ICANN79 | CF – Get to Know the ICANN Community: SSAC and RSSAC Members Networking Event
Sunday, March 3, 2024 – 4:30 to 5:30 SJU

SIRANUSH VARDANYAN: Thank you very much. So we are at our last session for today with fellows and next geners to talk about the security and stability aspect of ICANN and this is a great pleasure that I have the interest from RSSAC, which is the Root Server System Advisory Committee and SSAC, which is the Stability and Security Advisory Committee. They came to us and we have the pleasure of talking to them in this atmosphere, like almost one-on-one. We will have 30 minutes for RSSAC to start, then 30 minutes for SSAC, and then 5-10 minutes of presentations, and then Q&A.

Ram, please come and join us. So the first presenter will be the Chair of RSSAC, Jeff, who is the Chair of RSSAC and he will be talking and explaining what this community is doing and how important you are, and also how newcomers can get engaged. Thank you very much. The floor is yours, Jeff.

JEFF OSBORN: Hello. Can you hear me? This is the first of the audience interaction part. That was the easy one. It gets harder. My name is Jeff Osborn. I am the Chair of the RSSAC, which is the organization that ICANN that advises on matters surrounding the DNS root server system. And we had been tasked with trying to explain what it is we do in a simpler way. And then I heard I got to speak to you folks. So we spent a bunch of time

putting together this presentation to try to explain to people who don't do DNS all day long what DNS is and how the root system works. So at the end of this, you will be experts. You'll have a new page on your resume. You'll probably be able to get much better jobs immediately, probably by tomorrow. But you're going to have to pay attention because it's a tad complicated.

So, today we're working from a slide that's three days old. There are a couple of minor things in here I'm going to go through because, like I said, this is a first time through. We're going to learn here how DNS works, what the role of the root server system is in DNS, and we're going to learn about the significance of the root server system. The Root Server Advisory Committee is made up of the people who operate the root servers. There are 12 organizations that operate the root servers. There's a member and then an alternate member, so that's two times 12, which is -- You guys are good at this. You'll get those jobs. 24, and then we have liaisons to and from a couple of other organizations like the ICANN Board, or if you've heard of it, the Internet Engineering Task Force, the Internet Architecture Board, a number of technical organizations out there. So we try to get expertise from wherever we can get it, and here's what we do.

So let me introduce to you the DNS. Now ICANN, one of the first things I noticed about ICANN is how come it's not just spelled I-C-A-N? It would be pronounced the same. There are two Ns in ICANN for a reason. That's because one N is names and one N is numbers. So if you were trying to remember a long string of numbers, you'd have a harder time than if you were trying to remember a word. So DNS uses human names to find computer addresses. The computers themselves have no idea

what amazon.com is. What they understand is a string of letters. So humans know a domain name, amazon.com. I'd like to go there. The computer expects to hear 18.239.62.181, which is literally Amazon's IP address today.

So the DNS translates that. You type it into a browser. It is translated by the DNS system, and for the most part, you can use that same name no matter what's changing in the background. The numbers can go to different places, different countries, different servers, but you're always just using one name. It's a very convenient system, and at its basic root, that's the domain name system. So when people talk about the DNS, it's about using words to get to numbers. Anybody need me to go over that again, or does that make sense? Words to find numbers. Excellent.

SIRANUSH VARDANYAN: Thumbs up.

JEFF OSBORN: Thumbs up. Good. So most connected devices actually use the DNS to get to anything on the internet, and the internet, you may have heard, is large. So it's hard, and this is how we do it. Obviously, there are things like computers and servers that use this to connect, but increasingly cell phones are probably a larger user, and God forbid, it looks like washing machines, refrigerators, doorbells, and sort of everybody else out there. We're all using the DNS to find our way around the internet. The benefits of this system are you just have to remember the name. If your computer breaks and you get a new one, it still can be referred to as Jeff'sComputerHut.com. It doesn't matter that my server changed, I

moved down the street, or my building burned down. So you have service portability.

This is a really important thing in terms of not getting stuck with one vendor. You don't get stuck with one ISP. You don't even get stuck in one country. It's how the people who have a domain name, now here's another inside ICANN thing. You are not a domain name owner if you have a domain name. You're a registrant. So who here is a registrant and didn't know you were a registrant? That's something else to put on your resume. You can put, I am an actual registrant. I guarantee that's another 10% in pay at most places. It's a little-known secret, pro tip. The primary benefit of DNS at the end of the day though is it's a human easy identifier. It's easier to know that I am jeff@isc.org than jeff@118.36.12.92. And I can literally remember there was a day when you had to do that kind of nonsense. So this is a better system.

SIRANUSH VARDANYAN: Jeff, if I can request to slow down for our interpreters, that would be great. Thank you.

JEFF OSBORN: Is there an antidote to coffee? I will try to slow down.

SIRANUSH VARDANYAN: Thank you.

