ICANN79 | CF – Joint Session: RSSAC and SSAC
Sunday, March 3, 2024 – 10:30 to 12:00 SJU

RUSS MUNDY: Welcome everyone to the Joint Session of SSAC and RSSAC at ICANN79. This is an open forum this time. And can we have the next slide, please? Okay, thank you. And here's the agenda we have planned for today. And I see we have both our Chair of the SSAC, Ram Mohan, and Chair of the RSSAC, Jeff Osborn, present, and we have others from both groups present.

And I'd like to, if we could go around and introduce, have the SSAC and RSSAC members introduce themselves, especially those in the audience in the room there, if they don't know who people are, they'll put names with faces, and hopefully get even more interaction later on. So Ram, would you mind doing around the table, the introductions, and getting things started?

RAM MOHAN: Thanks, Russ. Appreciate it. I'm Ram Mohan. I'm the Chair of the SSAC and my day job, I work at a registry company called Identity Digital.

JEFF OSBORN: I'm Jeff Osborn. I'm the Chair of RSSAC, and my day job, I'm the President of ISC.

TARA WHALEN: Tara Whalen, SSAC Vice Chair and day jobs at CloudFlare.

JOE ABLEY: Joe Abley, SSAC, also CloudFlare.

LARS-JOHAN LIMAN: Lars-Johan Liman, member of RSSAC, I work at Netnod, root server operator.

ROBERT CAROLINA: Rob Carolina, member of RSSAC and General Counsel with ISC.

MAARTEN AERTSEN: Maarten Aertsen, member of SSAC, and I work at NLnet Labs.

MATT LARSON: Matt Larson RSSAC, ICANN.

DUANE WESSELS: I'm Duane Wessels from Verisign. I'm the Root Zone maintainer liaison to SSAC.

RAM MOHAN: Thanks. Ondrej.

ONDREJ FILIP: Ondrej Filip. I'm late, I apologize. I'm SSAC, and my day job is at CZ NIC. That's it.

RAM MOHAN: Wes.

WES HARDAKER: Wes Hardaker USC/ISI appointment to RSSAC and the RSSAC liaison to the ICANN Board.

WARREN KUMARI: Warren Kamari, day job is Google and SSAC.

HIRO HOTTA: I'm Hiro Hotta, RSSAC, WIDE and JPRS.

KARL REUSS: Karl Reuss, RSSAC, University of Maryland, root server operator.

SUZANNE WOOLF: Suzanne Woolf, dual citizen, and my day job is for PIR, the TLD registry.

KENNETH RENARD: Ken Renard, RSSAC Vice Chair, and day job is Army Research Lab RSO.

PETER THOMASSEN: Peter Thomassen, SSAC member, and I work for SSE and deSEC.

BARRY LEIBA: Barry Leiba, SSAC, and I work for Futurewei.

OZAN SAHIN: Ozan Sahin, SSAC support staff.

BRAD VERD: Brad Verd, Verisign, RSSAC.

JAAB AKKERHUIS: Jaap Akkerhuis, SSAC, and also for NLnet Labs.

MATTHIAS HUDOBNIK: I'm SSAC and I work for Europol.

RAM MOHAN: Any other SSAC, RSSAC members in the audience who are in the room who haven't spoken up? Terry. Come on Terry.

TERRY MANDERSON: Terry Manderson, RSSAC, ICANN.

RAM MOHAN: Okay. And shall we go online? Are there any members online other than Russ? Okay. Back to you, Russ.

RUSS MUNDY: Okay. Thanks very much, Ram, and thanks to all the RSSAC and SSAC members for the in introductions. And I'd like to go on to the next slide if we could please. There'll be an approach that since this is our first joint open meeting in the past, they've been closed, we felt both ACs agreed that it would be good to have this be an open meeting and give the community a chance to hear what we talk about in these meetings. In general, all of the discussions on the regular agenda will be between the RSSAC and SSAC members. At the end of this session, if there's time permitting the chairs from RSSAC and SSAC might invite comments from those either online or in the room. But if there isn't time, we have listed here on this slide several times later on in meetings and open sessions that are set up specifically to include interaction with the audience.

So with that, if we could have the next slide, please. Oh, and I apologize, I don't think I introduced myself. My goodness. I'm Russ Mundy from TIS Labs. Been around ICANN and been the RSSAC liaison from SSAC for quite a number of years. And so I'm remote this time, I'm not there in the building, but I will be online for most all of the meetings if anyone wants to get back or holler at me or do whatever. So we've done our introductions, so if we can have the next slide, please. Okay. Now, let's pass this over to Jeff Osborn to lead and give us a description of what's been going on in the RSSAC realm on publications and work parties. Over to you, Jeff.

JEFF OSBORN: Thanks, Russ. I'm gonna need help without a list of them up there. The work parties, I think Ken perennially is leading a work party. Do you wanna start?

KENNETH RENARD: Okay. I just wanted to give a chance because we do have our work party leader online. Robert Story, if you can hear me. Do you wanna present this?

ROBERT STORY: No, I'll let you do it. Go ahead. Thanks.

KENNETH RENARD: Okay, thanks. All right. So the RSSAC has been studying security incident reporting. This is a way to grow trust in the root service system by being transparent and disclosing security incidents that may be relevant to the public. So that's the main purpose. We are approaching this by defining generally, very broadly the types of incidents that would be reported, any little small blip in numbers, things like that is under the radar that really does not need to be reported. So we're approaching this from the security perspective of availability, confidentiality, and integrity. So we're kind of defining those sections of very broadly, very generally describing the, if you will, threshold that this is the type of thing that would be reported. The one thing that we came up that we thought we wanna at least let you know that we're working on is a statement about confidentiality.

And the document right now states that unauthorized disclosure of query data would be a security incident that would be reported. Makes sense. So here we go. This kind of got off track into a greater discussion of confidentiality, but for the purposes of this report, we're coming out and stating exactly that unauthorized disclosure. So questions to be answered are who authorizes disclosure and anonymization? So the big sticking point here is that we supply in general data to the day in the life, the diddle that our work does.

So we do that once a year, 24, 48, several hours of raw PCAP data. Some are anonymized, some are not, so the whole confidentiality discussion, if there's any off the cuff thoughts right now, we'd appreciate those. As far as the rest of the document, the work parties, what about six months in or so, and planning on about a year, but we're coming to this document and we'll hopefully have the resolution here in another couple months. Questions?

