SAN JUAN – How It Works: Root Server Operations
Saturday, March 10, 2018 – 17:00 to 18:30 AST
ICANN61 | San Juan, Puerto Rico

UNKNOWN SPEAKER: March 10th, 2018, How It Works: Root Server Operations, Room 209 BC, 5:00 p.m.

CATHY PETERSEN: Good afternoon, everyone. Welcome to How It Works. This session we'll be talking about Root Server system, a tutorial on Root Server operations, and our first presenter is Andrew McConaughey. Andrew?

ANDREW MCCONAUGHEY: Thanks a lot, Cathy. So, my name is Andrew McConaughey, I work for -- is everyone hearing? Does the mic sound okay? Okay, so my name is Andrew McConaughey, I work for ICANN policy department supporting the RSSAC, and I'm here today to tell you about the Root Service System. I'm going to be doing the first three sections, and then my colleague, Carlos Reyes who hopefully will enter the room at some point, will be doing the final section.

So, I'll be talking about the overview of the DNS, the Root Service System today and its features, I'll walk through how Anycast works, and then Carlos will take you through the RSSAC as an organization and some of their recent activities.

So, here we go with the overview of the Domain Name System and Root Servers. First off, a bit of a recap on how identifiers work on the internet; we've got IP addresses, and IP's are kind of the fundamental identifier on the internet, they go in packets, it's how packets get places. All hosts connected to the internet do need to have IP addresses, whether that's through NAT or IPv6 or IPV4, we won't go into that, but we see in the bottom right there we've also got two different kinds of IP addresses; IPV4 and IPv6 addresses.

So why do we need DNS if we've got IP addresses? Well, the original problem was, IP addresses are kind of hard to remember; people tend to like words, they like names, IP's are just numbers, and also that IP addresses change a lot. There's a bit more of a modern problem that DNS is also solving, and that IP addresses may also be shared. I talked about, I mentioned NAT previously on the other slide, also that we could have multiple IP addresses for a single service. So, you can kind of think of this green box as with a one to many and a many to one

problem. So, DNS kind of has, it still has the original problem, but we've also got the more modern problems.

So this is probably, I mean, this section on DNS is probably a big review for a lot of people in the room who were probably in the earlier sessions as well, but just quickly, this slide is talking about how DNS is a hierarchical system, we've got a rooted top, and then we've got some TLD's immediately under that, EDU, MIL, UK, and then second level and third level. And then over on the right we see this kind of named IP address mapping that I'm sure you're at least somewhat familiar with. And then there's a bit of a discussion as well about other types of records you might find in DNS for mail servers, or for reverse lookups or for IPv6 records.

So, there's a lot on this slide, and this slide is at first pretty confusing, but I'm going to kind of walk through it. So, this is how your average user using a computer would interact with the DNS and the process. I have to stay in the camera. And the process by which all the various packets and queries would go forth so that that user can resolve a name.

And we see on the right hand side we have a computer and a user, and the user, this is also called a "stub resolver," which is a term you might encounter, and the user wants to find a web server for www.example.com because there's some great cat

pictures there, so our user is going to go to a recursive name server first, or that users computer is going to hit a recursive name server, and that computer is going to ask a pretty basic questions which is, "What's the IP address for www.example.com?"

And that recursive name server, assuming that recursive name server doesn't already have anything cached and this is the first time it's been asked where .com is or where example.com is and where www.example.com is, it needs to go to the root; it needs to figure out where .com is. So, it does that, and it asks a root name server, "Hey, where's .com?"

It gets a response, and along with that response comes a signature, because this slide is showing not only the queries and the responses, but also kind of how some basic background of how DNSSEC works. So, the recursive name server is going to ask a root name server, "Where is .com?" And he's going to get a response along with a signature, so he can verify that the information is valid. And then he's going to go with that information he received on where .com is, go ask the .com server, "Hey, where is example.com?" Get that response back with a signature, validate it, then he's going to go, "Okay, now I know where example.com is."

So, he's going to go to example.com and he's going to say, "Where's www?" And he's going to do the same thing; he's going to get the response back, he's going to get a signature, he's going to validate it, and then finally when that recursive server knows where www.example.com is, he's going to send that response back to the stub resolver; our computer with our user who wants cat pics, and then that user can finally look at her pictures of cats.