JEFF OSBORN: This is hard. I like talking fast. I will try to slow down. Devices get these addresses from something you call a resolver. You've probably heard the term resolver around here. There are millions of these around the earth. The familiar ones to you might be 8.8.8.8 is a resolver that Google put all over the place. And if you point your machine to it, it'll tell you where things are. There are millions and millions of these things out there. And they basically act like they can read phone books. So the idea is the phone book is a thing called an authoritative server, and the numbers in it are zone data. And when you ask it, hey, where is www.amazon.com, they can look it up in here, and there's the address. This happens in milliseconds. It is incredibly quick, and it happens a staggering number of times.

We estimate in the average day, this happens 500 trillion times. 500 trillion is a lot. If you started now and counted to 500 trillion, you'd never finish. It's a big number. The resolvers themselves get their addresses from a thing called an authoritative server. So we're going to start now with the process of how these things actually connect on the network, who talks to who, and where it all comes from. Resolvers get asked all the time, how do I get to Amazon?

Well, in their memory, they keep the address, because you don't want to have to go look it up every time. So if you ask them simple questions like where is amazon.com, where is tinder.com, where are these things, they can just hand them to you out of their memory. This is called caching. They store it in a thing called cache. This is where most answers come from. When you ask for information on how to find something over 90% of the time, this is where it's coming from, from the memory.

Once in a while, you need to get a new number or confirm an old one and it has to go off to a server and look for it. Warning, here's the technical part. There's a three-level process it has to go through. Depending on how much information you need to know about this address, you're either going to go ask the domain names authoritative server or the domain names authoritative server and the top-level domain server.

The top-level domain is the thing to the .com, the .uk, the .nl, whatever. Or the domain name server, the top-level domain server and a root server, depending on how much your computer doesn't know. So we're going to walk you through the chart. This chart is either the best thing going for learning how this works or it is so confusing you're going to tell me why did you put all those letters up there? I can't see it. So work with me. We're going to try to figure this out.

In the first pass, this is up there. Can you read those letters? Good. In the first pass, the gray box is the resolver and we're asking it a question. We're starting from the outside and we're saying, hey, I'm trying to get to the address for www.amazon.com. Do I know where it is? Click. And the resolver goes, yes, I do. It's in my memory. So here you go. Here's the address. And so we go back out into end. And over on the right, you can see where it says here's the routine. This is over 90 out of 100 times. This is over 90% of the time you just get it from memory. But there are hundreds and hundreds of millions of addresses out there so sometimes you don't know them. So what if I don't know the answer? The first question you ask is do I know where the server is that might know the answer?

And if you do, you send the question, hey, I need to know the IP address for www.example.com. We end up with an authoritative server because example.com has its own place to look for all of the letters. What comes in front of example.com? It could be www. It could be mail. It could be a series of things. We ask it and it comes back and goes, here's the address. We remember it and now we ask, do we know the answer yet? Yes, we do. So we've gone through the two simple cases. Do I know the answer yet? Yes, it's in memory. Or do I know the answer yet? No, but I know where the server is who knows the answer. It tells me, now I keep that in memory and I've told it.

The only thing that makes this complicated is this is hundreds and hundreds of millions of devices so they can't all live on one machine so it's got three whole layers to it. So if you don't have the way to get to example.com, you've got to go a step further down and go, well, heck, how about not www.example.com but example.com? That server is a separate one. It's called a TLD server. If you get the information from that, it'll come back up around and say, we know where this is. Do we know where we are now? No, you just figured out where example.com is. We have to go back and add on the www and now we know both of those. This is a circuit that's going to go three steps down and once you get to the third step, you're done. So if you figure that out and follow it along, this is how DNS works.

If you know absolutely nothing, if your computer just came out of the box, it's got no memory and it doesn't know anything, it knows one thing. Down in the bottom left corner, it knows the addresses of the root servers. That's us. There are 12 addresses for the 12 root server operators and you can find one of them and say, hey, where is the list

for .com? I don't know anything. I just need to know what are the things to the right of the dot? I'm looking for the .com one and then I'm going to ask it where the example is and then I'm going to ask that where the www is. You get this answer, you drop it in your memory, you go back and look for example, you go back and you look for www. You have built the name by finding all of the addresses. You've got it in your storage. That's it. You're now all DNS experts. That goes in your resume. Congratulations. You can drop your tuition at the door. We take cash, tips.

The interesting thing about this is the frequency. If you see the letters in red over at the side, the number of times this happens out of memory is over 90%. The number of times you've got to drop down and go, oh, shoot, I don't even know what example.com is, is about another 5%. Having to go out and say, where is.com and what does it know? It's about 2%. Going all the way to the root, all the way to the root server is about one in 5,000 tries. We're really important if you forget everything, but it's not like every single time you go out and do something, you have to hit a root server. We're your backup. Does that make sense? Should this be a T-shirt? I was thinking about it. It might be too hard to read. I have another T-shirt slide coming up, so we'll figure that out.

