ICANN71 | Virtual Policy Forum – SSAC Public Meeting
Thursday, June 17, 2021 – 14:30 to 16:00 CEST

KATHY SCHNITT: Hello, and welcome to the SSAC public meeting. My name is Kathy and I am the remote participation manager for this session. Please note this session is being recorded and follows the ICANN expected standards of behavior.

During this session, questions or comments will only be read aloud if submitted within the Q&A pod. We will read them aloud during the time set by the Chair of this session. If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, you will be given permission to unmute your microphone. Kindly unmute your microphone at that time to speak.

This session also includes automated real-time transcription. By clicking on the Closed Caption button the Zoom toolbar, you can view the real-time transcription. This transcript is not official or authoritative.

With that, I'm happy to turn the floor over to our SSAC Chair, Rod Rasmussen.

ROD RASMUSSEN: Thank you, Kathy. Welcome, everybody. Good morning, good evening, good afternoon or whatever, wherever you are. I hope you all have had a good week here. This is one of the final sessions for the ICAN71 meeting this June 2021. It's been a busy and, for those of on the west

coast of the U.S., rather challenging meeting but I think, overall, a good meeting. This is the public SSAC session.

Could I have the next slide, please, Kathy? The agenda for today is going to include kind of our normal progression of events that we do at these meetings. We have not released new advice since the meeting in March. So we're going to spend some time beyond the normal overview for those of you may be unfamiliar or need a reminder of what the SSAC is and to spend some time with a more in-depth presentation on SAC 115 in which we will go into a lot more of the nuts and bolts of that. Then we'll follow that up with an update on what's going on with the Name Collision Analysis Project or NCAP. Then we'll talk a lot about the things we're working on and things we may be taking a look at and talk about our membership and where we're looking to add to the SSAC and then open it for Q&A. We'll also talk about just a general update on things that are happening with past work and things like that along the way.

Next slide, please. So, digging in, the Security and Stability Advisory Committee is what SSAC stands for. We have 30—is that correct, actually, Kathy? Well, we have 33 members. We have a few new members that we're inviting now which were just announced internally. And those have to be approved by the Board. So we have 33 members currently.

We advise the ICANN community and, in particular, the Board, on security and stability matters that affect the naming and numbering systems and beyond. We have looked at overall threats to the Internet

and also various things that are occurring across it that are emerging threats, etc.

We have a wide range of—oops; not yet [great] yet; thank you—of experience to draw from, as you can see in that lower-left, orange box—that and more. We'll actually go into that a little bit and the skills that we are looking forward to supplement. But we do have those skills areas. Not all of our members have all of the skills. We vary in where we maybe have more expertise, etc., but we are expected to have at least familiarity with all of these things. And then we have published 116 documents since the founding of the SSAC in 2002.

Next slide. Security and stability is one of ICANN's core areas. It's in the mission statement. This committee is positioned to try and provide advice and expertise in that area to keep community and, in particular, the Board informed of areas of concern and to anticipate things that might affect the system going forward as well as provide some advice on current events and issues in front of the community.

The way we work internally is that, on various topics that come up, we'll have a list of topics and we will form a work party based on the numbers interested in that and also the things like public comments, etc. If there's some sort of current demand for information or feedback, we will form that work party. Those people work together, along with their staff, to research that and start writing a report within the work party. We'll review that and debate that back and forth and then present that to the full SSAC for further review. Any updates or changes will be kicked back to the work party to work through. There'll be a final

approval process, and then that passes the full SSAC and that will be published. There is an opportunity for members to provide alternative views as well in that process, which may be published.

In those advisories or reports, we provide—often, not always—specific recommendations to the Board and perhaps recommendations to other parts of the ICANN community or even to the broader Internet infrastructure of law enforcement or whoever may be affected by various security and stability aspects. If it's advice to the Board, there's a formal process for that. That gets sent to the Board, there's a back-and-forth process where the Board acknowledges that, and we establish a common baseline by communicating back and forth on what the understandings are.

And then the Board may take action on that, which could include many different paths. If there's advice that affects ongoing policy or potential policy, they'll get sent over to the GNSO and perhaps the ccNSO if there's something to do with the ccTLD. It would get forwarded along to the correct parties or maybe sent directly to ICANN Org for implementation. Org and the Board, of course, can decline the advice, provided there was some feedback to the SSAC on why that advice was declined. So there is a formal process and some of those advice items can stay open for quite a while, as they might [deemed] implementation, etc. We have several of those. And those are all tracked. Then the final result is also tracked and that advice gets [closed out].

Next slide, please. The last three reports that we published right around … The last two were published just right about the previous ICANN meeting. We talked about them in our last public session. Those are those three that are listed on the top. There's some information on how to contact us.

So that is the background on the SSAC. Next slide. I'm going to hand this over to Jeff Bedser to dig in on SAC 115. Jeff, I'm going to hand that over to you.

JEFF BEDSER: Thanks, Rod. Good morning, good day, to everyone. Thanks for taking the time to attend our public meeting and hear about our latest work. I want to first give a shoutout and thank you to the … We have some invited guests work on this work party with us and this document. We had Kristine Dorrain from the Amazon Registry, Alan Woods from Donuts, and Chris Lewis-Evans from the Nation Crime Agency all contribute a significant amount of effort and text to the document that took a good 18 months to get through a full process but put out some really good effort, I believe, and some good results.

Next slide, please. So let's go about the scope and the purpose of the report.

Next slide. So it distills down to this. Our goal was to look at ways to reduce victimization of Internet users. The strategy was taking an interoperable approach based on universal standards for DNS abuse handling. What this simply means is that the Internet works on an

interoperable approach. The reason an Internet works is because everything is on agreed-upon standards for all the communications to work. For most of the last 20 years, there hasn't been an interoperable approach to dealing with abuse between the different parties and the stack of the ecosystem in terms of service and such. So interoperability was really core to what we were looking to solve. The desired outcome is that SAC 115 act as a catalyst to challenging ongoing efforts in order to begin establishing some universal standards and approaches.