A bit more about root servers; root servers only know who needs to be asked next, as we saw in the last slide. They only have the address of the TLD servers, so the .com and the .org. And what we didn't see in the last slide is that if that recursive name server had a whole bunch of caching, and it already cached that response from the root server, it wouldn't have to go ask again. So, there's a lot of caching that happens at recursive servers, and so when they get these responses from the authoritative side, from the example.coms or from the root, they cache it, and so when the next query comes in they can serve it without having to ask again.

And here we're talking about some modern refinements to DNS. So, I mentioned DNSSEC earlier about the validation, and that the validation is something that happens in the recursive servers. Well, on the authoritative side, they serve the

signatures, so they respond with the signatures and then the recursive servers validate the signatures and the responses.

There's also some very recent privacy enhancements to DNS like DNS over TLS; this is very recent. As well as Q and A minimization, which is also pretty resent. And so, this slide says there is standards being created to reduce this; that's still true, it's still very much a work in progress.

And then we've got on the bottom we've got Anycast, and I'll be talking about Anycast later; that's the third section, and there's a whole section where we'll discuss Anycast. But briefly; multiple servers share a single IP address. It helps to improve latency and resilience and primarily it helps against, protects against DDoS attacks; Distributed Denial of Service attacks.

So, the root zone versus the root servers; the root zone is the data, you can think of it as the data that's being served by the root servers. So, the root zone is basically the starting point, it's the list of the name servers for the TLD's. And it's managed by ICANN per community policy. It's compiled and created or distributed by the root zone maintainer to the root server operators, and you can think of it as the database content of the root servers. It's what the root servers are serving.

The root servers on the other hand, they respond with data from the root zone. And currently there are 13 identities, and over 900 instances; 900 Anycast instances at different physical locations worldwide. That number is kind of slowly going up all the time, so I say, "Over 900," I think there might be over 950 right now but don't quote me on that. It's over 900, last time I gave this presentation it was over 800, so now we had to change it to over 900. And they serve a purely technical role; they serve the root zone.

Root Server Operators; operate root servers, kind of obvious but it's worth saying. There are 12 different professional engineering groups, and they're really just focused on reliability and stability of the service, with the service being serving the root zone. Accessibility for all internet users, lots of technical cooperation, and a professional attitude. And they're a diverse group of organizations and operations and you'll see the list later, but they're diverse in different ways. They're diverse technically as well as different types of organizations, as well as geographically dispersed around the world.

Some of the things that operators don't do; they're not involved in policy making. They're publishers, not authors or editors of the information. They publish the information, they serve the root zone. And they are of course involved in some careful

operation evolution to service. They work at the ITF to work on standards and they expand their service as the internet expands. They evaluate and deploy suggested technical modifications to standards and whatnot, and they work to ensure stability, robustness, and reachability.

So that was kind of a basic overview of the DNS, and the Root Server System. Now we're going to get a bit more into the Root Server System of today, and some of its features.

A bit of history; the root server system has been growing since the early 80's, since the invention of DNS and there you can see the timeline, various dates, and how many addresses there were at those dates. These changes have been due to technical demands on the internet, not on the DNS. Nowadays, Anycast is really used as kind of the super solution for scaling issues. We can also see that IPv6 addresses were added in 2008, and currently all operators have IPV4 and IPv6 addresses, and once again serve from over 900 instances.

Here are some of the foundational principles of the root server system. Primarily to provide a stable, reliable, and resilient platform for the DNS, and it operates for the common good of all the internet. The IANA is the source of the DNS root data. Architectural changes have been made based on the results of technical evaluation and demonstrated technical need; notice

the word “technical” is there twice. And technical operation and expectations of the DNS is defined by the ITF and root operators do participate often at the ITF. And if you would like some more history on the root server system, please do check out RSSAC023 which is the history of the root server system.

So, here’s the list of the root server operators as well as their manager, and the IP address payers. You can see every root server operator, or, every identity actually has an IPV4 and IPv6 address.

And this is a screen capture from rootservers.org so every time I give this presentation I update this to capture the latest number of root servers around the world, right now there’s over 900. These numbers, this is a graphic, so I don’t want people to think that means there are exactly 15 in Australia or something; it’s a pretty graphic on the website, it gives you an idea of how geographically dispersed the root server instances are, they’re all over the place, they’re all over the world. And if you’re interested into drilling down into exactly where they are with more information, I recommend going to rootservers.org and then you can look at the individual locations for each operator.

This is a slide showing how the root zone is managed, so the flow chart for how changes make their way into the root zone and eventually end up being served by the root server operators.

We see on the left-hand side we have a change request comes in from a TLD operator, so let's say a TLD operator wants to change their NS records in the root zone, they want to host their DNS somewhere else, so they would send in that change request, and then IANA would accept that and then that's the provisioning side.

And every once in a while, I think, I don't know exactly how frequently, but let's just say every once in a while, IANA sends the root zone to the root zone maintainer who prepares it, and then distributes it to the root server operators. And then we can see on the far right the DNS root server operators are serving the root zone data.

Do you have a question? Can we keep questions until the end, is that okay? Thank you.

So, these are some of the various features of the root server operators. Again, focusing on diversity, there's a diversity of organization, there's a diversity of operational history, there's a diversity of different types of hardware and software platforms that they use, and there's a diversity of funding. But they do have certain shared best practices, so various types of physical system security, so who gets into the data center; that kind of thing. Making sure to overprovision capacity in the case of

attacks, or in the case of spikes. And once again, a professional and trusted staff.

And so throughout that diversity, there's still a lot of cooperation. Like I mentioned earlier, industry meetings like ICANN, ITF, the various NOGs, RIPE DNSO work, a lot of collaboration between the operators, as well as a number of -- there's a slide later on transparency, which Carlos should go through.

And there's also lots of coordination, lots of permanent infrastructure that is arranged between the roots of our operators, so they can communicate with one another, mailing lists, credentials, phone bridges, that kind of things, to either deal with regular activity or some kind of activity that might happen during an attack so that they can communicate with one another and respond. And they also coordinate to establish internet bodies.

As the internet evolves, so new requirements are placed on the DNS system, so the root service system and DNS has been around since the early 80's and obviously there's been a lot of changes since then, so the introduction of IDN's, DNSSEC, IPv6, and root system analyzers have to analyze and adopt new users and protocol extensions to the service. And this has been mainly

to increase the robustness and the responsiveness and resilience of the root service system.

So, here's a slide with some of the myths that people sometimes have about the root server system, or about root server operators, and so I'll just go through each myth and reality, and maybe during question time if you have more questions about these myths, certainly bring it up.

The first myth is; root servers control where internet traffic goes. It's not really true, routers control where internet traffic goes, although it is true of course that if you're resolving a name you will get an IP address, and then the router will act upon that IP address, but in the end it's routers that decide where the packets go.

Another myth is that most DNS queries are handled by root server; this is not really true because most DNS queries are probably going to be handled by the cache in a recursive DNS server. At the very least for the TLD and then for the second level maybe the recursive has to go out and find out where the second level is but caching in recursive DNS servers really cuts down on the amount of traffic that needs to go to DNS root servers.

So, the third myth; administration of the root zone and service revisioning are the same thing; well, as we saw on the slide with

the flow chart of how the root zone gets updated, there's a pretty clear distinction between the administration root zone, and the service revision ending root zone, right? The people who prepare the root zone are not the same people who serve it and respond to DNS requests. The root server identities have special meaning; no, they're all the same. There's no special meaning there.

Another myth is that there's only 13 root servers, well, there's over 900. There's a lot of Anycast instances. The root server operators conduct operations independently; I mean, that's true to some point but there's also as we've been seeing repeatedly there's a lot of cooperation and coordination, so it's not entirely independent.

And then the final myth; the root server operators only receive the TLD portion of a query, and that is not currently true. And there's a little asterisks there, so right now as we were going through the whole -- you know, in a previous slide as we were talking about the recursive server resolves something for the user that wanted to look at those cat pictures, in that instance, the root servers received the entire name that was being queried; we call it a QNAME.

So, it would have been www.example.com in that instance. There are some new protocols under works; there's one called

QNAME minimization which attempts to change that, but that's still not really deployed yet; that's still being worked on.

Okay, so now I'm on to the third section of the presentation, this is going to be a bit more technical; it's going to be a bit more of a dive into Anycast.

So, Unicast versus Anycast. Unicast you can think of it as basically just one sender, one receiver, and it's always the same sender and it's always the same receiver based on IP address, so Unicast is pretty much how the internet worked for a very long time and it still does, I don't have any specific statistics but I imagine most traffic on the internet is still kind of going Unicast.

Anycast is kind of allowed to happen through a quirk in what's called the "Border Gateway Protocol," and it has a lot of really interesting functions. So the reason, well, some of the reasons why Anycast is so appealing when you're serving DNS data and for the root servers is because sources get the data faster and it can also act as a DDoS sync. So, it has this dual purpose where the sender of the traffic gets to a server that's closer to them, and also if in the case of an attack, traffic can be sink holed and it won't affect other people; it won't affect the whole system.

So, here's a very brief explanation of Unicast, and you can see on the left we have a green source and on the right, we have a blue

destination, and these little arrows represent packets; so traffic takes the shortest route to single destination, and the little circles represent routers, and the lines represent links.

Now here's Anycast, you see, we still have that one green source, but now we have three blue destinations. And we see that the traffic, instead of going all the way over to the right side of the screen to where that one destination was before, now that we have three different destinations, the traffic can go to the closest one. And this is handled through intermediate routing policies, I think I mentioned the Boarder Gateway Protocol, and the big upside of this is that the path is shortened, and the data is delivered more quickly.

Now here's the other benefit of Anycast; is that when a DDoS is underway; a Distributed Denial of Service Attack, the traffic which here it's in red, will go to an A destination, which in this instance we call a DDoS sync, and because we can call it "bad traffic," is going to this one destination while the good traffic in green is going to the other destination, they're not affected. The good traffic, the actual query that needs to be handled isn't affected by the DDoS traffic, so this is the other great benefit of Anycast.

When it comes to the root server system, if you're a network operator and you want to make sure you're always getting, like

let's say you're running some recursive DNS servers, you want to make sure that you're always getting good responses from the root server system.

These are some of the things you might want to think about, you always want three or four instances nearby, that Anycast instances, and you might want to consider hosting a root server instance yourself in which case you could reach out to a root server operator and ask them if you can host in Anycast mode.

Another thing you might want to look into is an RFC called 7706 for hosting your own local route, and there's more information if you read that RFC. You might also be interested in turning on DNSSEC validation, and if you're really interested in root server system, you're really interested in RSSAC, you can always apply to join to the RSSAC Caucus.

And now we're reaching the fourth and final section of this presentation, and my colleague Carlos Reyes is going to take over from now.

CARLOS REYES: Hi, thanks. Hi, my name is Carlos Reyes and I'm on the policy team at ICANN and I also support the RSSAC. For the final section today, we're going to give an overview of the root server system advisory committee.

So, the root server system advisory committee; this is one of the advisory committees of the ICANN community that has been chartered by the board and if you look at the bylaws, the role of the RSSAC is to advise the ICANN community and the board on matters relating to the operation, administration, security and integrity of the internet's root service system. So, this is a very narrow scope, and in the next two slides I'll explain a little bit more about this particular scope and how that informs the work of the RSSAC.

So, what does RSSAC do, and what does RSSAC not do? First, the RSSAC is a committee, obviously; it's in the definition. And it produces advice primarily for the board, but also to the community and other organizations and groups within ICANN, and this particular function is particularly scoped around the issue of, if you go back to the first slide here, the operation, administration, security and integrity of the internet's root server system. So, anything outside of that, the RSSAC as a committee does not involve itself in those matters. Now, the members of the RSSAC ay comment on these issues, but as a committee it's really focused on its work in advising on this particular issue.

We'll talk a little bit about the composition in a moment, but the root server operator organizations are represented within the

RSSAC, but RSSAC does not involve itself in operational matters as they relate to the work of the operations. That is a very fine line, and it's something that I think if you look at the history of RSSAC and how it participates in the ICANN community, there's been a confusion over the years about this, but the RSSAC is very much intending these days to clarify that and to bring some clarity into the relationship between the root server operators as their own work outside of RSSAC, and the root server operators participating the work of RSSAC. So, I'll explain a little bit there in a moment.

So where does the RSSAC fall into the overall structure of the ICANN multi-cycle per model? As I mentioned, the RSSAC is one of the four advisory committees. It has a non-voting liaison to the ICANN board and it participates in all of the supporting organization advisory committee activity of the ICANN community.

So, who is part of the RSSAC here? As I mentioned, there are pointed representatives of the root server operators, so there are 12 organizations, so each organization appoints a representative to the RSSAC, and they also appoint an alternate. And that just allows for continuity of service and representation. Sometimes because of the operational responsibilities of the organizations, they might not be able to make certain calls or

participate in certain activities, so this allows them to be part of conversations and be present at the table regardless of their commitments.

The RSSAC also has liaisons, and there's a slide here next that explains the liaisons, so I'll leave that for that slide.

And then finally, after 2013 the RSSAC created a caucus, this group is a broader group of volunteers and subject matter experts that expand the base of volunteers for RSSAC and provide input and actually work on the documents, the technical documents that the RSSAC produces. These members are confirmed by RSSAC, so I'll explain that momentarily.

So, the current co-chairs; Brad Verd, representing VeriSign and Tripti Sinha from the University of Maryland. Tripti's term ends later this December, and Brad was recently re-elected. So, both of them, the terms are staggered to allow for new leaders and also experienced leaders to continue the operation of the RSSAC itself.

The liaisons; earlier I mentioned the liaisons, the liaisons are there to inform RSSAC from different perspectives, but also for RSSAC to communicate with other groups. So, the first one is the IANA functions operator, currently that is public technical identifiers; PTI, which is an affiliate of ICANN. The Root Zone

Maintainer, that is VeriSign. The Internet Architecture Board; they appoint a liaison to RSSAC. The Security, Stability Advisory Committee; that's one of the four advisory committees at ICANN, they also appoint a liaison to RSSAC.

And then as I mentioned, the Bylaws, the RSSAC has a liaison to the ICANN board. The RSSAC also has a liaison Nominating Committee. The Nominating Committee at ICANN is the group that appoints board members and other external candidates to certain organizations at ICANN, so this function allows the RSSAC to bring in technical perspectives into the deliberations of the NomCom.

The Customer Standing Committee; this organization was set up after the IANA stewardship transition, and it monitors the performance of PTI in performing IANA functions, so the RSSAC has a liaison to that group since 2014. Excuse me, 2016. And then finally the Root Zone Evolution Review Committee; this is another group that was set up after the IANA stewardship transition, and the RSSAC has a liaison there as well.

Any questions about the liaisons?

The liaisons participate actively in RSSAC work, they don't vote as the RSSAC members, but they do inform the conversations and they allow for the RSSAC as I said earlier to engage with

these particular groups and bring in those perspectives into their work.

So, the Caucus; the Caucus was established after the first organizational review of the RSSAC, and what this group does, as I mentioned earlier, it expands the volunteer base for the RSSAC. If you look at the membership composition of the RSSAC, it's a fairly limited group, but the work is really never ending, so if you bring in a broader pool of experts, it allows them to tackle different issues. So currently, there are 88 technical experts, they all have public statements of interest online. This is part of the membership application process. And as they contribute to work, they are attributes in the documents that RSSAC produces, so this allows them to meaningfully contribute to the work of the RSSAC.

I've already alluded to what the purpose is, but some other points to consider; this actually brings transparency to the work of RSSAC. Because it's a broader pool of experts, they bring in other perspectives beyond the roots of our operator organizations. Some of them have experience with root server operations, but they also have broader DNS expertise.

Now, if you're interested in applying, the e-mail address there; rsac-membership@icann.org, we have a membership

committee, they review applications, and they're always welcoming new members.

Recent publications; so, I won't go into detail too much, because there is an RSSAC public session later this week where the RSSAC members will explain some of these documents in more detail, but just a high-level overview of some publications since the last ICANN meeting.

First; RSSAC 029, this is a report from the most recent RSSAC workshop. So, a few years ago the RSSAC started a process of undertaking certain workshops to advance work that they couldn't necessarily do at ICANN meetings. So those started, as I mentioned, two years ago, and that's just a report for the most recent workshop. The next workshop is coming up in May, and there will be a report published after that.

RSSAC 030; this is a statement that the RSSAC produced. Actually, it's also an outcome of a workshop that essentially defines what, excuse me, the statement provides a definition for what actually defines an "operator," and if you look at the statement, they elaborate on the three sources, DNS root sources.

RSSAC 031; so, this was a request from the GNSO; the Generic Name Supporting Organization policy, development, process

working group. So, if you're familiar with supporting organizations and the advisory committees at ICANN, the supporting organizations are the groups at ICANN that develop policies, and because the RSSAC does not develop policy, if an organization is in the process of developing a policy, they will consult with the advisory committees, and that's exactly what happened in this instance.

The RSSAC responded to this particular working group about questions regarding root scaling, so this particular working group is looking at the experiences learned from the new GTLD round, and how any changes to the policies may be pursued for the next round.