Anyway, in review, the root server holds a copy of the root zone data, and the root zone tells you what the TLDs have. You just want to know, tell me everything that ends in .com. Tell me everything that ends in .uk. Tell me everything that ends in .nl, and it'll give you all of those. It'll give you the one word to the left. The TLD authoritative server, the second step down, knows all the addresses for the next step. Once you've got Amazon and .com, it knows www, and it knows mail, and it

knows name server, and it knows returns, and it knows John dot whatever. Anything else on that side. The resolver is the device that goes out there and knows how to stack up these three processes and get you the information you need so you have the address for where you're trying to go.

In the millisecond world of this going on, queries to the root servers are actually rare. We run into this all the time where because the root server is absolutely necessary if you forget everything, people think it's absolutely necessary every step of the way, and it's not true. It's a really important thing to have, but in the grand scheme of things, you don't always operate through the root server. Every time your grandmother looks up how to make sugar cookies, the root server isn't involved. It's involved very infrequently, and I have some people who argue it's one out of 10,000 times. It's at least one out of 5,000. It's just not something that happens a lot. Does that make sense?

This is the slide we pulled this morning because Ram liked it, and we accidentally put it back, so it's kind of funny. I'm going to get in trouble for showing these numbers because they said you can't prove it, but we're pretty sure it's right. It is arguable that the number of lookups that happen, these requests, in a day is about 500 trillion with a T. It's a shocking number. It's an unbelievably high number, but you think about all the times that you have a URL going on in your computer and it's got to go figure something out, and everybody on earth does this. That's about the number. The root servers, we believe, are about 100 billion questions a day, which is still a huge number, but it's only about one out of 5,000 of the total number. So we're talking about very large numbers going on, and that's what we do.

It's important to remember, while the volume is big, the frequency is not. What do we do to make sure this valuable resource always works? A couple of things. One, it is massively redundant. The root servers each contain all of the information. If you go to something like .com, it's got to remember 140 million things to put to the left of .com. The root servers only have to remember the 1,700 TLDs. There are only 1,700 words currently to the right of the dot, from .com to .edu to NL to all the country codes, all the commercial codes. 1,700 of them. In text, that's not a very big piece of information. It turns out every one of these root servers by itself could pretty much handle most of the demand of earth. We have 1,700 of them. When you have two computers that are dual purpose for a backup and one of them dies, the other one still works.

We've got, as of last month, I think it was 1,758 of these all around the globe, and 1,757 of them aren't even really necessary. You could imagine a situation where you lost that many and the whole thing still works. This is a shockingly resilient system. We have diverse hardware platforms. Everybody uses different computers to run the stuff on. We use different operating systems. Everybody uses different ones. We use different DNS applications. In my day job, I'm president of a company called ISC, and besides operating the F-Root server, we developed and maintained a piece of software called By9, which for the last three and a half decades has been one of the main pieces of DNS software.

But lots of other people don't use it. I saw Andre in here. There are makers of other -- nope, he's not here. There are makers of other software that runs all the time. What's great is if every Dell computer worked, the root server still works because we all use different stuff. If

an operating system collapsed, if Windows died tomorrow, we're running a bunch of flavors of Unix. If By9 died tomorrow, we have all kinds of other applications. It is a shockingly diverse and resilient system with no single point of technical failure.

In addition, there's no point of institutional failure, where you worry about, well, what if this company has a problem and it has a court case? What if somebody sues you? What if somebody says you're no good anymore? What if you run out of money? It doesn't matter. We're 12 separate organizations. We spend a lot of time together. I probably see some of these people more often than I see my kids, but my kids are adults, so it's not that bad. We cooperate, but we're all separate. We're separate organizations, so one act harming one organization isn't going to hurt the whole group of us. There's no single point of institutional failure.

Now, in addition, and this one's a little hard for people who are new at this, but there's a widespread rumor that we, the roots earn operators, can control the data of what you see. We can control who you can get to and where you can get to. It couldn't be farther from the truth. If you have a domain, you get to say where it connects to, and you pass that to your registrar, and they pass it up to a group called IANA. IANA puts it in a packet, and that guy right there is involved with an organization called -- he is not looking. He is involved with an organization called the Root Zone Maintainer. They pack it cryptographically, so it's absolutely safe. Pass it off to us, and we put it out to all 1,758 or however many it is right now.

Root servers around the world, and you all have that information. We don't change it. We don't address it. We don't adjust it. We don't edit it. We take the information that comes to us and then serve it to the world. For 40 years we've been doing this. This is operating for 40 years this year without a failure. The system has not gone down systematically in any way in 40 literal years. If you've been around computers, you realize that's kind of an amazing number. It's also sort of interesting if you want to really get down a rat hole, the addressing system we use is called Anycast, and it sort of prevents DDoS from working. You just can't DDoS in Anycast instance in ways that you could take down normal stuff. We're not asking people to try, but a lot of people have tried, and the system just has not gone down.

