
PANAMA – Community Dialogue: A Proposed Governance Model for the DNS Root Server System

Wednesday, June 27, 2018 – 13:30 to 15:00 EST

ICANN62 | Panama City, Panama

TRIPTI SINHA: Alright. Why don't we get started? Call to order, please. Come on in, folks. Welcome to RSSAC's community dialogue on a recently ratified document, RSSAC 037, whose title is A Proposed Governance Model for the DNS Root Server System. I invite our audience to sit at the table, if you find a seat, if you would so like. We're small enough that I think it would be nice if we introduce ourselves. So, I'm going to ask the gentleman on my extreme left to start.

MARK SVANCAREK: Mark Svancarek, Microsoft.

RYAN STEVENSON: Ryan Stephenson, RSSAC, DOD.

TERRY MANDERSON: Terry Manderson, RSSAC, ICANN.

KEN RENARD: Ken Renard, RSSAC, Omni Research Lab.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

WES HARDAKER: Wes Hardaker, RSSAC, USC, ISI.

FRED BAKER: Fred Baker, RSSAC, ISE.

MATT LARSON: Matt Larson, RSSAC, ICANN.

DUANE WESSELS: Duane Wessels, RSSAC, from Verisign.

BRAD VERD: Brad Verd, RSSAC co-chair.

DAVID CONRAD: David Conrad, ICANN CTO.

TRIPTI SINHA: Tripti Sinha, University of Maryland and RSSAC co-chair.

KAVEH RANJBAR: Kaveh Ranjbar, RSSAC, RIPE NCC.

LARS LIMAN: Lars-Johan Liman, Netnod, RSSAC.

JONNE SOININEN: Jonne Soininen, ICANN board.

LITO IBARRA: Lito Abarra, ICANN board.

EDMUN CHUNG: Edmun Chung, dot-asia.

MARIO ALEMAN: Mario Aleman, ICANN support staff for RSSAC.

RUSS MUNDY: Russ Mundy, SSAC liaison to the RSSAC.

SUZANNE WOOLF: Suzanne Woolf, USC ISI.

JIM: Jim [inaudible].

DANIEL MIGAULT: Daniel Migault, IAB, RSSAC liaison.

UNIDENTIFIED MALE: [inaudible], JPRS.

RICK WILHELM: Rick Wilhelm, Verisign.

UNIDENTIFIED FEMALE: [inaudible], USA DOD.

GUSTAVO PAIVA: Gustavo Paiva, NextGen 62.

CARLOS REYES: Carlos Reyes, ICANN staff.

TRIPTI SINHA: Alright. Thank you very much. This is a 90-minute session and we've broke it up into 30 minutes, three sections. The first section will give a background on the work we've done and the second section we'll actually walk through some scenarios that teases out and plays out the model that we will present and we leave the last 30 minutes for a Q&A.

So, with that said, a little bit of background on what this work is. I think, I will assume, make the assumption that everyone in the

audience is familiar with what the DNS root server system is. This is a proposed governance model for that particular system.

I would like to start by talking about the staggering growth of the Internet. When the Internet was first conceived way back when, no one expected the group to be as staggering and as huge as it has been and it continues to grow at quite a pace.

The root server system itself, the 12 operators who operate the service have their roots in organic growth. In other words, when we were assigned and selected to be operators, never in our dreams did we expect the infrastructure to be as large and the user base to be as large.

So, with that context, we decided it was time to address some rather important and pertinent questions as it relates to the service.

A few acronyms that you'll hear while I speak, and if you read the document, RSSAC 037, RSS stands for the Root Serve System. RSO is an operator, a server operator. SAP stands for Strategy, Architecture, and Policy function. It's an element within this model. PMMF is Performance Monitoring and Measurement Function. Once again, another entity within the model. RS is simply root server. SF is a secretariat function. DRF is Designation Removal Function and FF is Financial Function.

I'd like to set the context. This is a three-year effort of work that was done by the Root Server System Advisory Committee. We are simply advisory to the ICANN board, and thus far, it has only seen our eyes and it is to be treated as an initial draft model. It is simply a starting point.

Our initial impetus to workshop – and what we mean by that is the work was done through a series of workshops that RSSAC conducted and taking the time back – and forgive me, my eyes are bad. I think I can read it. I've done this slide deck enough times that I think I know what's being said in each component of that.

But, in 2014, the IANA stewardship transition was announced and there was a tremendous amount of energy and activity in this community. Global eyes were on this work and multi-stakeholderism took foot and was grounded deeply in this community. RSSAC recognized that we, too, needed to come to terms with answering some rather difficult questions that had been asked for many, many decades.

They are accountability. Who are the RSOs accountable to and how can they assure continuity of service? And who are the stakeholders and what are we accountable for?

So, it was imperative that we start to answer these questions, and in September 2015, RSSAC decided the best way to do that

was to [inaudible] ourselves in the form of workshops and begin to work on these series of questions, and look into the future to see what would our legacy be in turning over this system and service to future generations.

So, the workshops began in 2015 I think is what that says. Forgive me. I can't read what's on the left. Essentially, that's the timeline. We had a series of six workshops. They were held in the Washington Metro area at the University of Maryland and Verisign shared hosting these two workshops and we began to address the question.

The metaphor that I use, as we began to peel the onion to understand what exactly our problem statement was and put together our thoughts, which eventually came together as a model. And you will hear much more about that.

This essentially displays the global DNS root service. The triangle depicts DNS. We are in a DNS community. I think everyone in this room understands how that works. The bottom of the triangle is where the client ecosystem lives, and as you travel up the triangle, resolution occurs. And at the very top is where the root server system lives and that's where the top level is resolved.

There are 12 operators that operate the service and scattered across various entities. There's a smattering of different flavors.

There are universities. There are private industry. There's the US government and other entities. There are RIRs as well. Most of our growth is very organic in nature.

Root server systems. When we came up with a model, we realized it was important for us to maintain some root server system principles, and I must say, I cannot read that at all. Can someone? Thank you, Brad.

It was important that we maintain the eleven principles that we identified that would move into the future. Most of them, the majority of the principles, we embrace today, there were some that we did not and we believe that we must embrace them going forward. So, I will walk you through these principles, as they are very important.

Principle one says to remain a global network, the Internet requires a globally unique public name space. In other words, the DNS name space should remain unique, should always remain single and never fractured.

Moving on to principle two, there is simply one source for the data and that is the IANA source.

Moving on to principle three, the RSS should be stable, reliable, and resilient in order for DNS services to work.

Principle four, moving on, diversity of the root service operations is a strength of the system.

So, we firmly believe that the reason why the system has worked as well as it has worked for the past 30-plus years is because of the diversity of the system, and diversity is defined as we are all very different operators. We have architected this system differently internally. At a very fundamental level, however, the service that is delivered remains exactly the same, no matter who delivers it. So, the diversity is strength in many ways.

Moving onto architecture, we say that if there is some kind of technical evolution, we need to demonstrate the need to imbed that new architectural element into the RSS.

The steward of the protocol itself is the IETF. Principle 7 says RSOs must operate with integrity and ethos, demonstrating a commitment to the common good of the Internet. I can firmly speak for all the 12 operators. We have done exactly that. There is no money to be made in this service. We have simply done it as good global citizens to the greater community.

RSOs must be transparent in the past. The RSOs have been accused of not being transparent and we believe that going forward, there should be transparency and openness in how the service is operated. Of course, within certain boundaries in that we do not want to compromise the security of the system.

Principle nine says that we will collaborate and engage with the stakeholder community.

Principle ten says that we will remain autonomous and independent. In other words, that will not be influenced by any political forces that may be at play and that could be defined as the government of a nation as well.

Principle seven says that we remain neutral and impartial. In other words, we do just simply one thing, deliver a service. A technical service. We resolve to the top level of the domain name system and that is all we do. Just advance it for me. Sorry. I should have brought my laptop.

So, this essentially is the proposed model, the scoping of the model. So, let me describe this diagram from left to right. This is root zone provisioning distribution resolution. So, what happens on the left is where the blue circle talks about TLD operations, and the white rectangle describes the IANA function, and that function is the one that decides what [inaudible] changes that occur in the root zone. That root zone, in turn, is fed to the operators who in turn feed it to their individual constellations, their infrastructure.

What happens in extreme right is resolution. So, that's when the resolvers are talking to the root servers and queries are being resolved.

The scope of our work is what you see in light blue. That is all we are addressing in this particular model. It is important that while we came up with our new model we imbedded in it model design principles, and the principles are we needed to construct different functions that operated around the model. So, we would ensure that there was a separation of functions so that there was like activities and affinity groups clumped together. We wanted to ensure that there was an avoidance of conflict of interest, so there was no undue influence in any particular function. And all the functions needed to be transparent and auditable. At the same time, we want to ensure that we maintain the security of the system as well.

At some level, the 12 operators know who their stakeholders are. We are certainly beholden to our own individual operators, organizations, and we resolve for any end user of the Internet community. But, at some point, you do need to abstract this to a body that has some level of influence over what the RSS is doing and the operator is doing. So, we looked at this question quite hard and we ended up on three entities.

I'll start with the IETF/IAB. We believe they are the stewards of the protocol and we need to ensure that we do exactly what the protocol tells us to do. So, they indeed are stakeholders here.

We put ourselves in there, the RSOs and future RSOs. They as well are stakeholders and of course the ICANN community. We looked at various and sundry bodies across the globe and there is no other body that has as well represented to represent the Internet community across the globe. We think ICANN, the community and organization, does a very good job of representing just about every kind of entity in the world that needs the Internet. So, we feel that this was a good place to land.

This diagram essentially is a very high-level snapshot of the model, what we are calling the model. It's a very delicate interplay of three constructs that operate in tandem. So, what you see in the top in the pale pastel green is what we call governance. So, the highest level you have the three stakeholders that I just talked about, the ICANN community, IETF, IAB, and the RSOs.

Below them operate five different functions. There's a secretariat function and I'll get to the details of that in a minute. Then there's strategy, architecture, and policy function (SAP for short), a financial function (FF), the designation removal (DRF), and the performance monitoring measurement function.

What you see on the left-hand side tall triangle in light gray is just plain, simple DNS operations. This is just a day-to-day humming of the service by the 12 operators.

On the right-hand side in pale pastel orange or peach is what we call the onboarding and offboarding of operators. So, there are certain performance metrics that the operators must meet and their performance metrics that potential new operators should demonstrate that they can meet.

The way you identify a bonafide operator is in three sources: the root zone, the root [inaudible] file, and the root [inaudible] net zone. Activities occur in those three locations to remove and designate an operator. So, if you need to check and see who an operator is, that is where you need to look. Thank you.

So, governance is an interplay of these separate functions. A lot of care was taken into identifying these five different functions and all the model principles were applied to those. So, governance is in the middle, and you've got the five different functions interacting with each other. Next slide, please.

I will very quickly describe what the different functions are. The secretary function is something that happens today. It's a loosely organized meeting ... I shouldn't say loosely organized, but we're referring to it as a loosely organized secretariat function that occurs today where the root operators currently

meet three times a year, and we would like to add some form and function to that activity and we call that a secretariat function.

This function would simply coordinate and support operational meetings. They would hold common assets and they would be a point of contact to speak to the RSOs.

One of the big complaints we get today is we don't know how to reach the 12 operators and this would be a good place to begin in the future, once this model is implemented.

The next slide talks about the SAP function strategy, architecture and policy. There are three distinct streams here. Strategy essentially says we would coordinate with other stakeholders a strategic vision for the RSS, so it's essentially strategic in nature. It's looking over the horizon. It's looking to see what emerging technologies are around the corner and what needs to be incorporated into the RSS and what needs to be sunset as well.

The architecture says that this group of – a functional architecture element that says this is what the system should look like and should be designed to these parameters and operated in this matter. And it would define what we call externally verifiable metrics to demonstrate that the RSS as a

whole is online and also each individual operator is operating according to those statistics.

The last one says policy, which is we take this information and we build policy into it. For example, there could be a policy that we would create to handle grievances because the RSOs will be held to a particular standard. So, there should be avenues to handle grievances as well. So, there are three distinct activities that occur within SAPF. Next slide, please.

The designation removal function is simply an activity that says it's time to designate a new operator for some good reason, or there might be a reason to remove them, and at any given time, a set of operators looks – you compute that by looking at the number of designations and removals and that gives you the total set of operators.

The next function is about performance monitoring and measurement function, as I said earlier. We would hold the RSOs to a standard of service and the overall RSS would be computed in aggregate. So, this is just a smattering of things that could happen.

On the left-hand side, we say system capacity is something that should be measured. And by the way, we do do this today, but we will add some more form and structure around this. Bandwidth. Queries. Process. Just examples of that.

On the extreme right, you aggregate all that and it gives you a sense of the health of the RSS.

What you see in the middle is what we call the non-technical parameters, as ethos. In other words, we speak to the integrity of the operators and their commitment to offer the service and the financials as well, which is how healthy are their finances. Next slide.

The financial function, thus far the root server system operations is an unfunded mandate that the current operators bare. There is no funding. It is born by our parent organizations and it is what we call the forgotten cost of the DNS system. We do believe that this is a service that ought to be funded. It's critical. It upholds a billion-dollar enterprise that spans the global community, the global Internet.

So, there should be an option to receive funding with service-level expectations attached to it. The funding should be sourced from stakeholders and related parties, those that depend upon the service. The funding should support four different activities, RSS operations, emergencies, R&D and of course implementation of this model.

This breaks the pie into five different areas where the money would get spent, and the first gray slice says there has been a [inaudible] born thus far of getting to this model. A small

amount of money, but nevertheless, it needs to be accounted for. Then set aside for emergencies in the future, the big 800-pound gorilla is the RSS operations itself, and implementing the model will come with a cost and set aside for R&D.

A question that the ICANN board asked us is could we cost the system out? How much do we believe it will cost to implement the system in aggregate? Rather than give them a number, we say what if we give you a methodology and then tease the methodology out and then potentially use it to come up with a number?

So, this is something. It's a new indicator where I'm christening it and calling it the BPQ. It's a capacity indicator for the RSS. Essentially, it says take the bandwidth, add to it the packets per second, and the queries per second and you get a number called the BPQ. That we call the indicator for the RSS. And to determine the cost of that value, let's look at some industry standard cost determination methodologies, add those formulas, and come up with an actual dollar value for that BPQ.

The next slide essentially says the ICANN board is committed to ensuring that this service runs well, and they would like to be able to defend the methodology for the amount of money that's put into operating the service and also explaining it to any government or any entity that would like to know how and why

and how we came up with the number and what level of risk was accounted for.

So, we proposed that you take the cost for the value of BPQ and add to it a certain cost of risk. In other words, by adding the cost, you may increase the amount of bandwidth or you extrapolate the size of the queries process per second.

The green circle, essentially, it says that's your cost of risk. You add that and you come up a potential estimated cost of the model.

Essentially, this brings together what we call the governance model, a proposed governance model. Next slide, please.

We proposed three steps to the ICANN board to implement the model, what to do with it next. This has been submitted. So, we say that the ICANN board initiates some kind of a process to come up with the final version of this model. In other words, RSSAC 037 is presented to this community as simply a starting point. We firmly believe we need more eyes that need to look into this model. There are still some questions that remain unanswered and there might be gaps in it. There should be a process to look at this more deeply and then come up with a final version.

Recommendation two is to look into section 5.5.3, which essentially says costed out. We've provided a methodology and we costed out. Recommendation three says one these two steps are done, we highly recommended that this model be implemented.

As I said earlier, the model was teased out with various scenarios and I'm going to turn it over to my colleague, Brad Verd, who will walk us through that. Thank you.

BRAD VERD:

Good afternoon. This is where you could say the rubber hits the road. We spent a significant amount of time on the model and then the fun part about the scenarios here is this is the evidence that kind of tested the model itself. So, we ran through a number of scenarios about adding and removing root operators to see if the model worked.

You have to remember, when this work started, all this work was based upon that question. How do we add, how do we remove, a root server operator, since those processes were never documented and not in place, and what we ended up with is what you see in RSSAC 37.

So, the group produced five scenarios, designation scenario, which is the obvious one. How do you designate a new root

operator? Then, a couple of different scenarios on how to remove one. These are not inclusive. These aren't every scenario that we could think of. We just wanted to cover a couple of them. And based upon time, we'll see if we cover all of them, but I think it's important that we certainly cover a designation and a removal, just so you guys can see where we ended up. Jump right into it here.

The first one we have a designation scenario. You have to assume that there's a need for a root operator. So, that's just a foreground conclusion, and this scenario could be for any number of different sometimes, but let's just say we need a root operator.

Using this model, the strategic architecture policy function basically would identify that. Is there a need for one? If there was an opportunity to add one, they would task what we've identified is the DRF function, the Designation Removal Function. You can kind of see we've made this as a kind of flow chart for you. Is there a need to designate, yes or now? Obviously, we'll go through the yes. There's a call for applications that's sent out from the DRF to the people, any interested parties who want to be an RSO. The RSOs apply, obviously. Applications are received back by the DRF.

Those identities, those candidate RSOs, are essentially turned over to the PMMF and based upon criteria previously defined, which is not in here but those details would be in RSSAC 37. You don't see them here in the scenario, but based upon criteria given to the PMMF, measurements, qualifications, testing, what happened to see if a candidate RSO qualifies.

Then, obviously, we go back to the DRF. All the evaluations are done. Reports are aggregated. And the DRF evaluates its candidate pool and then makes a recommendation.

The recommendation would essentially go to the ICANN board for approval. Based upon the ICANN board's approval, a contract would happen between ICANN and the candidate RSO. The candidate RSO would be added to the secretariat function as they would now become a root operator.

Then, the most important part, IANA would be notified and the root zone would be updated. Those three files that were identified on the model, the root [inaudible] file, root servers [dot-net] zone and the root zone itself.

So, that's a simple designation scenario as it fits in with the model. Let's jump ahead one more and then I'll open it up to some questions. Go forward here. This is voluntary resignation. This is the idea that a current root operator would like to step down, so they would notify the strategic architecture policy

function that they plan on stepping down. There's some back and forth on confirming the intent and they sign a formal letter, some formalities back and forth.

The formal request to resign is accepted, and then there is ... So, the SAPF, there's a whole number of things that are kicked off here. First and foremost, it will notify the DRF that we're in need of replacing or adding one. So, we go right back to the designation piece that I just walked you through. But there is a back and forth here getting the existing one to step down.

So, there's a notification sent out to the secretariat that one is leaving. Obviously, notification of the contract holder that one is stepping down. Then, to the PMMF that they will no longer be needing to monitor this root operator once they transition out.

As we go back to the SAPF, the SAPF, Strategic Architecture Policy. This is the group that will be defining those age-old questions that we've been asked of how many root operators type of thing.

So, in a number of different scenarios, there might not be a need to add one. We might be safe. Everything's fine. We can keep on humming along. In this case, let's just agree that there is a need, in which case they would task the DRF to initiate the designation process.

Going back to the resignation, the DRF – again, they’re a designation and removal function. So, if somebody is stepping down, we need to remove them. The DRF would make a recommendation to the ICANN board and essentially, everything we saw with the designation piece is done here, but rather than adding them, we are removing them. So, we are ... The ICANN board would instruct to terminate the contract with the root operator. the secretariat would remove the root operator and obviously instruct IANA to remove them from the three root zone sources filed. So, that’s a really quick run through of a resignation.

We have three other scenarios for performance. Catastrophic shut-down and rogue operator that touch the other functions. I’m happy to run through them really quickly and then hopefully there are some feedback and some dialogue here.

We wanted to show how all the different functions were engaged through the different process. So, stepping ... I’m sorry. So, performance. This one, as you might imagine, would be, the PMMF, the Performance Monitoring and Measurement Function, that function, again, in RSSAC 37 is continuously monitoring the health of the root server system and the health of the root server operators. The PMMF identifies that there’s a clear sign that an SO is failing. They send a notification with supporting data to the SAPF where [inaudible] has happened.

Again, lots of detail in the document. This is a high-level flow of how it would work. But, obviously, if the PMMF had basically there was no possibility for curing the problem with the RSO and the strategic architecture policy group agreed, then it would be sent to the designation and removal function to study the data, try to cure that. You can see, are they in a critical state, yes or no? There's some back and forth trying to cure that whatever is amiss with the root server operator. Can improvements be made? And bring them back up to snuff.

As we end up with the DRF again, you can see in the bottom, flow decision there. Has the RSO improved by any certain date? Obviously, we don't want any root operator being poor-performing in perpetuity. Obviously, the decision is no here and we jump right back into making a recommendation to the ICANN board and everything that was identified earlier is gone through again.

So, you can kind of see that once you've been through one or two of these, there's a lot of repeating functions. But, with that, I'm not sure I need to go through all of them in detail and say it over and over again. They are in the slide deck. They are in RSSAC 37 for your review. With that, we will open it up to questions. Let me jump through here. David?

DAVID CONRAD:

David Conrad, ICANN CTO. So, speaking a bit personally, I've been involved with the root server system since like 1998 when I was working what's now Internet Systems Consortium. I've observed the Root Server System Advisory Committee since its inception. In the past, as some of you may know, I have been somewhat critical of the root server operators and the root server system and I take it all back. I have been absolutely truly impressed with the cooperation, the collaboration, the work that has been done by the root server operators and the Root Server System Advisory Committee and am just quite blown away by the quality of the work and the quality of the working relationships that have been formed.

I think this is, as Tripti has said, a starting point. I think it has some amazing aspects to it that are very innovative and look forward to the community discussions as the community figures out how to move forward with putting in a governance structure on something that grew sort of organically over time.

But, I just have to say for the record and publicly that you all have done an amazing job, and I apologize for the nasty things I said in the past.

BRAD VERD:

They were motivation, David. Thank you.

LARS LIMAN: Apology accepted.

UNIDENTIFIED MALE: Thanks, Brad. I [inaudible] publicly critical about the RSSAC, so I can't just say what David said, and agree with it. But, yes, I do agree this is a very, very impressive piece of work and it's long overdue and I'm delighted to see something of this standard being presented in public. This is great and a great step forward.

I've got one or two small [inaudible], though. Just [inaudible] questions that can be looked at as we discuss this [inaudible]. The first thing is this idea of the designated [inaudible] root operator and the root operator being bound by something of contract with ICANN. It might be worth considering that that contract has got some kind of term limit, rather than something that becomes like the birthright that's been talked about before in the context of the existing RSOs, so that if a contract is awarded to a new root server operator, it will not just continue indefinitely and maybe it would be subject to the view [inaudible] every five years, ten years, whatever. But maybe that should take place [inaudible] not the root server, the new root server operator is doing a good or a bad job. So, maybe that's something that might be worth considering going forward.

The second thing I want to talk about is the issue of the root server operator [inaudible] that's just [inaudible]. The paragraph starts "if one ISO has intentionally misbehaved," I don't think intentionally is the right word to use there because there may be other kinds of RSO misbehavior which is not deliberate or intentional. Existing root server operators are all honorable, competent people. They're not going to do bad things by default. But it could happen by inadvertent accident.

Say, for example, a root server operator is partnered with an Anycast provider and the Anycast provider has screwed up. Or perhaps we've got a scenario that was [inaudible] security hole, somebody hasn't updated their firewall or hasn't updated a router or whatever. So, it's not so much a deliberate act of malfeasance or bad behavior, but it's an accidental thing that's happened. Ultimately, the root server operator is accountable for that, but it may not be a deliberate act that's caused the root [inaudible] behavior.

BRAD VERD:

If I may, really quick, just to come back at you. I don't disagree with the statement. However, the way that the group viewed it – and please, other members, please chime in here – is that a root server operator going rogue is intentional. Using your by definition, using your example of a partner who screwed up or

did something wrong, that would be caught in the performance monitoring and they would be given time to cure, and if they were just not curing it because they couldn't do it for whatever reason, then they would go through a removal. But, a rogue operator in this scenario was somebody intentionally trying to do that.

TRIPTI SINHA:

Just to add to what Brad just said, those scenarios are not exhaustive. So, we just took a couple of obvious ones and the scenario you described is certainly possible. Yes, when this model is teased out, that would be accounted for.

The first point that you made that this should not be a birthright once the contract is in place, we completely agree the details are to be determined in the future, so we made no assumptions about that. So, that's a very good point. And when this is vetted by the community, those are the kinds of things that need to be build in. So, thank you.

MARK SVANCAREK:

Mark from Microsoft. Regarding BPQ, I had some questions about BPQ. I had trouble. There's this lack of knowledge on my part, I think. What is the range of values of BPQ? Are there some that are very small and very large? I wasn't sure why it was being

quantified, what it would look like, and what action you would take based on it.

LARS LIMAN:

Lars Liman from Netnod here. The BPQ thing is, in my view, I should say, is a symbolic thing. We need to find some values that we can use for measuring. This is not a cooked idea. So, these are things that we know that we look at when we try to assess volumes and traffic patterns and so on.

The thought is to bring those and other parameters into a monitoring model that we can use. So, it needs a lot more discussion to find out exactly how to use this, but these, for example, are parameters that become use for monitoring. So, not cooked yet.