MAARTEN AERTSEN: Yeah. Russ shared the paragraph that you just referenced, and it has a sentence about supporting crypto transports and compromise of keying material, but it was not specific about the key material being like related to the authentication aspect or the confidentiality of the traffic, which I suppose with adoption of protocols that don't actually compromise confidentiality is a relevant difference. I was wondering if this was a conscious choice or that I'm just needling at the level of detail, which is not relevant to the current discussion.

KENNETH RENARD: So it was definitely --

RUSS MUNDY: Excuse me for just a quick insert, if we could say our names before we ask the questions, so it'll get in the record who was asking what. Thank you.

MAARTEN AERTSEN: Yes. So this was Maarten.

KENNETH RENARD: So the intent there was keying material for confidentiality. So that's a good point. I'd like to make sure that that's clear in the document. Thank you.

JOE ABLEY: Joe Abley. I'm curious if you're aware that DNS OARC has a privacy work group that's been set up as a committee of the Board, but it's also got other participants in too, where they're looking at their general privacy provisions and the access agreements by which data such as digital data is stored. But I think it'll be probably worth talking to them. João Damas is running that effort at OARC, so if you haven't already, it may be worth touching base and seeing if there's any saved work that could happen.

KENNETH RENARD: Great, thanks. I forgot about that group. Yeah, all right. Yeah, so that's the security incident reporting. And Robert, if you wanna add anything else. Anybody else? All right, thanks.

RAM MOHAN: Thanks Ken. Again, without a list, I can't recall what other active work group leaders are here. Sure, Rob.

ROBERT CAROLINA: Ken, just a quick question. So incidents get reported. Does it stay in some kind of a long-lived archive?

KENNETH RENARD: Thank you. That was another big part of the document is how do we publish? So this document is meant to inform the eventual governance structure. It is not mandating or describing this is how it needs to work, it's advice to the governance structure because the governance structure of the root server system should have this capability, and that's part of there. So this is the RSSAC caucus, just coming up with that advice. So yes, it would be published by some part of the governance structure is what we think. And that would be a public archive, so. It would be sanitized appropriately obviously, but again, the idea is to be transparent and inform the public and build trust.

RAM MOHAN: Thanks. Ozan, that's it for working groups, right? Okay. So back to you, Russ.

RUSS MUNDY: Well, yes, the other working group that's sort of a working group actually has Jeff Osborn heavily leading the chair and a lot of other people doing work in that. And Jeff is going to give us a preview presentation, if you will, and of the slide set that he plans to use later today. But this is the overall effort that SSAC has heard about previously, but not seen a lot of detail, it's been referred to as RSS messaging. So Jeff, with that, over to you.

RAM MOHAN: And Jeff, just before you start, question for you. Can you tell the SSAC members, is this to inform, is this to inquire? What do you want from the SSAC in your presentation?

JEFF OSBORN: Thank you for that softball lob. Let me see if I can take a swing at it. About a year ago, we were talking about the difficulty of explaining to our parents what we do for a living. And this sort of bled into-- there are a lot of organizations around the world that increasingly are looking at legislating and regulating things involving internet governance. And there are people who are not dumb, they're in fact, highly educated, who do not understand our subject matter well. So the thought was, let's figure out a way to explain to a non-technical audience what it is we do, how the root server system works, so that

they can more intelligently act upon it. So that's the overwhelming thing. So what you're gonna see here is a very simple and relatively short answer to a very big long question that half of us have written master's programs on.

I'm actually giving this presentation this afternoon to the newcomers group as sort of a practice. We gave it to the messaging session an hour ago. It's about an 18-minute presentation with roughly 20 slides. And then this summer, this is going to be a presentation of the GAC. So the idea is this is us trying to explain to a non-technical audience how the root server system works. By definition, we have to start with how DNS works. And then backing this up, this is simply the cartoon version of something Rob and Wes spent a huge amount of time on. How long is that, guys, 6 pages, 7, 6, 11, 12? Okay. A lot of edits. A 12-page document that in great detail explains what this is trying to do at a high level. So I'm really looking for input SSAC, this is such a sibling organization.

When RSSAC and SSAC gets together, I feel like I really have to explain myself. And so, I'm really hopeful that you can provide us some meaningful feedback. All right, thank you. With that being said, I'll act like I'm doing-- my name's Jeff Osborn, I'm the Chair of the RSSAC and this is an introduction to the DNS Root Server System.

Now for the newcomers, I want you to act like you're not technical experts and you're listening for the first time. So you're gonna need to learn some things. You're gonna need to learn how the DNS works, you're gonna need to learn the role that the root service system plays in it, and learn the relative significance of the root server system. So,

RSSAC is simply the community of root server operators that advises the ICANN community on all things involving the root server system.

We're made up of two representatives, one member and one alternate for each of the 12 groups. Plus, we have liaisons to things like the RSSAC Board, the IETF, that's the Internet Engineering Task Force, and other organizations.

Next slide. DNS. So some of you will be brand new to this and some of you will have heard about it, but the name ICANN is kind of funny because it's got two N's in it. It would be pronounced the same if it didn't. It needed two Ns though, because one means names and the other means numbers. So names and numbers are the method by which you find things on the internet. DNS basically uses human names to find computer addresses. I tell people when they ask, what do you do for a living? And I explain something that computers don't understand human speech, they understand numbers.

So humans can remember a domain name like www. amazon. com. The computer needs an IP address like 18.239.62.181. At the simplest level, DNS translates the Amazon into the number. For the most part, numbers change and names don't. So connected devices of pretty much any sort, computers, servers, smartphones, need the DNS to find things on the internet. And you can pretty much say that questions use a domain name and the answers come back in addresses. Fair enough. And this is where I'm gonna start looking for heads nodding, but I think everybody in the room understands this pretty well. Yeah, Ken looks quizzical.

Next page. What are the benefits of DNS? Glad you asked. Service portability is really an important one. If you have a name, you can take it with you to the end of earth. I've had Osborne Capital, Osbros, I've had a whole series of things for just decades and I've moved them around.

They refer to different servers, they refer to different places, they've even referred to different countries because I am the one who has the domain name and it doesn't matter what my number is, DNS will follow me to whatever my new online home is. So, DNS is the way that registrants, that's an ICANN word, meaning the person who has the domain name, get to use their domain name. This is basically part of the world's largest distributed database. It's shockingly complicated and large and it involves hundreds of millions of directories. So again, the benefit of DNS primarily is that it's a human friendly number. You can remember example.com better than you can remember, 192.168.45.99. But the possible section of the people in this room, I must say in paren-- did you have a question?