So in summary, the root server system is really important, if infrequent, part of address resolution that gets used many, many, many times today by pretty much every internet user on earth. Most of the queries come out of memory, and most of the rest of them never get to a root server. The root server operators don't control the content. The system is remarkably strong, resilient, diverse, and it simply works, and that is how the root server system works. Thank you very much.

SIRANUSH VARDANYAN: That's very impressive, Jeff, very. You have one more DNS expert now sitting next to you. Thank you very much.

JEFF OSBORN: Put it on your resume and ask for a raise.

SIRANUSH VARDANYAN: Thank you very much. It's really very interesting. Any questions? Yes, please. Filomeno, you use the mic over there. So we have about seven minutes for questions to Jeff, and then we'll move to SSAC.

ZOE FILOMENO: Hello, my name is Zoe Filomeno, ICANN79 fellow. I come from the interdisciplinary background of narrow technology, brain computer interfaces for patients with degenerative diseases.

JEFF OSBORN: Can you get closer to the mic?

ZOE FILOMENO: Yes. Perfect.

JEFF OSBORN: You sound much better. I'm sorry.

ZOE FILOMENO: So intersectionality of narrow technology, brain computer interfaces, and emerging technologies. My question is in regards to emerging technologies and how a lot of the systems seem to follow the human connection within the system of our own creation. So basically, what I'm stating is that everything that this system is our own nervous system. Just give me one second. So the nervous system, the memory,

how we access is how the root servers access. So in future implications, my concern is with patients that have brain computer interfaces that might be connected to the internet. Would there be a future where you, Jeff, or anyone deals with kind of like the policy of that and how all of that is working? I hope that makes sense. I'm sorry.

JEFF OSBORN: That's the best question I have ever gotten at ICANN. I might not be smart enough to answer it adequately. Holy mackerel.

ZOE FILOMENO: Yes. If anything, IANA already told me that they're not going to be dealing with the issue unless it becomes a present problem.

JEFF OSBORN: That's pretty much what, unfortunately, what I'd have to say. It's really, really interesting trying to look into the future and see what happens. And I've been involved with the internet longer than I care to admit. And I've been wrong a lot of times in guessing what was going to happen. I never saw social media coming. That was just a complete shock. So I'd probably guess wrong on this. What we're doing that is so much fun is we basically, most of the people I work with, have been keeping this internet thing going.

And we're just thrilled with what people do with it. What we do is try to keep it reliable. We try to keep it fast. We try to keep it available to as many people on earth as possible, running as well as possible. And then we sit back and watch people like you do incredible things with it and

go, oh my God, we were part of that. So I will be watching what it is you do, but I don't know how much I'm going to do other than try to make sure it's always there and available for you and fast.

ZOE FILOMENO: Yes. And if I can add one of the main situations when you were talking about DDS attacks, I specialize in the region of how the mathematical algorithms inside music affect neurochemicals. So that can be used to basically infiltrate or biohack a human in order to create certain situations. If this is true of how you guys operate, that would possibly mean that if you guys have no problems with the DDS attacks, that may be true for patients with implants. So I'm kind of...

JEFF OSBORN: The parallels are interesting. I have a son who's trying to attend medical school. I have discussions along this line and this is really hard to do in front of a hundred people, but very interesting question. If you spot me in the hall, I will buy you coffee and we can talk about it. It's very interesting.

ZOE FILOMENO: Thank you.

JEFF OSBORN: Sure. Thanks.

SIRANUSH VARDANYAN: Shanelle, please go ahead.

SHANELLE MCPHERSON: Yes. Good afternoon, everyone. My name is Shanelle McPherson, ICANN79 follow and I'm speaking in my own capacity. First of all, we have all just graduated with honors. Thank you, sir. That's first. My question is, and I'm just curious. My question is for the new gTLDs that are about to come out, I want to know if these are already embedded in the server or is it a case where they are created from the root so if a new gTLD is out, how does the server know that it's already there? Is it created within the root? That's my question.

JEFF OSBORN: That's an excellent question.

SHANELLE MCPHERSON: Yes. Thank you.

JEFF OSBORN: The root zone is basically a list of all of the TLDs and then an address for them, actually one in IPv4, which is the version of addressing that I was using here, and then also an IPv6 generally, which is the longer, newer addressing way. Generally, about twice a day, a new root zone file is generated. So when they come up with the new TLDs, it will literally just be added to the list and away it goes. As soon as that's generated, twice a day, it goes out. They can send them more. Somebody was just telling

me the other day they did five pushes instead of two. It's really fast and it'll all just become part of the root zone. Good question.