Next slide, please. So defining the problem.

Next slide. DNS abuse in SSAC 115 refers to the use of domain names of the DNS to perpetrate abusive activities. We basically looked at the five primaries: the malware, botnets, phishing, pharming, and spam, with, of course the caveat that it's spam used to deliver another type of malicious content. There are other types of abuse acknowledged. They should be addressed. We took these core ones first as ones that are standard, that are agreed-upon, and are noted. There are new types of abuse commonly being created and, of course, frequency waxes and wanes over time. Again, the work here does not preclude them or ignore them. It just acknowledges that they do exist. No one individual list of abuse types will ever be comprehensive, and we support the concept of a regular community-driven review of both definitions and classifications of types.

Next slide, please. So what's currently being done about DNS abuse? Well, as a session this morning on RBLs talks about, there's blocking and filtering going on, where domains are being blocked based on

certain types of content. Yeah, it's quick to implement. It's difficult to maintain its scale, though. There's a high number of false positives, depending on how you look at it, based on timing. And blacklists can go stale. And there's a possibility of collateral damage that's been well-noted in the community when action is taken at the domain level for something that's a small component of the domain's content. There are notifications and takedowns. Many take a long time based on the processes and the interoperability between two parties in the ecosystem and in the stack, which gives it inconsistent outcomes and, of course, still the possibility of collateral damage.

Leading efforts. The APWG is doing quite a bit. That's the Anti-Phishing Working Group. The Messaging, Malware, and Mobile Anti-Abuse Working Group, which is M3AAWG, also is another one. And then the list goes down there, including things like the Internet Jurisdiction and Policy Network and the PIR DNS Abuse Institute—all initiatives that are working toward defining and tackling abuse. And of course, there are notifier programs that are used to expedite DNS abuse for mitigation, explicit networks of trust. Scaling is difficult by its nature, of course, when it's network of trust. And each program sets its own definitions and standards.

Next slide, please. So framework for interoperable approach.

Next slide. It boils down to this. We believe there's a primary point of responsibility for abuse resolution, wherein if that party is the party that should resolve that type of abuse, the report should go to that party to void extra time for the victimization. There should be known and

standardized escalation paths, where if a party does not respond or refuses to respond, the abuse reporter knows where to go next to get the resolution of the abuse removed.

We believe it needs to be evidentiary terminology and standards, that the evidence of, for example, a phish should be the same, no matter who is looking at it. This is what evidence is required to demonstrate that it's phish so that it can be acted upon with the terms of service of different providers.

There should be a timeframe for reasonable action, where the domain, again, should not continue to perpetrate abuse on longer periods of time because of different policies that slow down the reaction to the domain.

And of course, to have all this happen and to make this all happen, you have to have good availability of contact information so that the parties can be reached to resolve the abuse.

Next slide, please. So, in the primary point of responsibility, you of course got the manifestation of abuse and the primary party to resolve that, and you've got secondary escalation parties to take it to, as I discussed in previous slide. Each incident of DNS abuse should have a reporting entry point in the ecosystem where the abuse is resolved by policy and process.

Next slide, please. Escalation paths. Again, pretty self-explanatory. But when a reporter either reports to the wrong party or does not get a response from the proper party, there's a path of meditation where the

reporter can take the report of abuse and still get it resolved in that way. So the evidence of the abuse and the time of the report can be conveyed as part of the escalation. Standardized paths will eventually allow building of automation towards these processes [where] we don't include proposed escalation paths beyond Appendix B in the report, but there are opportunities to expand on that. And escalation paths and standardized documentation should be determined by the stakeholders.

Next slide, please. Evidentiary terminology and standards is important, we believe, because if you're requiring an infrastructure provider, whether it be a hosting company all the way up the registries, to take action on an abusive domain, you have to be able to demonstrate evidentiarily that the domain has been used for abuse or has been used for abuse.

So we have four categories here. We've got the temporal relevance: when did it happen, how long after the registration did the abuse occur, and how long after the abuse was detected? Did the evidence get logged or captured. Visual. Was there an A-record? Was there a DNS record logged for the domain? Was there content hosted on a domain that was not a parked-page record, and was it captured via screenshot or some other means? Behavioral. Are there logs or activities regarding the domain itself? Are there records in the zone, changes in delegations, WHOIS records, passive DNS, something else that just demonstrates the behavior of the domain? And of course demonstrative. What was the abuse for? How did it violate the terms of service that would support raid action against it? What is the impact of the abuse? Of course, what

are the anti-abuse policies of the responsible party that can be enacted?

Next slide, please. So reasonable timeframes for action. There's been some criticism of this. It's reasonably so because 96 hours is a very big window for abuse to continue to exist. This comes about because of the research we did with various abuse entities that handle abusive registries and registrars and hosting companies that for the most part do a standard of notifying a party they have 24 hours before they'll take action. So, if it's a registrar, they'll notify the registrant they have 24 hours to take action. Many times the registry will notify a registrar and give them 24 hours to act. This of course does potentially give a full window of 96 hours. That should be the maximum time. We believe there's opportunities to minimize this significantly in a process. Where it's well-evidenced and well-reported the right party, it should be less than 24 hours. I think we all should shoot for a standard that victimization time of abuse of domain should be as small as possible.

Next slide, please. So availability and quality of contact information. There are policies at the gTLD level of course where it's required, but this is a full stack ecosystem that's not all controlled under the ICANN remit. So we're looking for an accurate [inaudible] accessible [content] information for entities in the DNS ecosystem toward mitigating abuse. