

---

KOBE – RSSAC Information Session  
Tuesday, March 12, 2019 – 09:00 to 10:15 JST  
ICANN64 | Kobe, Japan

FRED BAKER: ... Is an information session, so the target audience is the people sitting over there, and there's a good chance that it will be a long dry session if you all go to sleep. So, I'm looking forward to dialogue and questions as we go through this. If you're looking at a slide, and you find yourself saying, "What in the world was that?" we need to know.

So, now to get started, let me go around the table which is the people that are involved with the RSSAC and get people to just introduce themselves. Jeff, can we start with you?

JEFF OSBORN: Jeff Osborn, ISC. I'm on the RSSAC for the F-root.

KARL REUSS: Carl Reuss, University of Maryland for D-root.

KEN RENARD: Ken Renard, Army Research Lab, H.

---

**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.**

---

RYAN STEVENSON: Ryan Stevenson, Department of Defense, G.

LARS-JOHAN LIMAN: Lars-Johan Liman, NetNode, I.

OZAN SAHIN: Ozan Sahin, RSSAC support staff.

CARLOS REYES: Carlos Reyes, RSSAC support staff.

MARIO ALEMAN: Mario Aleman, RSSAC support staff.

FRED BAKER: So, Fred Baker, ISC along with Jeff, and also RSSAC co-chair.

BRAD VERD: Brad Verd, Verisign, RSSAC co-chair.

KAVEH RANJBAR: Kaveh Ranjbar, RIPE NCC, and RIPE NCC root server operations.

WES HARDAKER: Wes Hardaker, USC ISI.

---

TERRY MANDERSON: Terry Manderson, ICANN, ICANN Managed Root Server.

DANIEL MIGAULT: Daniel Migault, IAB liaison to RSSAC.

DUANE WESSELS: Duane Wessels from Verisign as the root zone maintainer liaison to RSSAC.

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

FRED BAKER: So, we have a slide deck to run through, basically discussing what the RSSAC is involved in right now and what we've done recently and how the system works. Let's go through that. Mario, can you give me the next slide?

So, we're going to give your four basic pieces of information. One of them is an overview of the RSSAC and with it the root server system, discuss a document that we recently produced, RSSAC 37, which is a proposed governance model and talk a little bit about the conversation that we're having with ICANN right now about where that goes.

---

The third thing we'll talk about is the current RSSAC caucus work, and of course in the RSSAC overview I'm going to tell you what the caucus is, but this is some projects that we're working on. And then in the fourth section we'll talk about community interaction.

So, RSSAC, who are we and what are we about? The Root Server Advisory Committee is created in the ICANN bylaws and our job is essentially to advise the ICANN community and the board on matters relating to the root server system.

That's actually a very narrow scope, and so we fairly frequently get questions that come our direction and we say, "That's not about the operation of the root server. That's not what we do." But the people around the table told you what companies they were from and kind of where they fit in things. Understand that they are actually representatives from the operators of the root server system which is to say the distribution of the data that is in a repository at the Internet Assigned Numbers Authority. Next slide.

So, we are appointed representatives from the 12 root server operators, and we have alternates to these, so nominally 24 people there. And then a number of liaisons to and from different organizations. I can say a little bit about liaisons. People, as mentioned in the recent review of the RSSAC, that people kind of had a hard time figuring out how to talk with the RSSAC or

---

whatever. Well, you know, this is a way or a common way of how one does that. So, for example, Russ is a liaison from the Security System Advisory Committee. Daniel is a liaison from the Internet Architecture Board, and we welcome having liaisons from organizations that are interested in what we do.

The RSSAC Caucus, on the other hand, is a body of volunteer subject matter experts. We are all members of the caucus, but most of the members of the caucus do not work for an RSO. They're people that work in the DNS technology in other ways. The statement here is that they are appointed by RSSAC. I have a little bit of frustration with the word "appointed." It sounds like we thought of them and said let's drag this guy in. No, that's not how that works. They come in and they say, "We'd like to be here." And we say, "Oh, okay, that's great. Yeah, yeah." So, they're accepted and so they're adopted by RSSAC. Next slide.