JOE ABLEY: Just a quick comment if you're inviting them. I presume this is not a tutorial, you're asking for comments on the flyer.

JEFF OSBORN: Sure. I don't wanna consume the whole time though. I kind of wanna run-

JOE ABLEY: I can wait as well. It was a quick one.

JEFF OSBORN: Go ahead.

JOE ABLEY: All right. Well, if it turns out to be annoying, you can just shut that down and nobody else can do them.

JEFF OSBORN: Sure.

JOE ABLEY: So one thing I noticed around DNS abuse is that conversations around DNS abuse can be difficult because to a big subsection of this community, DNS means domain names, it means registries, it doesn't mean DNS.

JEFF OSBORN: I don't think DNS abuse exists. That's a whole nother discussion, that's called fraud.

JOE ABLEY: My point is there's this community who thinks DNS is domain names and registries and there's a minority almost in this community that thinks DNS is anything to do with the actual protocol that we understand by DNS. And I'm just thinking when you talk to

newcomers, it may be worth clarifying exactly what you're talking about because I think people have preconceived notions of what DNS is, and it's not the same meaning that we understand.

JEFF OSBORN: Good point.

RAM MOHAN: Who's taking notes?

JEFF OSBORN: Rob.

PETER THOMASSEN: Hey, this is Peter.

JEFF OSBORN: Thanks. I think let's keep them to the end. If you could take notes, it would be good because this has devolved down so many rat holes in the last year of doing this because we all know this, we all have an opinion, and so if I can keep the questions to the end, I think-

PETER THOMASSEN: It's a procedure question.

JEFF OSBORN: Sure.

PETER THOMASSEN: So Ram put a suggestion in the chat about like some term. I think my suggestion would be that if we have like wording suggestions, we all just put them in the chat and you save the chat later or something.

JEFF OSBORN: That'd be great. That'd be great. Rob, do you mind helping me remember that, grabbing things out of the chat that are comments?

ROBERT CAROLINA: Yeah, I had suggested that copy edits go into the chat and more substantive interventions come here on the microphone.

JEFF OSBORN: That's a really good idea. Thanks. Next slide. So devices get their addresses from resolvers. There are millions of resolvers around the world, and here we're trying to make this simple, we're describing it like a phone book. The new members unfortunately are all under 22 and may not know what a phone book was, but most people in this room remember what a phone book was. The phone book is an authoritative server basically, and the listings are zoned data.

So if we say what is the number for Amazon, we can look it up and we get the number that it's currently at and it is some strength. This happens in milliseconds, and it happens, we believe about 500 trillion times a day. That's an interesting number, don't ask me questions about it until the end. But any input on a better number, I'd love to

hear. And if anybody can swear, they know the exact perfect number on a particular day, I'd really love to hear it. Next slide.

The resolvers get their addresses from the authoritative server. So resolvers remember addresses, this is called caching. This is where answers come from most of the time. Most of the time when you talk to a resolver, it remembers. The odds of you going to a resolver and saying, hey, where's amazon. com, and it hasn't already answered that question is pretty slim. So it's in memory. So most of the time resolvers remember the addresses and give you from a cache.

Once in a while they need to go get a new number or they need to confirm an old one. And depending on how much the server doesn't know, it will either need to go ask the domain names authoritative server or the domain names authoritative server and the TLD, the top-level domains authoritative server. Or if it knows nothing, it's gonna need the domain names authoritative server, the TLDs authoritative server, and the root server. Next slide.

So we're gonna walk through a diagram of how this works. We'll start at the top, and I wish we had a bigger screen, but I think this is legible from here. Basically, the first question that gets asked is, Hey, what is the IP address for www.example.com? And once we get into the box, we're in the resolver. So the resolver asks itself the question, do I know the answer yet? Next slide.

If the answer is yes, then it simply gives the answer back. And over on the right we can see with our estimate of frequency that over 90% of the time answers are returned and all they needed to do was come out

of the cache. So that's the simple and primary case. But let's suppose we don't know the answer yet.

Next slide. If we don't know the answer yet to the question, what is the IP address of www.example.com, then we ask, do I know the IP address of the authoritative server for example.com, not www but example.com?

Next slide. If we do, we then ask it, Hey, authoritative server for example.com, what's www.example.com? The example.com zone data is there, next slide, and has an IP address. So it sends it back to us, we remember it, and now do I know the answer yet? Yes, I do. So off we go to the end box on the top. Did we follow that? This is where I'm looking for nods from the room and hoping that works because if you get that first circuit, then you'll understand the additional circuits. And trying to do this without using the word recursive is really, really hard, but I was told not to. So next slide.

Let's suppose you don't even know the IP address for the authoritative server for example.com. Now you're gonna have to look at the big file, the one for .com. There are 140 ish million entries in here, so it's a big thing. We're gonna send the question, what are the IP addresses for example.com, .com's authoritative server looks in the zone data, finds the address, and sends it back to us up through the loop.

We go to cache memory and we store that, do we know the answer yet? No. All we've learned is where .com is. So we have to go down to look for the IP address of example.com, send the question, get the answer, it goes to memory. Now do I know the answer? Yes, I do. So again, we're at the end. Next slide.

In what's called a cold cache or a brand new out of the box computer where it doesn't know nothing, it will always know one thing. When I start and say what is the IP address, and I don't know the answer to any of where those servers are, I will always know where the root server system is. It's a very short list of addresses and it's just there, out of the box. We will send the question, what are the IP addresses for the .com authoritative server.

The root server will then look in the root zone and give the IP address for the .com authoritative servers. This is a short list compared to these other things. We were just talking about 140 million, there are something like 1700 top level domains that are referred to in the root zone file. This is not a big list. So again, you have an IP address for .com, you store it in memory. Do you know the address? No, I need example.com.

You get the address for example.com. Do I know everything now? Nope. I need it for www.example.com. Do I know it now? Yes, I know it, and it's all in memory. Now looking over to the right, what's interesting is, the frequency with which you get this out of memory is high, it's over 90%. Having to go to an authoritative server, we estimate is about 5% of the time. Going to the TLD server, about two.

And down there with that small number is the rare instance we believe is about one out of 5 or 10,000 times that you look for an answer, you need to go to the root server to get it. And if you understood this slide, congratulations, we have job openings for you. Next slide.