UNKNOWNN SPEAKER: [Inaudible - 00:28:22], ICANN follow. My question is, what kind of techniques and systems in the root server system from the protection and from the monitoring point of view, for the health check and for the performance of the all root server systems -- I know that there is a protection in case of the failure, but what are the systems which are utilized for the performance point of view?

JEFF OSBORN: One of the interesting things about the root server and why we try to remind people how infrequently it's used is it is not a part of the latency that persons will notice in their average day. If you think about the last time you looked something up and you entered a URL. Imagine how long it would take you to get to the 5,000th time you did that. Let's suppose a root server closest to you had failed. The next one is probably within several milliseconds away. On the 5,000th time you looked up something, there is an additional 5 milliseconds in your lookup time. What are the odds you're going to notice that? It's literally so infrequently used and it's something that once you use it once just ends up getting put in the cache that we've had a really hard time trying to understand how an end user could ever notice not just a slow root server but a failed one.

Because there are so many of them. When I said the DDoS attacks don't work well, we use a really interesting thing called Anycast. Imagine you

had a restaurant with 1,000 waiters all named Jose. If you put your head up and go, hey, Jose, there is always a waiter there. If he goes to the kitchen, it doesn't matter if you say Jose because there's another one. Because all of the F-Root server instances around the world, there are almost 300 of them from my company, are all with the same address, then when you say, hey, F-Root server, it doesn't matter if this one's dead. There's another one 50 miles beyond and 50 miles beyond and 50 miles beyond.

The way the internet works, it will just keep going until it finds the one that works. Literally, the additional time involved is milliseconds. We do an awful lot of things independently and each organization has different ways of ensuring that everything is up. We're very proud of how we do things, but we all do them differently by design because what you don't want to do is have a monoculture where a single flaw, a single bug or a single problem takes down more than one thing. Does that make sense?

UNKNOWNN SPEAKER: Sure. Thank you.

JEFF OSBORN: Good. Thanks to help.

SIRANUSH VARDANYAN: Guys, we don't have time for more questions for now because we need to go for the SSAC presentation. Do you want to wait until the end of the session if we have time for more questions and if Jeff will stay here?

OLORUNDARE JAMES KUNLE: So after me or…?

SIRANUSH VARDANYAN: No, you're also included, but if you can have 30 seconds to go, then do it.

OLORUNDARE JAMES KUNLE: Oh, absolutely. Good. So my name is Kunle.

SIRANUSH VARDANYAN: I will make sure our next gener will be the first to ask the question.

OLORUNDARE JAMES KUNLE: The question is this. I'm sorry, perhaps I missed it when I went out to Eastman myself, it's about who can join RSSAC and who cannot join, and what are the conditions for joining? Thank you.

JEFF OSBORN: The RSSAC is made up of the two representatives from each of the 12 organizations that serve it, and then there is an RSSAC caucus, which is made up of DNS experts from around the world. To be a member of RSSAC, you have to be operating one of the root server letters. To be a member of the RSSAC caucus, I'm actually also the chair of the RSSAC caucus membership committee. You just have to show a working knowledge of DNS, because we're not doing a training. We're getting a

group of experts together who do most of the work. Most of the work performed by RSSAC is performed by the caucus, and this is a diverse group of people from around the world. There are a bunch of them in here. I'd start looking, but the lights are in my eyes. All you need to do is have a good working knowledge of DNS and apply.

It'll literally come to me and six other guys who look through and determine. If you are not skilled enough, we'll recommend, hey, show us where you think you are skilled. If that doesn't work, we can make you a mailing list only member, and you can serve until you feel like, hey, I've learned enough, and now I know how to do it. It's a very inclusive group. We really look for a lot of help, but it is not a training organization. We kind of need experts when they get there, because this is how most of the work gets done. Does that help?

SIRANUSH VARDANYAN: Thank you, Jeff Osborn, RSSAC chair. Thank you very much. Your applause is to Jeff, but Jeff, please stay with us. And my great pleasure to introduce now the SSAC chair and vice chair, Ram Mohan and Tara Whalen. Ram, it's so nice to see you coming to the fellowship program. When I was a fellow, you were the one who did the presentation for the community, and I really enjoyed it, so thank you for coming. The floor is yours.

RAM MOHAN: Thank you so much. And a good evening for all of you here. It's really a pleasure to be here. Thank you for inviting Tara and I here. So before I get into the presentation, I want to tell you a story. And this is the story

of the SSAC. Since the 1990s, before many of you were born, but since the 1990s, there have been concerns about attacks on critical internet infrastructure, terrorist attacks on critical internet infrastructure. This has been a concern for quite a long time. And people gathered would talk about it, and then they would disperse. They would say, well, this is important. Maybe we should keep track of it, we should continue.