So, the caucus is over 100 people and we changed this to over 100 because every time we look at the number the number changes, but it's over 100 people that are experts on the various aspects of the DNS and do things in the DNS. When people want to become a member they make a public statement of interest that kind of says who they are and why they care, and then we, in turn ... It would be easy for us to say, "Well, this work was done by the caucus and the RSSAC will take credit for it," or something like that. We don't do that. If people are involved in doing work, we

---

want to give them credit for the work that they do. So, they get credit for their work.

The purpose of the caucus is in essence a pool of experts. Later on, I'm going to talk about one of the caucus committees that I'm involved with which is trying to figure out how resolvers really work, and how best to use them in the Internet system. The people that we are bringing in are people that are involved with resolvers and are working with the software in that area and are trying hard to make things happen, to make measurements happen. So, bringing in people that can help us, frankly, do a better job.

So, we try to be transparent as to who does the actual work but we provide a framework for getting work done, which is to say that we expect something to happen and we say so in the form of deadlines. If you'd like to become a member of the caucus – and you're welcome to do that – please drop a note to [rssac-membership@icann.org](mailto:rssac-membership@icann.org). You will get a reply that kind of says, “Who are you?” And you'll be asked to provide one of those statements of interest and identify the expertise that you bring to the party. I don't think we've actually refused anybody. We'll think about it for five minutes and then go ahead. So, next slide.

The caucus meets at every other IETF meeting, the even-numbered IETF meetings, and in-between them the caucus meets

---

here at ICANN meetings. So, the caucus will actually be having a meeting in two weeks at Prague at IETF 104, and there we will also have at least two sessions where two of the work parties get together and talk about what it is that we have going on.

BRAD VERD:

One minor correction there, and this is at the request of the caucus, we meet at every even IETF, and then there is a formal caucus meeting at the AGM here at ICANN. We do have periodic work sessions that will take place, based upon the work party that is happening that would include the caucus here at non-AGM meetings.

FRED BAKER:

Thank you. Okay. Next slide. So, what I've just done is gone over who are we, why are we here and what extensions do we have. Let me ask: do you have questions? I'm seeing none.

So, what we have been working for the last three years is a proposed governance model for the DNS root server system. The DNS root server system, and there is a document that describes the history, originated in 1984 with the development of this brand new protocol called DNS, and I suppose we needed a way to make a run, so we as a community put up a couple of root servers and had people start asking into them. Within a year, that became

four different servers. By 1998, it had become 13 individual servers, and the companies, the people that operated those servers were basically friends of Jon Postel. Now, then in 1998 – I guess it was 1998 – Jon died and so one of the questions has been how do we and why would we add an operator or change operators in any way, do something to evolve that model?

And, well, I suppose we could get out a Ouija board, but having a séance to decide whether somebody should be a root operator just seems like a bad way to approach things. So, the question that Steve Crocker asked the RSSAC at its first meeting, over dinner, and which we've been basically working on since is, “Okay, what should we do instead?”

So, we've been working for three years on a process – a governance model – and trying to sort out exactly how a next generation root server system should work and where we fit into that. So, that includes what I'm about to describe, the 11 principles for the operation of the root server system. How do we do it and why do we do it?

It proposes an initial governance model for the root server system and I want to underline the word "initial". It's how we think it should work, but we'll probably learn something somewhere along the way and we'll need to make that adjustment, whatever it is. And then it describes how that model works through a set of



---

scenarios based on the designation or removal of operators. Next slide.

What the system is that we're talking about is that blue thing called resolution on the slide in front of you. I found myself literally drawing on a piece of paper and explaining this to somebody the other day. I wish I had this slide at the time I was doing that. But the guys question was, “My favorite country operates ccTLD. Do we have an authoritative root server that I should be going back to for information?” And, well, that's not how it works.