This is just to show-- somebody said we should do a pie chart instead. We're trying to show the idea of the frequency of the events, and they

sort of were an asymptotic curve down, there are fewer and fewer of them as you go down the stack.

Next slide. So in review, the root server holds a copy of the root zone data, and that will get you to the list of all of the things that end in .com or end in .nl or end in .jobs. The TLD, top level domain authoritative server knows the addresses for the next step.

So it knows all the things that end in .com for instance, amazon. com, tiktok.com, gmail.com or that end in .nl, Google, Amsterdam, all those things that end in .nl. And then the domain names authoritative server knows the answer to the questions about the things to the left of, for instance, amazon. com could be mail. amazon. com, www. amazon. com, nameserver. amazon. com returns anything else. So the resolver finds and returns the answer. Next page.

This is the t-shirt, and we didn't know where to put it, so I put it on the slide here. In the millisecond, world of a resolver, queries to the root server system are rare. This is so in parent to the SSAC. This is an important thing because I think a lot of times, it seems like the first thing that happens is you go to the root server, and it just seems like it's a little too much a gateway to systems. We thought it was important to show you in a well-working system. It is rare that you actually do a query to the root server on parent. Next slide.

So three things about the root server operation, it is massively redundant. Last time I checked in January, there are 1,758 instances of the root server around the world. Each server instance holds a hundred percent of the root zone content. In theory, any of those 1700 could pretty much handle Earth. So it is a massively redundant

failover system. They are operated on diverse hardware platforms, diverse operating systems, and diverse DNS applications.

In my day job, I'm president of a company called ISC, that has a product called BIND, which operates some huge number of these DNS servers out there. But because the different operators use different products, if BIND were to all fail, end users wouldn't notice because we have other products from other fine companies, several of them represented here that would just keep chugging right along and you literally would not notice it as an end user.

Same if every Dell platform that we use had some horrible bug and they all went up in smoke, it simply wouldn't matter, same for the operating system, same for the routing we actually route in diverse ways as well. So as a result, it's important there is no single point of technological failure in the system. Next page.

Similarly, there is no single point of institutional failure. The 12 root server operators are independent of each other completely, we all are in very different businesses and organizations. We collaborate continuously with each other. Most of us know each other's children's names and birthdays, we're together a lot, we talk a lot, but we're very independent, and there is no linked connection between them. A force majeure event one could imagine that would take one of these organizations out would have no effect on the end user experience of the internet. It would simply happen and everyone else would pick up the slack. So no single point of institutional failure. Next.

This is something that's important because I think it's misunderstood. The root zone operator do not choose or edit or change or write the

content of root zone data. Root zone data basically comes up from the bottom. You have the right as the registrant to control the zone data for your domain. That gets passed up to the registries, that gets passed up to IANA, IANA puts this document together, the root zone maintainer generates it, cryptographically signs it and makes it available to the root servers who then serve up exactly what IANA wrote. IANA is entirely responsible for the content of the root zone data and the root server operators can have nothing to do with it. Next page.

We are very proud to have been doing this for a long time, and because I'm addressing the newcomers group this afternoon, we have been doing this since before you were all born, which is kind of a funny thought. I was there, you weren't. I'm hoping your parents were all born because at some point, I'm gonna feel very old about this. It's never suffered a service outage, DDoS attacks have been tried, and it's a fascinating side discussion of how the addressing works as any cast, which makes it particularly difficult for a DDoS attack to even get traction. Next slide.

In summary, the roots of our system is an important if infrequent component of address resolution. Most DNS queries are answered out of cache memory. Most of the remaining queries go to the authoritative servers. And we use, like we said, in about one out of five or 10,000 queries. The root server operators don't control the content of the root zone and the root server system itself is massively redundant, technologically diverse, and institutionally resilient. In summary, the root service system simply works. Thank you for your time. Any questions, sir?

RAM MOHAN: Jeff, I think this is a good presentation. There's a bunch of comments that are rolling through the Zoom chat. What struck me was that this began as an educational piece and it seemed to end up as an advocacy piece. And I'm not sure what you're trying to accomplish. Are you trying to say, we run a system that is great, resilient, works a hundred percent of the time, don't mess with it. Is that the underlying message? Or is the message also it's a pretty complex system and it works quite well and we have good processes, and other things in place to keep it stable and resilient. If it's the latter, you don't have enough, I think to convey that only speaking about uptime isn't enough, I think to convey a resilience message.

JEFF OSBORN: I'll reserve my comments to the end. We have other comments. Peter.

PETER THOMASSEN: So, you used the phone book analogy and you also say that it's like resolvers can read all the phone books of the world and then you introduce caching. So that's all correct and makes kind of sense, but if you wanna take that a little further, the way I sometimes explain it is that it's in the old days with phones where you call like information service agent. And let's say I'm in Germany and I'm asking for the number of a French company, then they would say, oh, well I don't know that because that's in France, but I'll tell you the service agent number of the French phone book company. And then next time somebody else asks for a French number, he can say that directly

without looking it up and that's the cache. So I don't know if you wanna make that point, but…

JEFF OSBORN: That's a really good point. And again, to be clear, you are getting this fresh out of the box. This is not the 15th run through, this is like day zero. So to answer Ram's question, the problem is, the answer to your question is yes. We're trying to say this is important, but it's not so important you need to do anything about it, and that's the hard part. The natural inclination, I think is that the root server system looks like something that if I am a sovereign nation state, I need a letter. And we're trying to make the point that it's important, but it's not a gateway. This is not the thing that's stopping you from getting anywhere.

This is not something whose failures are harming you. This is a thing that it just works. This is like the pyramids. It's a thing that's out there, it works just fine, it's not worth messing with, and so to try to both say it's important enough, be careful not to mess with it, and it's not important enough that you wanna mess with it is hard. So you saw through it, that's exactly what we're trying to do, I think. And I might not be putting that well, but we're definitely trying to do both those things. Maarten?

MAARTEN AERTSEN: Yes sir, this is Maarten. I very much like the idea of trying to express how this works. One analogy that may help people, I'm not sure, is that there are a lot of centralized systems where the root of the tree or

the most centralized element is actually the one that does most of the work. And I think you're touching upon that in the diagram, but you don't explicitly contrast it with an example of a system that is that way, and maybe that will be useful to get an analogy that has this property and then saying, so this is not that. So I'll think about an analogy to-- so basically it's a system where that curve is flipped, where most of the work is done centrally instead of decentraly.

JEFF OSBORN: Yeah, Maarten, if you come up with one, please, I'll be around all week. Thanks. Joe.

JOE ABLEY: So I'm just curious. It's very nicely presented, good job. And there are lots and lots of presentations that are like this that have been written and presented by lots of different people over the years. The first one I remember seeing was written by the gentleman to my left I think long time ago. But I'm curious, we have so many presentations about this already. What is there about the existing presentations that means there's a gap that needs to be filled? I understand this has got a little bit more root server centricity about it, but I'm curious as to why did you go about doing this work? What is the gap?

JEFF OSBORN: Rob's got his hand up first, so I'm gonna.

ROBERT CAROLINA: I'd actually like to address that question if I might.

JEFF OSBORN: Please.

ROBERT CAROLINA: Okay. It's a very, very good question, and as many of you have noticed while Jeff has been talking, I've been engaged in a lot of individual discussions with about half a dozen of you in the chat trying to address some of the points that you've made and questions that you've raised, and this has come up in the chat window as well. The challenge that we faced is that every set of teaching tools, every diagram, every pyramid, every description with one exception that we located, now there might be other exception, but there's only one big exception we found all present an understanding of DNS from the perspective of the logical hierarchy of how the namespace is organized.

So a typical description of DNS, and one of the folks in the chat has said, when I teach DNSI, I go completely in the opposite direction, you ask the root server how to find the TLD, you ask the TLD how to find the domain authoritative. We specifically wanted to invert the pyramid because the logic hierarchy of how the domain space is organized is the inverse of operational frequency and the inverse of impact if something goes wrong with the system. And the only time we found that inverse flow being demonstrated, which you could call the end user journey because it's, I ask a resolver, and a resolver then does these things In this order, was an article published by the Internet Society about 22 years ago, 23 years ago.

And it did not try to chart what we just said, it described it, but using pseudo code, which made it a little bit inaccessible to the non-technical audience we were looking for. So understanding both the importance of the root server system, and let's face it, the relative unimportance of the root server system in terms of impact of difficulties, proceeds from that inversion of the hierarchy. So that's why that picture exists and the larger 12-page document that Wes and I have now dedicated person years of our lives to, it feels like, except it was very productive, also proceeds with that kind of description. That's the reason for the difference in presentation.

JEFF OSBORN: Joe, if I can add, Rob obviously put that really well because that's what he does. There is a concern, well, hell, you know ISC really well. Okay, the two things we do open-source software and maintain a root server are both under I would call it threat from legislative and regulatory actions, primarily by our friends in Brussels and in Washington DC. And so we kind of we're in the position to realize Maarten and Robert talked to the EU folks a lot, we talked to the US government a lot. If they get this wrong, they could really, really damage the internet, I think. If they mask with how the root service system works on a bad day, if somebody had a hangover and a bad attitude in the US government, they could do more damage to us than a boatload of hackers.

And so we kind of felt like we're gonna need to go in and explain what this thing is and then explain what your actions would do and explain further. So we had Jamie and Elena in here before from the

government interactions folks, and we're suspecting we're going to be involved with them. The basic impetus was though that you've gotta be able to explain to a legislative aide what this stuff is and why not to mess with it. I don't know if you've seen some of the legislation that is threatening to come down on open-source software.

I literally saw a document that would've required ISC to continue operating to put up over a hundred-billion-dollar bond, that would cost us our annual revenues per year for the premium. So, I asked the Board what we do about it, and they say have a huge party in big bonuses, and then you just go away. So a lot of it was, we're in kind of a protective mode, what do we do in the legislators and regulators, get a hold of this. Does that make sense?

RAM MOHAN: Question, Jeff. A couple times you have frequency, you had on that chart that where frequency was mentioned, and then you are ending in summary, you have frequency again. What's the goal? Why is frequency important to show and mention?

JEFF OSBORN: I think it's because what Rob was saying, where the usual way this is demonstrated, the primacy of the root, I think is artificial because it's where you start. And we're trying to say that the root is important, but it is not the primary thing that gets dinged every time you open your laptop.

RAM MOHAN: So then maybe another way of visualizing that instead of providing kind of bland percentages might be to show actual volume that goes, if that's available.

JEFF OSBORN: We took out the slide that had the numbers an hour ago because everybody said you can't show the slide with the actual numbers. So it's pretty funny that we then get into this group going, it'd be great if it had numbers.

RAM MOHAN: That's why I'm not making the decisions. From where I sit, it's fine to show even the percentages, but if that's a key point, you have this as an important if infrequent. So frequency is an important part of the message you're trying to convey. I'd encourage you to think about how to make that more front and center than something on the side. Even on the chart, it's off on the side. But it seems like it's actually a key point, a key thrust of the messaging that this this is important, it's exercised, but there's a whole bunch of other stuff ahead of this that gets exercised a whole lot more.

JEFF OSBORN: Wes.

WES HARDAKER: Thanks. I guess coming back to a couple of points, I think that the middle that I wanted to land in that we're not there yet is actually to

leave one large number, which is how many DNS resolutions happen in the world in a day, but not how many get to the root, because that's still a really large number, because a percentage times a really large number is still a really large number, and we don't want to emphasize that. My colleague wrote a paper along with some people that attend ICANN that actually measured the impact of TTLs and of what happens when an authoritative server goes down and you have a lot of stuff in your cache and you're able to write out DDoS events. And then the reality is in order to take out some of the top level TLDs and the route and things like that, most users would notice.

The reason why infrastructure has gone down is because they have taken out the site you were trying to get, not what's up above it. And so I think one of our goals is to sort of say, look, the way the root is designed as it is very stable because the number of times you actually need to go talk to it are rare and it is wide, it's diverse, we've brought about three different points of diversity that Jeff mentioned as well. Going back to Joe's question earlier, the point behind this document to me, and I've said this a few times, is to give enough of a background to people that have never seen this, that you can start to have the conversation.

You need to have them have the vocabulary so that when they start asking questions, they can actually ask a question in an intelligent way that then they can represent, and then you can go dive into the details then. Well, it turns out that there's more than just an address that comes back. There's name server records, by the way, they're all secured. We didn't talk about security in here which Robert will come back to. And so it's the point where then you can dive off in the

direction that they care about and have the deeper conversation after they've seen the introduction and can actually hold a conversation.

JEFF OSBORN: Thanks, Wes. Brad.