Current work at the RSSAC; there are two work parties that are currently underway; harmonization of anonymization procedures for data collecting. This is led by Paul Hoffman who's a Caucus member. Again, this is a working group, excuse me, a work party that was proposed by the Caucus, the RSSAC reviewed the statement of work, and decided to go ahead and charter it, and that work is currently underway.

Packet sizes; that's also another proposal from the Caucus, currently led by George Michaelson. And for both work parties, the RSSAC also has a shepherd from the RSSAC that tracks the

work, so as the Caucus pursues its activities, there's a porting into the RSSAC on the work and how that is proceeding.

So, some points about transparency; a few years ago, as I mentioned, the RSSAC established a Caucus to broaden its base of volunteers. But it also publishes minutes from all of its monthly meetings if you go to the RSSAC website rsac.icann.org. You'll see a list of all the meetings and their minutes.

The workshop reports obviously I mentioned earlier, any workshop that occurs, the RSSAC will summarize its deliberations and publish those for the general awareness of the community. Obviously public meetings, that happens at every ICANN meeting. Tutorials like this one, the RSSAC is actively engaging with the community

And then finally; the operational procedures. Every supporting organization in an advisory committee at ICANN has operational procedures, these are publicly available so that everyone knows how the RSSAC operates; how does it make decisions, how does it conduct its votes, how does it establish a work party, how does it review work products etcetera. So, all of that is available online.

Looking at the other column here about the root server operator organizations; the RSO's publish the agendas from their

meetings. They will be meeting I think soon at ITF 101 in London, so those agendas are published. The Root Server Operator Organizations also publish statistics that were requested by the RSSAC and RSSAC 002, which is an advisory that provides measurements of the root server system.

The public webpage root.servers.org, that is online, there's a lot of information about the root service system as a whole, but also about each individual operator, and there are different tabs they can explore including RSSAC 002 statistics.

And then finally, about a year and a half ago, the RSSAC essentially provided a door for the root server operator organizations, so any technical questions can come to RSSAC, they will then forward it on to the operators, and that e-mail address is: ask.rsac@icann.org.

So, this is the last slide here; as I mentioned, both the RSSAC and the Caucus have websites, and if you're interested in applying for the Caucus, please send an e-mail to rsac-membership@icann.org. I'll probably respond, and I'll direct you to the statement of interest template, and then at that point we can begin the membership process for the Caucus.

I'll pause here to see if there are any questions? Yes?

CATHY PETERSEN: If you could please provide your name and affiliation?

ABDALMONEM GALILA: I am Abdalmonem Galilafrom Telecom Regulator of Egypt, ICANN coach. I have many questions; the first question, could you go back to slide number 17? No, no, no, go to Anycast, because I talk about Anycast and DDoS attacks. Again? It doesn't run? Yes, it's this one. How does Anycast service prevent root servers from DDoS attacks? One by one, an attacker can attack one by one for example, maybe attack one root server, one instance of the root server, and after he realizes that this instance is already dead, it will continue for other barters with the same ID's, it didn't change anything, is the first question.

After the deployment of IPv6 worldwide, the number of main root servers changed to be more than 13 or not. Second question, this is the third one; could you elaborate more about the root servers maintainers?

The fifth one, what is the idea behind that root servers receive the entire query rather than the part of the query? So last question; what is the relation between BTI and IANA? I think that the IANA transition is already completed since 2016, thank you.

CARLOS REYES: Sure, so what I'm going to do now is I'm actually going to invite members of the RSSAC to help with some of these answers. And why don't we start with Wes? Thanks.

WES HARDAKER: This is Wes Hardaker from the University of Southern California; ISI. Excellent questions, so thank you for asking them, they were well thought out as clearly you were thinking about the problem. I will try to remember them and if I forget one you can remind me, okay?

First off, excellent question about DDoS attacks, but the reality is that a DDoS attack is an attack that's designed to overwhelm some aspect of service. So once, let's say the red blob up there on the DDoS attacker attacks the service on the upper right, that attack doesn't permanently damage it, it doesn't permanently take it out; that attacker has to keep up with that attack and actually take out that entire node.

The reality is, the rest of the internet starts, data send their queries, all go somewhere else, so you don't notice. In order to take out every instance, you would have to take out all 900 instances that Andrew talked about earlier, and that's a rather massive problem. You can't take them out one and go onto the

next; you've got to do all of them at once to have an effective attack, and that's not easy. It has not been done to date, let me put it that way.

