the top. These two documents define what you

should expect from us. If we don't, there will be a lot of shame, I can tell you. Things will happen and they will get fixed.

UNIDENTIFIED MALE: No, but there's no subject agreement, right? It's a kind of protocol standard you have to follow.

LARS-JOHAN LIMAN: It is. You're right. There is no formal agreement regarding that anywhere, but I-Root is one of those who have made a statement to the IANA about this relationship, and so have several others.

UNIDENTIFIED MALE: What about ICANN?

JIM MARTIN: Just to continue this part, we have Warren, and then Elise wants to contribute on the dialogue there. Warren?

WARREN KUMARI: This is probably not going to make me very popular with other RSSAC root op people, so maybe I'll escape before the meeting ends. The root really isn't that interesting and that exciting. It

only handles a small percentage of the queries. .com, .net, and .org handle way, way, way more actual queries.

The root, there's a bunch of people who run it. It works well, or at least I think it works well. As a whole system, it's never gone down. There are much other more important sets of things in the Internet that we should be worrying about and keeping running.

Even if the root servers got really slow, probably nobody would notice. You very seldom need to find a TLD. There are only of them. You need to find names within TLDs much, much, much, much, much more often. So this really isn't the exciting part. There are other much more exciting parts.

I'm really glad that people came along to listen to this, but this isn't where the fun stuff is.

JIM MARTIN: Elise? No? Okay.

UNIDENTIFIED MALE: [inaudible]

JIM MARTIN: Yeah. Please, John.

JOHN CRAIN: Yeah. To the ICANN question, we've also made a statement. This part of ICANN that operates this has a published statement on our website that says we will meet [inaudible] –

UNIDENTIFIED MALE: For the L-Root server?

JOHN CRAIN: Yes. We will meet all those criteria and much more.

JIM MARTIN: Okay.

ELISE GERICH: I guess what it comes down to – and a lot of people think that it's really important, and Warren kind of said that, "No, it's not all that great big a deal." But there's a reputation on line here. It's kind of just like your own family. You don't want to let your family down. You don't want to do something bad that will embarrass them. So the root server operators are sort of like a family, and each and every one of them really want to do the best they can for the Internet community. They don't want to embarrass their organization. They don't want to be embarrassed in front of the Internet.

I think one of us – I don't know who said it; Liman, maybe – if I-Root started to behave badly, Liman would be in a whole lot of trouble because it would reflect on his organization, who prides themselves on having provided this service for a very long time and doing a really good job.

So there's really a lot of reputation on the line to continue to do that, and I think that's far more powerful than any sort of legal agreement, where you're just doing it because the contract says so. This is because you're doing it because you're really invested.

I think the agreements and SLAs – this is part of being part of the community. I think that's important.

UNIDENTIFIED MALE: But, as per my understanding, there's a need of – okay, I agree. They have a reputation and they made it [inaudible] several years. That I agree with you. But the times are changing. The new economics are coming up. The Internet users are coming from the different language backgrounds. I don't think so. In the future there will be a need of some kind of agreement or some kind of performance metrics, I guess. [inaudible]

JIM MARTIN: John, and then I want to bring new voice in. Then we're going to wrap it up.

JOHN CRAIN: Yeah. Things will change. Always. Through systems like RSSAC and ICANN and public discussions, like this open, transparent system we have, things will change. Maybe ten years from now, we do need something completely different. Or maybe six months from now. We don't know. So we will have those discussions in the future, beyond doubt.

UNIDENTIFIED MALE: We many not even have [inaudible] root servers in the future.

JOHN CRAIN: We may not even have root servers in the future. Maybe we won't need them, and that's completely fine. Things will get [inaudible] –

UNIDENTIFIED MALE: We are really going to have it because IPv6 is [inaudible]

JIM MARTIN: Okay. Thank you. Alejandro, last word. You get the last word.

ALEJANDRO ACOSTA: Oh, really. How lucky. Well, hello. Alejandro Acosta from LACNIC. I'm going to take advantage of that things change. Is there any chance that, maybe in the future, there will be more than 13 root servers?

ELISE GERICH: Operators.

ALEJANDRO ACOSTA: Okay, operators. That's right. More than 12 operators and 13 original servers?

JOHN CRAIN: I'm not aware of any currently published studies that cover all the technical details of it. Somebody else might be. There's always a chance. There could also be less,. You have to look at this and see what the technical needs are for expanding the number of the letters, the number of servers, or even decreasing them.

But remember, it's got to be about the technical service, the service that is provided, and making sure that that is better.

Most of the time when we see, "I need a letter," it's nearly always for political reasons, and mostly it's, "You need to add one letter because I need one." But if we need to add a letter, it's going to

be for a service need, and it probably won't be one. It may be a complete redesign. We don't know.

If I look at the over-provisioning now for the service, we keep saying we could lose some. So maybe we need less letters. We don't know. We've not actually done that study, and that study needs to take place at some point.

JIM MARTIN: Jim from Q-Root, I think.

JIM: Q-Root? Yes. I'm Q-Root these days. No. So the thing to keep in mind – I'm pretty sure you know this, but I want to make sure that everybody who's listening knows this – there's a world of difference between adding another letter and adding another instance.

Adding more instances is entirely possible, and it happens daily that there are new instances coming up in the world. Those are instances are happening all over the world in all the various different regions. So adding root server infrastructure to serve the new regions, growing regions – anything along those lines – is entirely possible today.

The only thing that is ill-defined and would really spend much more time in the political world than in the technical world, is the number of letters. So, please, if you're interested in hosting or feel a need for an instance, there are many that can help you out with that.

JIM MARTIN: Wes?

WES HARDAKER: One last data point, which is, if you really want to look into studying these types of things, there are a number of organizations that are working together to try to see what happens when a number of large changing metrics with the root system happens. And there's a project called the Yeti Project, which was started by Paul Vixie and a number of other people, that is doing a number of interesting metrics.

They're adding as many servers as they can. They're rolling the DNS keys as fast as they can. They're doing a whole bunch of stuff to study this exact kind of problem. If you're interested, you ought to go talk to them because they are trying to look at this exact problem to see what they can determine the end cap might be.

UNIDENTIFIED SPEAKER: [inaudible]

WES HARDAKER: Yeti. Y-E-T-I.

[STEVE CONTE]: With that, I want to specifically thank Jim and Wes for the presentation. Thank you very much, as well as all the other root server operators that are hiding out here, as well as Elise, and especially Steve Sheng for allowing me to harangue him over and over again to get this presentation here and in Dublin. So thank you all.

Everyone else who's here, I've given you a piece of paper. It's a small survey. We are actively seeking input of these types of sessions; if this was a relevant session, what else would be interesting? If you take a moment and if you haven't already, please just fill it out. I read it. I bring it back to my boss and say, "Hey, we should do it this way."

With that, thank you guys very much for sitting in today. Enjoy your evening. Day One, so you still have more days of ICANN to go. So enjoy them. Really. Please.

[END OF TRANSCRIPTION]