BRAD VERD: Again, building on everything here and echoing what was said in a previous meeting is, I am concerned about putting a specific number down, and the specific number you put down today was not what the number was a year ago, it's not the number 10 years ago, much less a year from now. So that is a moving target. And in that same line, echoing what was said earlier, the fear is if you put a number down, that becomes fact, and we want to try to avoid that. So I think that's super important to pay attention to.

WES HARDAKER: And humans notoriously don't understand big numbers. They can't wrap their heads around it. I don't. Go ask somebody how big a million is. Like even people in this room can't answer that well, it's conceptually beyond us.

JEFF OSBORN: I think we're done, I appreciate everybody's time and indulgence. Thank you very much. And we'll hand it back to Russ.

RUSS MUNDY: Well, thank you very much, Jeff. I certainly hope it gave you additional feedback that you were looking for, and I thank all our SSAC members for their suggestions and the RSSAC members for their suggestions and responses. As you said, it's still an ongoing effort, more is welcome if you think of things later, and of course, the slides will be on the side. So now we shift over to the main area that the SSAC is focused on at this point in time. And I think that, Ram, did you wanna give any introductory remarks or should we just turn it over to Suzanne?

RAM MOHAN: Introducing Suzanne Woolf.

SUZANNE WOOLF: Yeah. Well, some days-- never mind. Thanks, Ram. There's only a few slides here and I'm gonna go through them reasonably quickly, both in the interest of time and because I think we can reasonably assume SSAC and RSSAC folks are comfortable with an understanding of what name collisions are and why they're bad. And the other thing to keep in mind is that we as used here actually refers to, there's a discussion group, across community discussion group that produced the report and did the work here, but it did include quite a few SSAC members, and I really wanna say thanks to all you folks have put in your time on this endeavor. NCAP Study 2 just proposes a comprehensive name collision risk assessment framework to safeguard the DNS from potential disruptions caused by name collisions.

The thing to keep hold of there is that we worked really hard to shift the thinking about name collisions away from an engineering or operations issue towards a risk management kind of framework kind of way of thinking about it. So the key features of the risk assessment framework, integrated risk assessment that is assessing the risks of name collision for an applied force string becomes part of the general process, which has, as everybody remembers, multiple steps before a decision is made about granting a string to an applicant.

It proposes a technical review team to evaluate proposed new gTLD strings based on empirical analysis. There's a lot of difference between different names and different ways that name collisions can happen or be remediated. So there's a certain amount of technical experience and background required to manage that. Enhanced data collection encourages the collection of additional quantitative and qualitative data from publicly available data sets. There's a lot of data out there, but there also-- because different name collision cases require different kinds of understanding to remediate, well to find and remediate, one of the things we found is that it would be really helpful to have more data and multiple assessment methods. Four different methods for collecting and analyzing data to assess risk. Next, please.

The goals, first ensure that name collisions can be assessed. Some of them are only visible under certain conditions or in certain data, and that doesn't mean they're harmless, it just means they're harder to find. And goal two, provides process for ICANN to evaluate mitigation and remediation plans for identified name collisions. And the thing to keep in mind there is that it does ensure a mitigation or remediation plan can be developed and assessed. We don't want to say to

applicants, oh, there's a problem out here, too bad. There's also, if you think you can do something about it, let's go ahead and consider that. Next, please.

And I am not going to spend a lot of time on the diagram, but part of what's going on here is the name collision assessment was done as after the fact in the new gTLD round in 2012. And by identifying it in advance as part of the evaluation, that has to be done for any application, it changes the parameters a little bit. It gives us more room to explore and to study what's going on, but it also complicates things a little bit. Next, please.

So for this audience, this is one of the important pieces, and we will get back to this. From the initial application, there's a certain amount of review and study that can be done, then assign the application to, well, the string to the TRT, which will be able to determine from initial data collection how far they have to go into additional data collection methodologies. One of the things that's characteristic of this is that a delegation in the root zone, under the custody of IANA, may be required before a decision is made about delegating a name to, well, granting the name and moving the delegation to the applicant, but we'll get back to that.

If there's no name collision concerns, and everything else is in order, go ahead and grant it. Otherwise, you might have to go back and talk about, okay, what can we do about this? What is really the impact and how do we manage that before a final decision can be made about granting the name. Next, please. Sure.

BRAD VERD: I just have a comment and the only comment here is because I had a lot of confusion around this. And so, not to say if I have it, other people have it, but the terminology used up there is a little misleading. Like when you say TRT delegate, it's already delegated, so you're just delegating it like with a different owner, like not IANA or somebody. And I think some of that language needs to be fleshed out a bit more because it's misleading. Because I was not under the impression that all applicated TLDs would be delegated to the root and go through this process until just recently. And I get it, I'm not arguing with it, I'm just saying that the language I think needs to be flushed out a bit.

SUZANNE WOOLF: That is really helpful, that's why the current report is labeled draft. Thank you very much. And next, please. Again, this is worth looking back at later, but again, a fair amount of the complexity has to do with the fact that this is being done as part of the normal flow of the decision process around granting a TLD. Another piece of it is that because of various technical changes, we all know QNAME minimization and local root and all the other things that people are doing to make their data slightly less visible on the internet. There is a longer list of tools and techniques being proposed here.

And, again, this is worth taking a look at, but let's go ahead, let's now. Right. So there are a couple of places that contain explicit references to the root zone and root server system. One recommendation establishes a process for removing a temporary delegation where the assessment process results an unacceptable risk to internet services. There has been talk about such a process or a procedure for quite a

long time now. So there's nothing new about this recommendation and there's nothing new about assuming that RSSAC would need to be consulted when defining this process. Next, please.

Updated assessment methods, develop and implement improved methods for assessing potential name collisions, incorporating data beyond root servers for a more comprehensive evaluation. That doesn't mean we don't need root server data. So again, there has always been some amount of RSO involvement that will continue and maybe expand just as far as providing data under whatever terms people are able to do so to support improved name collision assessments. There's also something that I just realized this morning, probably needs another slide, and I apologize for not making sure that got done.

But in some cases, as Brad suggested, we need to clarify, the delegation of an applied for a string may change one or more times before a decision is made on granting that name. This would result in multiple changes to the delegation. But those would not change the size of the zone by more than a few bites. You're changing an NS record, and would not require more frequent updates to the zone.