And then in 2001, 911 happened. And that was a concerted attack, not on internet infrastructure, but on other kinds of infrastructure, other things that happened. And it really raised global awareness that had the terrorists, in addition to the things that they did, had they also targeted critical internet infrastructure at the same time, it would have been an even greater magnification of the disruption and of the problems that happened then. So after September of 2001, in November of 2001, ICANN met in Marina del Rey in California, in the U.S. And at that meeting, some of the folks who were at ICANN, some of the folks who were also helping coordinate some of the overall governmental responses to terrorist attacks, came together and pulled a few of us who were directly operating critical infrastructure.

At that time, my company, and I was the CTO for that company, I was operating the .info top-level domain, which was a new domain, but still a significant part of it. There were folks from VeriSign who were operating not only the root servers, but also operating .com and .net and .org at that time. Some other folks who were operating country codes, brought us together and said, can you start to think about what may happen if there is an attack and if your domains get taken over or get taken down. Because until then, while there was some uneasiness, there was this idea, it's going to be okay. We're all good people working

in good faith and everything is going to be fine. Build decent defenses, build decent systems, but you don't necessarily have to worry a whole lot more.

In 2001, inside of ICANN, we created a new group. It was only six of us at the start, but we called ourselves the Security and Stability Advisory Committee. In 2002, it actually became formalized. The ICANN board got together and said, we do think that it's important to have a group that is composed of experts who understand the foundations of how the DNS technology, how the DNS operations understand those foundations, get together, and let's give them a clear focus. So that is what led to the start of SSAC. ICANN itself has a mission, and that mission says it is here in the world to ensure the stable and secure operation of the Internet's unique identifier systems. That stable and secure is what leads to the mission of SSAC itself.

So what is our role? We look at matters, issues that relate to the security and the integrity of the entire Internet's naming and address allocation systems, and we've compiled a group of experts who understand these kinds of issues. We come together, we work together, and then we provide advice. The advice is nominally aimed at the board, but really, it's aimed at the board and the community. Here is the really interesting thing about what we do. We are an advisory committee. We only provide advice. Our advice, the ICANN board gets to listen to it. They can just ignore it. They can throw it away if they so choose. There are no bylaws, there are no requirements for the ICANN board to act on our advice.

This is a really great feature because, let me tell you, what it means is that we, the SSAC, have a huge responsibility that when we provide advice, it is backed up by data. It is backed up by real things that we have observed in the real world. It's not just a bunch of folks who want to come together and say the sky is falling. Our credibility is based on the data-based modeling, the data-based analysis, or in some the experience-based analysis that we put together. It means that we are also not hobbling ICANN and the board to every one of our pieces of advice and our recommendations because in some cases, our advice might be very strong and may not be implementable inside the ICANN system.

In other cases, our advice might not go far enough. It provides some leeway, but the biggest thing for us in the SSAC is that because we're an advisory committee, without the power to force something to happen inside of this community, it raises the bar for us to provide advice that's actually decent, implementable, that is based on data and facts. This slide has quite a lot of words in it, but I will focus on just a few things. What do we actually do? We engage with the technical community, the technical community that is focused on DNS infrastructure, that is focused on security requirements. But when I say the technical community, I also mean policy people. I also mean people who are involved with the law.

So we engage with a bunch of folks who have a keen interest in DNS and DNS security. After we engage with them, we analyze the risks and the threats to the Internet's naming and address allocation systems. We form groups inside of the SSAC, and many times we don't have all the expertise that we need, so we will go and we will reach out and ask and

invite experts to come and work alongside us. That means that folks who are in the SSAC, who work with the SSAC, you are often rubbing shoulders, having deep conversations with people who are actually building the technologies, the protocols, the systems that make the Internet, the DNS as we know today, run.

Once we do that, we coordinate with the various technical and other entities, and we arrive at conclusions on whatever topic that we are looking at. So let me give you an example. We recently looked at what happens with blockchain names that are in the blockchain. What happens with other kinds of names that are evolving in the overall ecosystem? What is the impact that has on the DNS. And we put together a group of people who have interest, expertise. We spent about a year working on it, and we published a report. That report is public, and that has observations in it, recommendations in it, that is available for the community to look at and is also aimed at the board. So those are the things that we do on a regular basis. That's what the SSAC does, and if you look at who is the SSAC, it's really on the leadership team of the SSAC.

There's a few of us who got elected to be in the SSAC as the leaders. I'm the chair. Tara here on my right is the vice chair. Jim Galvin is our liaison to the ICANN board, and Jeff Bedser and Barry Lieba are also two esteemed colleagues who are part of the leadership team. But the way we look at this is the SSAC and the work that we do is not possible to be done without the folks on the other side of the screen, the amazing support staff whom we have. So there's Steve, there's Danielle, Kathy, and Moses, and between all of us, at the bottom of the screen, you see the skill sets and the expertise that we bring. There are another 30 or so

people in the SSAC who bring other kinds of skills, other kinds of knowledge and expertise to bear.

Just briefly, I'm not going to go into any of the details here. These slides will be made available to you, but this slide shows you both the breadth and the depth of the issues that we look at. As you can see, DNS security is a big preoccupation for us. We've got tons of things that we've done on DNS security, but we've also written things about the global namespace, dotless domains, name collisions, shared use of domain names, all of those kinds of things. We just published a report about having a name that should never be given to the public, because it should be used for internal purposes by networks.

So we just put up a report on that. We've been focused on registration data, on internationalized domain names, universal acceptance issues, variants, and routing security. So it's a broad range of interests, but if you look at it, they're pretty much all focused on names and numbers. We do not comment on is spam good or bad. We do not comment on social media, on things that are application-level items. We pretty much focus on things that are at the DNS layer of the internet.

So that is the SSAC. For a certain kind of geek, it is a lot of fun. And for a lot of other people who are on the receiving end of our advice, we work very hard to make our advice digestible, accessible, and explainable in a way that doesn't talk down to people. Just because you're an expert doesn't mean that you have to show off your expertise. In fact, the best way of being an expert is to make that expertise available and accessible to as many people as possible. So that's the

who is SSAC. Now I'll pass it to Tara to speak about, I talked about what is SSAC and what we do. Tara is here to speak about who is SSAC.

TARA WHALEN: Hello. As vice chair, I am also the chair of the membership committee, so I'm very happy to speak to what is potentially a pool of applicants who may become members in the future. We have as one of our goals to sort of diversify our membership, and we are certainly looking toward the future because we need to sustain our membership over time. So that's why I'm also happy to talk to people who are perhaps at the beginning of their relationship with ICANN. So depending on where your professional journey has come up to now, you may already have some expertise to provide. You may be in a position in the near future to be in a position, or maybe you're looking a little further on and you think maybe SSAC and the type of work they do may be of interest to you.

So we want to connect with you to make sure those opportunities happen. So today I want to talk a bit about our particular diversity goals we have set out and the sort of skills and expertise that we're looking for, and also why you might be interested in joining and what's in it for you, and then of course how you can connect with us and find a way to join us.

SIRANUSH VARDANYAN: Tara, just to remind you, 10 minutes notice.

TARA WHALEN: Ten. Thank you very much.

SIRANUSH VARDANYAN: Now we have 10 minutes.

TARA WHALEN: Sure. So here is our diversity goals here. We are really looking for diverse perspectives and unique perspectives. As Ram pointed out, we have quite a breadth of topics that we look at within the huge remit of security and stability. We have challenges all over the global internet, but remember we're talking about global and we have areas of geographic diversity that we feel we're weak in. We are highly biased toward North America and Europe, so we would really like to extend our reach to other areas of the globe, Africa, Latin America, and Asia Pacific.

If we have people here who are perhaps academic or research backgrounds, we may have people who are students who are used to doing data analysis, doing various types of internet measurement. You have expertise of the sort that we may find very useful for some of our SSAC work. We're also looking at our gender balance, so I'm going to be very happy if we could have some more women to join us in SSAC as well. Of course, we're moving to the future as time progresses again. We like people who may be at an earlier stage of their career who are ready and eager to join as well.

This word cloud gives an idea of the breadth of what we might be looking for in new members. Many, many topics here, very technical topics taken from some of our skill survey. You can see that there are

obviously sort of security and networking aspects. You may be someone who has worked in a particular area of networking, a particular form of operations. You may be having experience in cybercrime aspects of security or risk management. Many different aspects of the security stability questions are relevant to us and our work. You may see yourself in some of these areas of expertise. Have a look around, particularly at the documents and the sort of things we do, and see if there's some way that you feel you could contribute.

Again, why might you want to join us? I think Ram also discussed some of the fun bits and what you might want to do. We work on very crucial security matters. If you like to sink your teeth in and work on deep complex matters, we have no shortage of that type of work in SSAC. You will be able to broaden your knowledge into other relevant areas of the internet generally, sort of global internet issues, large scale issues, working with various security experts on these problems. You will certainly probably find this good for career enhancement and deepening your own expertise. There's visibility for a lot of the security forms because this work has such impact globally speaking.

And of course, the expertise is given to improve the way that internet systems work. And so you're advising the board and the DNS community globally to make things more secure and stable and help to contribute to the mission. So if this is the sort of thing that you feel like you might want to do, then we strongly encourage you to apply to be a member. We have materials on the website and an online application which lays all of that material out for you. But I will also point out that we're trying something new this year and that we have an open forum. So we have a drop-in session that is on Thursday, Block 1.

And it's an opportunity for you to interact with us directly, to give us a little bit more time for you to have a one-on-one, to ask more questions, to find out more about the particulars of the projects that we're working on and work we may be planning to and just generally get to know us and find out what we're like as a group of people to be around because we know that's also important. So we do encourage you to go to that session, but we're also available for you to talk to us at any time as well. So please do reach out because we would love to connect with you. Thank you.