BRAD VERD:

If I can add to it real quick, the cost of a service as far as bandwidth, as far as servers and whatnot, is a calculable cost. But, that's never the question that is asked of RSSAC or the root server operators when it comes to the service. The question is: can you survive a one terabyte, two terabyte, six terabyte attack? It's important. That is the multiplier, the risk multiplier. What is the risk that we are willing to accept? That is the

multiplier that comes up with the true cost. That's really important as the takeaway here.

LARS LIMAN: David?

DAVID CONRAD: So, just to clarify, BPQ isn't actually a single-value metric, right? It's actually a composite of multiple metrics that can be measured independently. So, given two RSOs, the Bs can be the same, but the Ps and Qs could be different, right?

UNIDENTIFIED MALE: Yeah. That was clear from the presentation, but I didn't have a sense that one operator would have a BPQ, say, ten times greater than the other, if that was likely or if that was significant. I was just wondering. I needed more context.

KAVEH RANJBAR: So, the idea is not to measure BPQ [inaudible] operators. This is for the whole system, for the root server system. So, this is something that we can say, okay, the current capacity of root server system has this [BMP and Q] and then we can also – that also gives us another capability to go and ask ICANN board or IETF or any group of experts, what do you think is a good BPQ for

the whole root server system? They will come up with a number. That will enable us to cost it. And also [inaudible] functions, we can say if we add this three new operators, we will actually get to that level, which we all globally think that's a good level capacity for the whole root server system. It's not about individual operators or comparing them. It will give a model to be able to cost where we are and where we want to be.

BRAD VERD:

We have Wes, Russ, and the gentleman in the back there.

WES HARDAKER:

I'd like to remind us that the whole notion of BBQ is that it's not supposed to be an exhaustive monitoring and measurement system. It is a high-level system in order to make some architectural and design decisions. I think if you ask the operators in the room, you'll find we monitor a whole lot more than that. This is just a way of how can we establish a framework for discussion around this system.

RUSS MUNDY:

Thanks. Russ Mundy, SSAC. One of the things that I think is really important about this document as a whole, and the BPQ sort of is a way to illustrate that, is that in putting it together, there was very conscious effort made to think about the system as a

whole. This is a measure, although it was established as a starting point primarily for costing, it also gives a way to achieve a quantifiable, measurable, over time – at least that’s the objective – way to look at the system as a whole. Is it getting better? Is it degrading? How much bigger does it need to be? So, it gives some possibility of quantification for the system as a whole, which we have never had in the past.

BRAD VERD: Right here.

SEBASTIEN BACHOLLET: First of all, thank you very much for presenting that to the community. Maybe the question I will ask will show my deep non-understanding of what your work day to day. But, I wanted to ask two questions. The first one is that all that is based on the fact that we are currently [inaudible] root server and it will stay [inaudible] number, or part of the scenario you are putting, if one day there is possibility to have more root servers at this scenario could be used to have new people to run those servers.

My second is a suggestion on the slide where you have, who runs a root server? You have 12. It could be useful when you present that to a group who has less knowledge than you to put the letter of the root server who they are running. Of course, there is

one running two I guess. It could be useful to understand better what is at stake. Thank you.

BRAD VERD:

Really quick, I'll go first, and then Tripti can add regarding this process for adding and removing. Right now, this does not address the question of more than 13 or less than 13. What is that number? What should it be? There's a lot of questions. Can we add more? There's questions of do we need more? Do we need less? Maybe 13 isn't the right number.

But, this presents a process that allows us to do that. Once the work is teased out in the SAPF function, the Strategic Architecture and Probably Function, as to what that number is, we can go forward and do that. Tripti?

TRIPTI SINHA:

Just to add some more context, we introduced BPQ. So, the approach we're taking now is to determine the overall aggregate capacity of the system. Let's set aside letters and operators and so forth. What is the desired capacity of the system for this ever-growing Internet?

Then, from that, say, alright, what kind of operational infrastructure needs to be put in place to deliver the service? The letters, by design, we have deliberately removed from our

psyche. I'll tell you why. It's an internal identifier. It's a piece of technical information that helps you build the infrastructure.

This was an organic growth, and for some reason, we began to call each other by each other's letters. We should not have done that. So, we're trying to correct the mistakes of the past.