The discussion group concluded there's nothing in the recommendation, which I believe is section 5. 5 in the report. The discussion group concluded there's nothing in this recommendation that violates existing on SSAC advice on new gTLDs. SSAC feels differently. There was discussion and decisions made about basically staying within the existing guidance, but if that needs to be revisited.

And next. And that's what we've got, and I think we've run out of time. Okay.

RUSS MUNDY: We have till the top of the hour.

SUZANNE WOOLF: Okay. I feel much better now. My apologies. Questions, comments? Brick, bats.

WES HARDAKER: Sorry. Zoom was buried. Thanks for the presentation. I think we and RSSAC have known this is coming, it's not a surprise the roots server operators, we published a document long ago that describes what our concerns were about the rate of growth, and nothing that is happening in this affects that. I know that I've spoken with NCAP participants that have wondered, are we gonna care about the rate of churn? The fact that it's gonna be inserted, possibly removed, whatever or inserted and changed. No, that simply doesn't affect us. So please do feel, unless the rest of our RSSAC disagrees with me, but I don't think that's gonna be the case. Please do take that back and say the root server operators are ready for as long as you prescribe to the rate of growth that we published in an RSSAC document, we're not concerned at all. So

SUZANNE WOOLF: Thank you. That's helpful to know. Go ahead, Ken.

KENNETH RENARD: This is Ken Renard, ARL. So yeah, echoing Wes' comments, one thing that was mentioned about the sort of emergency revocation, the propagation of a new root zone is pretty darn quick. In the context of TTLs and caching, it's almost irrelevant. My question is really about the measurements that we do and what we publish, so RSSAC002, we publish the, what is it, the 70th or 90th percentile of, what is the latency between publication and 90% of our instances having the new version. Is that sufficient data for you? Do you want better data? Again, putting in the context of TTLs and things like that. Thanks.

SUZANNE WOOLF: I think that's an implementation question that people will be keeping in mind, I'm quite sure.

RAM MOHAN: You have Wes in the queue.

WES HARDAKER: Yeah, the only thing I was gonna add was I guess one of the other comments I heard was if you insert things or if it looks like something is coming, the caching for real answers from the root, or of course two days, the caching for negative answers is quite a bit smaller. And I did hear one comment that was like, well, is there a concern that there's going to be an increase in questions about things that don't exist because they might have existed in the past if you actually remove it.

No, the number of NEX domains we get for everything else will far outweigh anything from this particular approach.

SUZANNE WOOLF: Yeah, I think it's worth pointing out that there was an extensive discussion of, no, we don't want to and we don't need to go against the existing guidance, but one of the things about calling it a risk management problem is it's trade offs all the way down, things are going to change whether we like them to or not, and there's always what are the costs and benefits of acting as opposed to acting differently or not acting at all.

So we did consult with IANA, with Kim, and he pointed out that he wasn't worried about this either largely because one of the things that has changed since we went through this exercise before is that IANA has much more highly developed processes for deciding something needs to change and then being able to determine what change is required and what protections are required and so on. So that's actually one of the things that in my view at least has changed for the better in that time is that there's a lot more room for IANA to study issues and be able to change their procedures.

BRAD VERD: Yeah, that's a perfect segue in my comment because, Ken, I don't think the question about recovating and the emergency revocation was written on the distribution of the root zone. I believe it was more around the IANA change and getting it into the root, which there are

very well-defined procedures and SLAs around an emergency change that are followed.

SUZANNE WOOLF: Anybody else? All right. Am I missing something in the Zoom room?

DUANE WESSELS: Suzanne, I've got one last minute question, sorry. You mentioned QNAME minimization as sort of a challenge in all this. Do you feel like you have good workarounds to that or is it still a concern?

SUZANNE WOOLF: It's part of the sort of general point that yes, we can get root server data, yes, root server data is valuable, but there is probably less information in root server data, which sort of drove some of the reasons for developing these other methodologies for collecting relevant data for name collusions. Well, you can call it a workaround, but basically, it's an issue of the data that you need being sort of spread out among sources and interactions.

WARREN KUMARI: On that, I mean there's actually a bunch of things that have changed since the 2012 round, like QNAME minimization, aggressive NSEC, local root, blah, blah blah. Specifically on the Q QNAME A minimization front, one of the potential things is if you delegate it to a name server to collect some further data, if you have a wild card created in that with a record type that nobody will ever look up like H Info or

something, you get to sort of break the QNAME minimization thing because the resolver will keep coming back and asking for the subsequent queries down the root, or sorry, down the name. So that kind of gets around that, but there is a huge chunk of information which just doesn't show up because of things like aggressive NSEC, but it's somewhat of a concern.

SUZANNE WOOLF: Any further comments from the class?

RUSS MUNDY: Suzanne, this is Russ. I have a question for you, and that is, at what point in this process will it be as the TRT, given that all these recommendations get accepted and move forward? At what point will it be visible whether or not existing data collected by root name server operators is sufficient, or if it's not, what specifically additionally needs to be collected? Will that be part of what the TRT implementation deals with? Or what point in the process do you expect that to give a better definition determination?

SUZANNE WOOLF: Yeah, to some extent I'm sort of, this is a personal opinion, I'm kind of going out on a limb, but what I would expect of the TRT is that if the initial investigation finds evidence that there are name collisions that need to be dealt with, they would be in a position to ask for additional data. And then whether and how they get that probably depends on the situation at hand. There's a bunch of factors that impact, for

instance, if you looked at the ICANN Org comment to the public comments proceeding on study two, there are quite a number of factors that influence what data's available and under what conditions. So I think that's gonna have to evolve. I like the question though, and what I can do is go back into the report when we're working on the final version and see if there's something we need to add to what's already in there.

RUSS MUNDY: Yeah, thanks, Suzanne. Because if it's a general change and what needs to be collected at the roots, then it would be good if it could be incorporated into the RSSAC documents or into the diddle process. But if it's a particular one-on-one, it's hard to know what should the RSOs need to make.

SUZANNE WOOLF: Yeah, and in fact, I think people were fairly disappointed that it was difficult to make more generalizations than we did, there was a pretty hard limit on how much we can generalize about any of this, which, oh, well, it's the internet, but it's still kind of not the ideal.

RUSS MUNDY: Yes, it's the internet world we live in. Okay. Does anyone else there in the room or online have any questions? Ram.

RAM MOHAN: Suzanne, could you inform the RSSAC of the timelines that we're working with and kind of what the next steps are gonna be?

SUZANNE WOOLF: Yeah, I was actually realizing I should probably do that. The public comment period just closed and the discussion group, there's a staff report that's due, I believe, next week. And the study group has to review the comments that came in, and we have several weeks to produce a final draft of the report. The due date for that is early May in order to fit in with the rest of the timeline over the other pieces of the new gTLD program that need to fall into place.

And I don't have it in front of me, but we can send to the mailing list. There is some time for the discussion group to meet here in the next few days. And we will be starting to discuss the comments that came in the public comment proceeding. And that meeting is open, so, if anybody really wanted to, there's a bunch of recordings of the meetings, of the study group and so on and so forth. So we can certainly advise for folks that might want to be available for that.

RAM MOHAN: Yeah, thank you, Suzanne. Jeff and the RSSAC members specifically, if there are terms that we're not using exactly the right way or if there are other specific things that you've already seen and mentioned to us here, when we finish the work, we're gonna end up submitting this as recommendations to the Board for them to go and implement, so accuracy is an important thing. So would really appreciate your folks spending some time looking at what the DG report is doing, coming to

the meetings that are gonna be held here. And if there are any follow up meetings on Zoom, come to those as well, because we want to get this right. We don't want to issue a document and then have to provide clarifying things on it so appreciate your involvement and assistance on this.

SUZANNE WOOLF: Matt, go ahead.

RUSS MUNDY: Matt. Go ahead, Matt.

MATT LARSON: Yeah, I think this is a good follow up to what you just said, Ram. Do you know yet the nature of SSAC's additions to study two? Like, what is SSAC going to publish relative to study two? My understanding is it'll be anything from a cover letter to something more.

RAM MOHAN: Thanks, Matt. We're having conversations on this, but in general, there seems to be quite-- so far, there seems to be quite strong consensus towards what the DG is proposing as recommendations. But I think the SSAC is probably gonna weigh in quite a bit more on the dilemma between privacy versus security, that part of it is probably where the SSAC is gonna make more specific comments. And some of that, we'll end up sharing with the Board as well, but it's still work in progress. But in the work party meetings inside the SSAC, the SSAC work party,

there has not been much divergence from what the DG recommendations have been.

But the work party also hasn't engaged on any of the public comments that have come to the DG because that comes to the DG first before it gets to the work party. So most of the discussions inside of the work party have been focused on the ICANN Org legal letter and looking at the privacy concerns that have been raised there and trying to phrase our responses in a way that allows the Board to actually look at this as a risk management exercise rather than picking a side.

SUZANNE WOOLF: Yeah, and that's actually a good opening for me to point out that actually, we really appreciate all of the substantive comments that were with the public comment proceeding, particularly including the Org because these are some difficult conversations and they should be held in a transparent way where possible. So we do appreciate that. Going, going. Anybody else? All right. I'm not seeing any hands. So back to you, Russ.

RUSS MUNDY: Thanks very much, Suzanne. And we have now reached our any other business part of the agenda and open discussions. So let me ask our two chairs of the advisory committee what the preference is, because we did say that it would be up to the chairs whether or not we took any questions from the audience or the non-members on, on the remote participants. So Ram, Jeff, over to you.

RAM MOHAN: Thanks, Russ. On the AOB, one thing that I wanted to just ask if Tara would take a moment, she's leading up an initiative on membership and adding to that. I think perhaps if you take a couple of moments, and after that, Jeff, I have no problem getting audience questions.

JEFF OSBORN: Me neither.

TARA WHALEN: Sure. Just to let people know that we, as Ram said, are looking for new membership in SSAC. We're always sort of happy to have people contribute to our general efforts towards security and stability, and you probably have community members who have skills and expertise, that can be can be brought into us for SSAC's benefit. So we're particularly looking for sort of areas to improve our diversity geographically. A lot of membership in sort of Europe, North America, so if you have people in other regions, we'd be looking for that. We certainly are looking for people that have research backgrounds, if you do data analysis, that sort of work that is very useful for us as well.

There's a breadth of skills as well, people have other interests that they can bring. There's technical ability, but lots of things around risk management and people have domain expertise that they can also add. So we're very happy to increase our reach on these areas. We've also now started changing up a bit of the ways in which we do the applications. So we're now making more of our materials easier for

people to apply online, to give you more information to know what the process is like, and all of that's a work in progress. So we're very happy to talk to you and to facilitate your entry in joining of the SSAC, if you would be interested in that. And Barry wants to add.

BARRY LEIBA: This is Barry Leiba. Particularly also, as Tara said, we're looking for researchers and people doing measurements and stuff like that, that often will bring in some people who are younger earlier in their careers who are still researching in college, university sorts of things. Please send those people because that's another aspect of diversity that we need is people with beards less white than mine to think differently.

TARA WHALEN: Thank you for that point, Barry.

RAM MOHAN: Suzanne.

RUSS MUNDY: Go ahead.

SUZANNE WOOLF: That's right. I put my hand down. They said more or less what I was gonna say.

RAM MOHAN: Thanks. And Russ, Jeff and I said we're fine taking questions from the audience.

RUSS MUNDY: Okay. Let me just thank you for that pitch for the SSAC membership, but also, I wanted to make a little pitch for the RSSAC caucus membership because RSSAC caucus is where most of the RSSAC work is being done these days, and so there's a process for applying there. And so both groups are anxious to get interested qualifying people to help, and that's the way to do it is raise your hand and say, I'd like to help here. Okay. So Ram, Jeff, back to you. I have Ozan put the chat room structure, and I haven't seen anything more there, but there may be people in the room that have questions.

RAM MOHAN: All right. Nobody wants to ask us any questions, so back to you, Russ.

BARRY LEIBA: Well, let me just ask. Ram.

RAM MOHAN: Yes, Barry.

BARRY LEIBA: Let me just ask if those of you back there found this to be interesting to just sort of give some sort of wave. You glad we had it open, glad you were able to come. Okay, thanks.

JEFF OSBORN: Early lunch.

RUSS MUNDY: Okay. That's what it sounds like. We give folks a few minutes back, and I wanted to thank everybody that came and participated. I think this was a very fruitful exchange between the two ACs. And we will do it again at the next meeting. And thank you for the chairs and thank you to all the members. And I think this closes the meeting. Ozan, do you wanna give the ending?

OZAN SAHIN: Yes, please stop the recording.

JEFF OSBORN: Thanks, Ross. Thanks all.

RAM MOHAN: Thank you.

[END OF TRANSCRIPTION]