SIRANUSH VARDANYAN: Thank you. I thrilled for this session and having you coming, both groups. Because we always, the fellows and next-geners have so much interest in security aspect and we always like we're facing for closed doors. But now I'm happy that your doors are open for us. We have real talents here. And I'm sure after 15 years of the fellowship, we finally have one member in SSAC, which is coming from the fellowship program, and we are proud of that. But we have the huge talent and capacity in this group who have interest to come and join. So thank you for the presentations.

And I also would like to thank Jeff for nominating selection committee member and the mentor for the fellowship program. I cannot explain how important it is to have your expertise to be in the selection and in the mentoring. I would encourage SSAC also follow those steps. I promised next-geners to have the first question and you might have the last one, but I would like to keep my promise. So the question comes from you. Go ahead. Take the mic.

RASHID HASSAN: Thank you for the opportunity. I'm Rashid Hassan. I'm a NextGen fellow. My question is to Mr. Jeff. I have two questions. So the first question is a technical question. In the DNS system, there is a record called CNAME record, and a lot of companies have CNAME record, and they just keep it for record-keeping purposes or some other services. However, after the service is done, they forget to delete the resource, but the DNS record is always there. And due to that, there's that attack called CNAME hijacking or domain hijacking. So hackers do that and tell that I hacked Google. You did not hack Google. You just used the CNAME.

So on the policy level, is there something, like, do you recommend something to companies or to ICANN that if I'm not using a CNAME, how much time should I keep it, and what should I do with it, and how fast I should destroy it? Because it's just a technical thing. And the second thing is if I suppose look forward to joining RSSAC, I am sure there are a lot of white papers and resources and a lot of things. Do you have anything in an infographic version or in a short summary, I can understand the work very easily if I go through it? That's. Thank you.

JEFF OSBORN: Thank you for the question. Let me answer the second one first. If you literally just Google search, join RSSAC Caucus, I literally think you'll end up right there. It's in the ICANN website, which is really it's a really great process. There's an easy application to fill out. It describes what we are looking for, what the work is. As I said, it's a good process where if you do not currently have the level of skills that we feel like is

important that you'd be a valuable working member, you're offered the position of sitting in watching and being on the mailing list and choosing at some point if you feel like, what, I've figured this out enough. Come back around and ask. But rather than a learning function, the actual membership is a group that does real work. We get real things done. And so observing it is interesting.

So I'd literally just, if you Google it, you'll be there in a second. I've tried to have other here's an easy way to find it. Google works. As for how long you stick around with the CNAME, that is way outside the remit of what the root server system does. And I would point you to, you could throw a cat and find a better expert than me. Wes, for instance, here, or Peter, you want to walk up to a mic and answer? Peter is a member of the RSSAC caucus as well as SSAC, I think.

PETER THOMASSEN: Very quick. I'm in one of the younger SSAC members. And so regarding the CNAME, when you don't use it anymore, it still points to the target because the CNAME is an alias record, essentially. So the problem only occurs if, without you noticing, that target hostname changes control and is controlled by somebody else. So usually that doesn't happen in a day or so, but if you forget to check on it, then a problem might arise later. So the recommendation would be, don't have things around that you don't need. Don't have doors in your house that you don't need. And just remove the alias once you don't need it anymore. But it's not like, if you let it sit, then it immediately becomes a problem. It only becomes a problem when the target changes control.

SIRANUSH VARDANYAN: Thank you very much. And the last question will be from your Esau, and we'll sum up.

ESAU TUPOU: Thank you. I'm an ICANN Fellow, and I'm speaking on my own capacity. So from my understanding, the internet started as a defense project. I could be wrong, but sort of like a defense project. And it wasn't considered with the concept of security by design. But then it was made for exchange of information as well as communication. But now it changes everything, like conventional crime. Now we have cybercrime and all the other stuff. So my question is for the SSAC, do you also work with the other parties for the regulation as well as human rights and the other for making sure that the security is safe? Thank you.

RAM MOHAN: Thank you very much. That's a great question. We liaise with governments, with government folks. We have some of our members actually work for government entities. But we also directly work inside of ICANN. We work directly with folks in the Government Advisory Committee, for example. And inside of the Government Advisory Committee, there's a subgroup called the Public Safety Working Group. That usually has people who are not just regulators, but they're also sometimes policemen or people like that, who are focused on making sure that there is stability and law and order on the internet as well.

So we engage quite regularly and strongly with them. We also have many of our members attend. There is, across the world, there are computer emergency response teams, CERTs. And so we have

members who engage with that. Or there's an umbrella organization called FIRST, and they work with that as well.

SIRANUSH VARDANYAN: Thank you very much. And let's give a round of applause to our presenters. Thank you very much for coming. And this is the great end of the second day in ICANN. So thank you very much. With that, the meeting is adjourned. See you tomorrow at the welcome ceremony. Thank you. Thank you, that was great.

[END OF TRANSCRIPTION]