When the SAP function, as Brad just said, determines what exactly the capacity of the system should be and what kind of technical identifiers will help meet that, then we will decide how many operators could potentially deliver this service. And it could be less than 12 and we believe we will likely end there. It could maybe be more. We don't know yet.

But, we are by design moving away from using letters. That's just an internal technical identifier and that's not the right way to look at this. So, we're correcting the model, the errors. The organic growth of the past needs to be behind us now and let's do this right. I hope that explains it thank you.

BRAD VERD: Jim, you wanted to add?

JIM PRENDERGAST: 8Yeah. Thank you. It's just going off on a slight tangent for that, but how does 7706 fit in with this model where we've potentially

got private instances of the root running locally and local resolvers or perhaps even running on my own laptop? Would this be part of this overall architecture that we're talking about here or would that be completely separate?

BRAD VERD: Terry?

TERRY MANDERSON: Terry Manderson, ICANN. To me, 7706 doesn't really play into that because 7706 is about distributing content, not actually dealing with root server, the root server operators themselves. 7706 does two things. One, it extends a caching ability. The other is it provides a privacy function. That doesn't necessarily contribute or take away from the efforts within the larger root server system.

BRAD VERD: Alright. Any other questions, comments, or feedback?

LARS LIMAN: Thanks for this, not being very technically [inaudible]. I've been very impressed by seeing how the root server community has come together to work towards [inaudible] more resilient Internet without losing the autonomous part of the root server.

For the BPQ [inaudible] I can see the [inaudible] for having any indicator. What I don't understand is how the BPQ would be a better indicator than having three indicators, one for B, one for P, and one for Q, because for me, I'm not sure whether you're counting apples, bananas, and pears or whether it's really all apples. Can you explain that?

UNIDENTIFIED MALE:

Yes. It is three different types of fruits. So, we cannot sum them up. That's why we measure them separately, because you might have, for example, a lot bandwidth, but you don't get that many queries, for example. And there is a lot of other kinds of traffic. So, we need all this [inaudible]. We found [inaudible] three separate critical pieces to count, basically, to be able to measure this.

BRAD VERD:

Yes?

SEBASTIEN BACHOLLET:

Thank you. I wanted to ask you, now that you are presenting that to an open meeting, how you imagine to deal exchange with the community at a later stage, when are you ready to do that, to go to SO and ACs to discuss this issue yet, and when it

will be the right time, how will you imagine this exchange?
Thank you.

BRAD VERD:

Right now, we're following the normal ICANN process, which is we've provided the board with advice and they will do their normal evaluation of it. Right now, it's in their hands as to what happens next with this.

We, RSSAC, expect this to lead to a larger discussion. That's what this RSSAC 37 was designed to be a framework for a larger discussion to tease out all these deep details that we are talking about here. But, right now, it's in the hands of the board to what happens next.

UNIDENTIFIED MALE:

To add a bit to that, yes, that advice was published ten days ago or so, and so now ICANN Org has received it and registered it in the [ARO] database and acknowledged by the board. Board Technical Committee has received it. Next step is Board Technical Committee will go through that, come up with a proposed what board plans to do. I guess it might come into consultation, but we will see. And based on that, the next steps will be decided, but it's not yet ... There has been no action from the board yet. It will happen in a short time and obviously it will

be communicated. But, right now, the next action is expected to be from the board.

SEBASTIEN BACHOLLET: If you'll allow me just a comment, I know that you are following the process, and okay, that's the process. Now, your goal, if I understand well, it's to have the community involved and the community involvement can't be just with the board. Therefore, I suggest, but it's just a suggestion, that the next meeting there are some exchange organized with the SO and AC who are willing to have the discussion with you, even with the change document because the board will have said we would like to have these types of changes. Thank you.

UNIDENTIFIED MALE: So, yes. Anyways we were planning to engage in Barcelona, but it won't be formal engagement like [inaudible] to affect how it will go on, but definitely it was in our plan to engage with different constituencies in Barcelona.

TRIPTI SINHA: Sir, if you were referring to socializing this with the community, yes, that is in the plan. But, the formal implementation of getting another set of eyes to look deeper into RSSAC 37, we're waiting for the board to give us a course of action for that. But, we do

take that into consideration and we had it in the plans. Thank you.

BRAD VERD:

Any further questions? Alright. With no further questions, there's nothing in the Adobe room, I assume. Okay. Then I thank you for your time and we are adjourned.

[END OF TRANSCRIPTION]