The way it works is the various TLDs, of which we have about 1,530 now, and 250 which are country code TLDs and the rest of which are generic, they are approved by the ICANN board as a function of the process of the ccNSO or the GNSO and the things that they do. They are identified by the board then to the Internet Assigned Number Authority, which is a secretarial office, operated by an ICANN subsidiary called PTI.

The TLDs will now tell the IANA, “Hi, my character string is this. I'm located at that address. If you have questions, call me at this phone number,” and all of that information. That goes into the brown bubble called the IANA function. They collect that information into a database, and then a contract person, called the root zone maintainer, pulls that information out of the

database, collates it into the root zone files, of which there are three, and makes them available to the root server system operators.

What we do, as operators, then, is that we go get those files, and they change, what? Is it six times a day, I think, or they potentially change six times a day, and so we have to go and make sure that we've got the latest one. Then we, in turn, distribute that among our own servers, and then from our servers to anyone that asks.

Now, to be very honest, the principle people that ask are DNS resolvers, which is open DNS, 8.8.8. There are probably a million different services that provide DNS resolution for people in an area or whatever, but they will come back to the root zone. What's the information for this name? Please give out the resource records so that we can give it to whoever it is that's asking us.

Now, I said earlier that the history of the root zone, we started out in 1984 and by 1985 we had four servers at four different addresses and by 1998 we had 13 servers at 13 different addresses. What has happened since then is that each of the operators, each of the RSOs, has then expanded their system to have multiple computers that are accessible at the same address. They also have their own private address which we use for managing them, but if you go ask at the address that it's associated with – to pick one of the top of my head, A – you know

---

you go to Verisign, and you say, “How do I get to example.com?” And you'll talk to whichever server router thinks is the best one for you to ask, and you'll get an answer. You would get the same answer no matter who you ask, because we're all working for the same source database.

You can also verify that the answer that you got was the answer that you should have gotten, the one that IANA thinks you should have gotten, because it comes with a signature and you can verify the signature against the DNS key of IANA, but now that means that we've got, right now, over 1,000 servers that you might as, “How do I get to example.com?” And each of them would deliver the same information and do so in a timely fashion. So, that's the structure of the root zone.

Now, we try to operate – or we do operate – by 11 fundamental principles, and the first one is that we think that there should be one root zone. Why? Because if we have multiple root zones, they will probably have different answers, and then things will get really confused. You'll ask to go to Google and you'll wind up at [BYDU] or something like that. So, we want to maintain a single globally unique public name space and then deliver that for everybody to use.

Second, the IANA is the source of that data and that's just gospel with us. That's just where get it and that's just what we hand out.

---

For the Internet to operate reliably our services must be stable and reliable for all users, whether we like them or not, whether they agree with us or not. We provide a simple, consistent, reliable and resilient platform for all users.

Now, one of the interesting things, as you go around the table, is that we do things in different ways. We use different hardware, we use different operating systems, we deploy different DNS software. We are from different companies and different kinds of companies, everything from the Army Research Lab which is obviously somehow related to the U.S. Government, to in ISE's case a 501(3)c corporation. There's ICANN, there's Verisign and it's a for-profit company. We're very diverse, and we consider that a strength. We consider that a good thing. In fact, sometimes we will find ourselves looking at, "Gee, I want to try this new thing. How do I want to go about doing that?" There's five ways, and we pick one that is different than everybody else's.

Why do we do that? Because we like pain ... I don't know. But we consider that a strength because if one of the systems can be attacked, if we find that there's a way to take down BIND, or Unbound, or FreeBSD, or whatever it might be. If we find that a particular operating system, or a particular hardware platform has a bug in the firmware, we don't want that to be a point of attack where one could attack the whole system. We want to be able to say, "Fine. That company, or that piece of hardware, isn't

---

working very well today. We'll go fix that and come back.” But then maintain that stable, reliable and resilient platform, guarded in part by the diversity of the systems and the companies that are at work.

Now, change control. Architectural changes happen. They happen, actually, often. We're forever talking about possible changes, and we tend to think that when change happens – and notice I don't say if change happens, it's when change happens – that it should be something that we're aware of, that we discuss, and that there's a demonstrated technical need about them. Where did these specifications come from that we operate on? The DNS specification is maintained by the IETF, and so we operate on those specifications. Next slide.