That was the first one, the second -- the deployment of IPv6. So right now, all 13 named entities have IPv6 addresses, and I think, is there any instances -- I don't know of any instances that aren't running dual V6 and V4 at each of the 900. I think they all are, but -- no, okay, so there's quite possibly a couple that don't have IPv6 at certain instances. But the number is still very, very high in terms of the number of IPv6 deployment. So, all 13 individual names do have V6 addresses associated with them now, and there's a couple of instances that don't have V6 deployed all the way to the end, does that make sense?

Next question? Yes.

UNKNOWN SPEAKER: [Inaudible]. So that's why I am asking. [Inaudible] the package would be much bigger.

WES HARDAKER: The original design decision many years ago was to try and keep packet size below 512 bytes. That design decision is no longer relevant with packets that are larger, and they are larger. In fact,

DNSSEC has made packets bigger than even can fit in a single frame when you query keys for especially organizations like .org which has larger keys than the root for example. So DNSSEC has sort of shown that the world is ready for larger packets, so no, that's no longer a design consideration. John? You're up.

JOHN CRAIN: There is one nuance to that; is that there are still systems out there that suffer from the 512-byte issue, but with the addition of version 6, we can still fit all of the V4 information and some of the V6 information, so you will get enough data even if you're over that 512 bytes. What was the next question?

UNKNOWN SPEAKER: [Inaudible].

JOHN CRAIN: Which details would you like to know?

UNKNOWN SPEAKER: [Inaudible].

JOHN CRAIN: Okay, so the maintainer is basically two organizations; it's PTI's which is the operator of the IANA function and there's a contract with VeriSign who basically generate the zone from databases. So, the PTI/IANA folks will send over the data changes, there's a whole process there obviously, and then the zone file will get generated, get signed, and then will be pushed out to distribution servers where then we, as the operators, will get that and then spread it out. And that's pretty well documented, and if you're interested in the whole system, you need to look at not just our side, but also what the RSSAC is doing as well, who are dealing with different parts of that.

The discussion about whether or not you get the whole query or just a part of the query; well that's just the way it's been designed up until now is that when you send a query, you send the entire thing, and there is work happening now within the internet engineering task force to talk about whether that's the right way to do things, or whether for privacy reasons we should only send the needed part of the query, for example, for the roots, you only really need to ask them which top level domain, because that's the information they know.

UNKNOWN SPEAKER: [Inaudible].

JOHN CRAIN: There is some advantage to having smaller data.

WES HARDAKER: Really, I think the reason, the original reason, way back when DNS was designed is that you actually don't know where the boundaries are between say, the root and com, and example.com, and it may be that one domain, one server can actually answer multiple depths of the tree for you.

And if you only ask one, so queue name minimization suffers from potentially asking more questions than you needed, because you may have to go back to the same server to ask the next question, when the reality is, they could have answered a whole bunch more for you. A lot of times multiple levels of a label, multiple labels are answered by a single server, so it's an expense that you help with privacy. Warren?

WARREN KUMARI: [Inaudible].

WES HARDAKER: So, [inaudible] minimization was defined in early 2016; it's a very recent specification, and resolving software now, or enabling

those options and pushing down in into software. So, if you download I think the latest and greatest of many things like Unbound and Bind; I don't know what their timeline is, but they're beginning to show the options that you can turn that on, or off, you have the choice as a resolver operator of making that decision.

ABDULKARIM OLOYEDE: Hello, my name is Abdulkarim Oloyede from Nigeria. I want to ask this question; I don't really understand what is the relationship between DNS and packet size, because you have been talking about packet size DNS and I've been trying to work it out in my head, now what is the relationship between the packet size and DNS? Thank you.

FRED BAKER: Fred Baker, Internet Systems Consortium. The deal there is that we have two trans protocols that the DNS can use. It normally uses UDP which is once and done; I send you a packet, you send me a packet. You can use, and we use this option occasionally, it can use TCP in which I can send something that is much larger, and I break it up into some numbers, I can send them to you, you respond. UDP has a size limitation if nothing else, because there's a path MTU that has a size limitation.

And so, we have to calculate how much we can fit in there, and if it fits, it fits, and if it doesn't then we use TCP. So, the response that I think you're looking for is, "Well gee, why don't we just use TCP more often or in certain cases or things like that?" And I think those questions are being asked, but right now the usual service is built on UDP, which is once and done.

UNKNOWN SPEAKER: [Inaudible].