Now, obviously, a lot of the integrity and the reliability of the system relies on the integrity and the reliability of the people and the companies that operate it. We believe that the RSOs have to operate with integrity and with a commitment to the common good of the Internet.

I want to underline that. That can get to the point of being fighting words. If you want to do something that is not to the common good of the Internet, you'll probably have a bunch of us pouncing on you saying, “That's a bad idea. Let's not do that.” So, we really

---

try very hard to operate with integrity and with an ethos targeting the common good of the Internet.

Now, an important thing is the RSOs have to be transparent, at least to a degree. Now, obviously, we're operated by companies and there are some things that shouldn't go outside any given company, and so we're not going to be completely transparent in the sense of “so we had our board meeting and foo and bar were agreed to” or whatever, but we do capture statistics about ourselves. We talk about the structure of our deployment and we are quite willing to talk with people about the issues that might happen. We really try to be transparent. This meeting is an example of us trying to be transparent.

It's also an example of us engaging with the stakeholder community. You guys are all stakeholders, so here we are, collaborating and engaging with you.

Now, at the same time, RSOs have to be autonomous and independent. They need to be transparent. We need to understand what's happening in the different RSOs, but it's not like somebody gets to say, “You're wrong. You can't do that.” The RSOs have to be autonomous. They have to be different companies. And they need to be able to operate independently in order to evolve and to deliver the service that we do with the stability that we'd like to deliver that. And then the RSOs have to

---

be neutral and impartial, and there we're primarily talking about it with respect to the data that we're carrying.

If you think in terms of pipes and water or wires and electricity, we carry the water from the IANA resolver and we just get it there. We're not sorting through stuff as it's going by. So, we must be neutral and impartial. Next slide.

Let me back up. So, do we have any questions about the principles that I just discussed? Seeing none.

So, RSSAC 37, the model that we have spent the last three years talking about and are now discussing with ICANN and the ICANN community, we started out with kind of a mental picture of how the system works and what the different parts are and tried to separate them into different components that we could talk about rationally. Part of that is who are we doing this for? Well, in a certain sense, we're doing it for the resolvers, for the average Internet user, and so on and so forth, but the representatives of those people are the ICANN community who is in turn, in a large part, represented by the board; the IETF and the IAB, from a technical side; and of course the companies themselves, the RSOs.

What we decided was that if we were to reorganize the system and go into something resembling a next generation, most of the function that RSSAC does, besides actually writing documents,

---

has to do with strategy architecture and policy, and we figured that needed to continue to be on the table and to exist, and so we described something that we call the Strategy Architecture and Policy Functions (SAPF) as kind of a central function to talk about, to discuss, to move forward the things that are going on in this part of the world.

When it comes to the question of the designation and/or removal of an operator, we kind of have to do so based on data. We need to have a reason for what we're doing, and so we started out with the assumption that there would be measurements of some kind, and they might be accounting measurements, they might be counting measurements, but various kinds of measurements that would be taken of the system.

So, we identified a part of things that we call the PMMF, and that's the box to the right there, the blue box, which carries on that monitoring, and then that data might be reviewed by the strategy people and they say, "Well, it appears that there is a problem and somebody needs to figure out what they're doing." And that might be that the country of Lower Slobbovia has a hard time getting access to the root. Or, it might mean that there's a particular operator that's not performing, or something else, and they have to go and think about it, and decide what they think about.



One of the options that they might have would be to say we really need to allow for a new operator, or we need to make that one, whichever one that is, go away. And so, we have a function called the designation and removal function which basically prepares recommendations for the ICANN board with community input with the community support, hopefully, and then they would go to the board and the board would add or drop an operator.

Somewhere in the course of that, there's money – surprise, surprise, there always is. So we have an accountant. We call that the financial function. And by the way, there needs to be a means of carrying messages around, so on and so forth, and so we have a secretariat function. So, RSSAC 37 basically focuses on those five components and says they must exist to find a way to do things.

Now, at that point, we go pretty much directly into the root operations. The RSOs are obviously the operators, and there's a question of performance metrics. How are they doing? Are they actually delivering the service that they're expected to deliver? Are they up to snuff? If they are, that's a good thing and we probably need to say that they're doing well. If there's a problem, they need to say that there's a problem. And that feeds into ... And as I mentioned, the designation and removal function might then designate a new operator or remove an operator.

---

A question that we get asked, actually pretty frequently, is, “Why can't my country ...” Whatever "my country" is, or my company, or my/us have a letter? Please send a letter, and the board gets asked that, we get asked that, and the answer really is that there's no fundamental reason that we have to have as many as 13 letters, as many as 13 companies. There's no reason that we're limited to 13. We could have more than that.

There's probably some upper bound – and I don't know that I can definitely say what that upper bound would be. But this is put together basically to allow the community to have the flexibility to deal with those questions, so should we add, should we drop, and how does that happen, what is the reason for that to happen. And then from the company's perspective, the RSO's perspective, we really hope that the reason is a technical one. Somebody walking in and plopping down a billion dollars and saying, “I want to buy a letter,” isn't a good reason to make them an operator. Next slide.

BRAD VERD:

Before you go to the next slide, two things. One, while we could do more than 13, the right answer might be less than 13. That might be the correct answer. I just want to put that out there. Secondly, I assume a lot of people have seen the model. They've read RSSAC 37 which goes into deep detail on the model, but I just

---

want to remind everybody how the work started. The work started with the simple question of: how do you add a root server operator?

While that seems like a simple question, when you dive into it, it gets very complicated because if you add one, you need to know how to remove one. If you're going to add or remove one, you need to have policies and procedures around doing that. You need to have a governance around it. You need to have a means of measuring. So, you can kind of quickly see how this grew into what it did, and I just want to remind everybody on the impetus of what it was. So, thank you.

FRED BAKER:

Thank you Brad. So, now, along with RSSAC 37, the document that describes the model that I just talked about, we sent a companion which was the RSSAC 38, and that, if you will, was an executive letter to the board, saying, “Please read 37.” And so we wanted them to do certain things, and one of them is, “Could we start talking about this?” And, well, RSSAC is right now in conversation with ICANN about that, and there are various things going on with that.

We needed to estimate the cost of the RSS and develop the model. Frankly, with the exception of Jeff, none of the people sitting around the room are money people. So, we needed to have

---

people that actually understood money to go look at this and make some sense about it, and then figure out how to implement that version and model, based on the principles that we discussed earlier. Next slide.

So, I've gone through now kind of the big work that we've been doing for the last three years, and which we're now going to ask the ICANN community to comment on and that will be coming down the pike in not too long. Do we have questions on that discussion?

[ERIN WELLING]:

[Erin Welling], USG. I have a question for you. You mentioned RSS costs in the previous slide. Can you tell me a little bit about the RSS costs, the costs of the RSS, can you talk about the scope of that, what are you including in that, in that estimate?

FRED BAKER:

Well, what are we including in that estimate? The real question is not what it is in the past, but what it's going to be going forward. What are we asking ICANN to spend? And they're kind of curious about that, as you can imagine. And that's not a number that we have crisply, right now, but we expect it to be part of the conversation.

---

[ERIN WELLING]: So, the cost [off mic].

FRED BAKER: The cost to implement and then to do the model. Yeah. Are there any other questions? There was a question over here, no? Okay.

So, now different activities that we have going on. That's where we've been and there's kind of an over-riding activity talking with the ICANN community about the RSSAC 37 model. Oh, by the way, we have a few other things that have been going on. One of them is that we recently had an organizational review, which is mandated in the bylaws for each of us every five years. The independent examiner concluded that review and made several recommendations, six recommendations.

We delivered a feasibility assessment and initial implementation plan for what we're going to do for that review. That is right now being discussed by a board subcommittee on organizational effectiveness. We expect to get an answer from them and have some dialogue with them a little later in the year. So, the OEC is considering the independent examiner's report and the feasibility assessment and the initial implementation plan and is going to come back to the board with next steps, whatever those turn out to be. So, slide.

---

Current work. Primarily in the caucus we have three work parties that are active in the caucus. One is looking at service coverage and is there a country, Lower Slobbovia, that doesn't have adequate service, and if so, what recommendations might we make? Liman, where are you? I don't see Liman.

But this is going on. So, we want to validate or discover indicators that might be used to measure the service coverage and figure out the procedures and tools to do that, and then make some recommendations to the community as to how to ... Improvements that might need to be made, the changes that might need to be made on service coverage. So, one of our fundamental objectives is that we want to deliver the same data all of the time, everywhere, and well, where is "everywhere" and are we in fact getting there?

There you are, Liman. I was looking around for you and must have looked past you. I'm sorry. Do you want to comment on this? You're the [stuck key] here.

LARS-JOHAN LIMAN:

Yes, I know I'm the [stuck key] but I think you gave the basic parts of this, so I don't think I have anything to add.

---

FRED BAKER:

Next slide. Another work item that we have going on in the caucus – and for this one I'm actually the shepherd – studies modern resolver behavior, and the question is how do they really work and are they really configured correctly? Do we have suggestions about how people might want to configure their resolvers, and so on?

So, we're investigating the behavior of recursive DNS resolvers and asking ourselves how they interact with the authoritative servers and with the RSS. To that end, we have two major contributors. One of them is with Paul Hoffman who is sitting over there, who works at ICANN and is doing some work in this area. He's active in the work party. The other is Geoff Huston in APNIC who periodically does different experiments and writes blogs about them. So, we're trying to investigate the behavior.

One of the possible outcomes, and I don't know that it's a probable outcome, but one of the possible outcomes is that we decide the DNS doesn't work all that well and we should even improve it. Or, where we should improve the implementation. We might possibly recommend changes to various and sundry things in order to improve the service itself.

And then, again, on another outcome that might happen – and isn't required to happen, but might happen – would be to craft advice to the ICANN community regarding the stability of the DNS

---

or the root server system. Paul, since you are in the room, do you want to comment at this point?

PAUL HOFFMAN:

Sure. Yes. So, just to be clear that the two parts that Fred had mentioned, the one where we're testing how resolvers work, sort of in a lab, and the part where actually go out and test resolvers that are out in the world, are actually very different things, and we're expecting to get different results from them, and it would be nice to know why.

That is, configuration seems to be very important, so this is sort of ... Some of this work has been done by various researchers in the past, but a very important part of this is that we want our research to be reproducible. And for either of the two parts, if you look at our research and say, "Hey, that doesn't seem to make sense to me," and you were a researcher and you had some capabilities, you would be able to actually reproduce the research yourself.

So, this would go beyond, and RSSAC has made some requests, but this will be an ongoing thing, not just by RSSAC, but by anybody. So, great. Thank you.



---

FRED BAKER:                    Okay, cool. Next slide. And Brad you are actually the shepherd on this, right?

BRAD VERD:                    I am.

FRED BAKER:                    Do you want to tell us about it?

BRAD VERD:                    Yeah. So, we just started. I think we've had one call, one face-to-face meeting on this new work party that was created. This is what I continue to describe as kind of the technical accountability to the root server system and the root server operators. This is where we want to define system-wide measurements that kind of just, quite simply put, are to define what "good" looks like for both the root server system and the root server operators.

This has been a long outstanding issue – or challenge, I should say – in the community, and it's a tough question to answer, but we're going to work exceptionally hard at this and try to come to a final product. This is, in my eyes, and what I've been saying is the second half of 37. So, RSSAC 37 was the governance model. This is the technical accountability piece that plugs into 37. Should it

---

be implemented? Should it not be implemented? This is still the technical accountability piece for service expectations.