FRED BAKER: So, the question is; what's the reliability of UDP, and that's really a question of the reliability of the path. If the path, let's say that it has a limitation, it will only take 1000-byte packets, and I'm sending him a 1500-byte packet, the reliability is 0; nothing is going to get through there. In the same case, if I'm sending 500-byte packets, it's probably pretty good. And so, it really depends on the circumstances of the path. In most paths in the internet, frankly the internet wouldn't work very well if the reliability wasn't very high. Okay so, reliability is generally close to 100%, but yeah, there are limitations and bad things happen.

LENDON TELESFORD: Hi, Lendon from Grenada. It was mentioned that there were over 900 instances of the root server, I was just curious, if I'm interested in hosting one of the instances what are the steps or requirements?

MATT WEINBERG: My name is Matt Weinberg, I'm from VeriSign. So, VeriSign actually has a solution whereby you yourself can help host a J-root instance, more information is available at a website labelled rirs.verisign.com. But basically, in summary, the way VeriSign does it and Mauricio here will talk about how ICANN does it, but the way we do it is; there's an application process, you fill out application process, you have to meet certain criteria and be able to provide bandwidth and do a couple other things, but if all those things are met then we can look into possibly basically shipping a server at no cost to you.

The only deal for you is that you have to sign some paperwork and be responsible for hosting the gear. But then, the goal is to have a J-root instance available at your location. Mauricio?

MAURICIO VERGARA ERECHE: Hi, Mauricio from ICANN. We also are operators, so one of the root servers, so we also would like to have the most instances that we could have. So, if you want to host one of the

ICANN managed root servers, you can visit dns.icann.org, so there is information about how is the process; it's very similar to what Matt was describing for VeriSign. And if you have any doubts, we are going to be here the whole week, so we can talk about it.

ABDULKARIM OLOYEDE: Sorry, I wanted to ask one last question; how many of these root servers do we have in Africa?

LARS-JOHAN LIMAN: This is very difficult to talk to everyone from the center of the room. Hello, my name is Lars Liman, I work for Netnod, and we also operate one of the root server clusters. Speaking to the number of instances in Africa, it is something where we at Netnod would like to improve the situation, and I would guess that the others would like to improve it as well.

We do have a bit of a challenge to find good spots to put the servers because as was mentioned by Matt before, there are a number of technical requirements that need to be fulfilled, and there needs to be a good environment, there needs to be people who can respond when we have problems with the servers, that we need people that we can contact and so on.

And at least we at Netnod, we don't have the right relationship for this yet, so please do come and talk to us. It's not a matter of not wanting to put servers in there; it's about creating the right relationships to make this happen. And I'm sure Allen is going to speak more to that, please?

ALAN BARRETT: Thanks, Liman. Hi, I'm Alan Barrett from AFRINIC. If you're in Africa and you want a root server, come talk to AFRINIC, we'll hook you up with one or two or three root server operators, and in some cases we're also able to sponsor the equipment that you need. What you'll need is a place to put it and bandwidth, thanks.

ABUBAKAR MUHAMMAD: I'm Abubakar Muhammad from Nigeria. So to say, you talked about where to put them. Where we are to put them is not really a challenge, I think the exchange points are places that you can place root servers, and we have them all over, so we've gone beyond that I think, thank you.

WARREN KUMARI: I'm Warren Kumari, I just had a quick look at the map, it looks like they are currently around 66 or so, so I mean, there are

some, but obviously it would be better if there were a bunch more, just as a quick data point.

CARLOS REYES: Anyone else? Any other questions?

BRAD VERD: Hello, Brad Verd, VeriSign. Just really quick; maybe not to give you a data point that's coming out in Caucus meeting, but most of these questions are reflected in a FAQ that has been recently published on the RSSAC webpage, so all the questions asked here, over many, many of these tutorials we've given, we've aggregated them, kind of created what the general questions are with answers, so if you have questions and don't want to come up to the mic and ask, obviously come and talk to us. There's also the FAQ that you can look at, and there's stuff there, and that's constantly changing as new questions come up and whatnot, so I just wanted to share that.

CARLOS REYES: One more call for questions? Alright, well thank you everyone for joining the Root Server Operator Representatives will be here for awhile if you have questions. I'm here, Andrew is here. Feel free to approach us and thank you very much for joining us.

CATHY PETERSEN: Thank you everyone.

[END OF TRANSCRIPTION]