And that is ongoing. We encourage anyone and anyone to engage on this. This will be impactful and so we want as much input and opinions as possible. So, right now, I think we're somewhere around 20 people in the work party, and there's a significant amount of work that has to be done. There's two co-chairs that have been identified, Duane Wessels and Russ Mundy, both in the room here, and then we have two shepherds from the RSSAC to help because of the amount of work, and that's myself and Wes Hardaker. So, I look forward to working with everybody on that.

FRED BAKER:

So, we've gone through now the review and the work parties that we have going on. Do people have questions about the work that RSSAC is involved with right now, RSSAC or the RSSAC Caucus? Seeing none, I'm going to move on then.

One of the things that I said we were really trying to do was be transparent. To that end, we have minutes that are posted and various publications. We have, I think, 41 numbered publications. We have a root server system tutorial, usually given by Andrew, who is not in the room. It happens twice at every ICANN meeting, so you can go look at the meetings and find it on the schedule. And then we also have something we describe as the RSSAC

---

Operational Procedures, which is, if you will our version of the bylaws.

Now, if I go to the ICANN bylaws, the ICANN bylaws say things like there will be two co-chairs for RSSAC and it has a certain purpose of its own and so forth. Our operational procedures then go on to talk about the elections and all of that kind of thing. So, those are all documents that are available. They're on Google or they're on the ICANN website. Next slide.

We also, and in this case, we use the root server operators themselves. The RSOs have a fair amount of information that they put out for public consumption. If you go to rootservers.org, there is for example a map there – an interactive map – you can find your favorite country and say what root servers are there here, and which of the organizations are they associated with. At the last meeting I was asked by somebody from Zimbabwe where are the root servers that are near me, and the funny thing was that there were none in Zimbabwe, but there were two in Malawi, which is the country next door and so they got involved in a discussion then of, well, where should they be and do we need something special?

So, you can go to rootsevers.org, you can learn about the system, and investigate it in various ways. You can see which of the operators has a number of systems. They have links to the various

statistics that are supported, delivered by those operators. That information, and of course the agendas for root operator meetings, that information is all at [rootservers.org](http://rootservers.org). If you want to know what's going on in the root, that's a great place to start. There is also a news link there and what's happening recently. You can see that. When there are major events, that shows up in the newsfeed.

I mentioned statistics. The document that we have that describes the statistics that are gathered and reported on the root server operators is RSSAC 2, and what we're describing there is queries per day that come in to the operator in question. That's invites and that kind of thing. The RSOs themselves have webpages and often have other information that they want post, so you can get all that starting from [rootservers.org](http://rootservers.org).

You can also ask questions. If you send an email to [askrssac@icann.org](mailto:askrssac@icann.org), you should get a reply that hopefully answers your questions, or it at least gives you a place to go look to understand the information that you're trying to find.

And let me ask you, we were dinged fairly heavily in the review for not being transparent, and we all kind of did a face palm. We were knocking ourselves out to be transparent. We're having this session, and we're having so on and so forth. Let me ask: are you aware of these transparency items? Do we need to do something

---

to make people more aware? What is missing from the list? Are there things that you think should be there that aren't? And, if so, I'll encourage you to talk to Brad or myself and we can discuss them. So, what would you like us to do to further improve transparency on our part?

So, one more. You can find out about us from the RSSAC main page which you'll find under ICANN.org. We have, as I said, 41 publications. We have an FAQ, and there is also a page for the RSSAC Caucus. At this point I'm going to shut up and somebody else needs to talk.

BRAD VERD: Any questions? Anything at all?

LARS-JOHAN LIMAN: I guess that questions will come from the audience at this stage, yes?

BRAD VERD: Absolutely.

FRED BAKER: You mean I have to? Anyhow, okay. The time is now approaching, I'm sorry, you have a question?

UNIDENTIFIED MALE: Is this slide available on the web? It's not uploaded, so I'd like to see the slide later.

BRAD VERD: Yes, it should be uploaded. If it's not, we'll get it there.

FRED BAKER: Okay, I'm going to give you back a half-hour of your day. I think we're done here, thank you.

BRAD VERD: Thank you, all. We're adjourned.

**[END OF TRANSCRIPTION]**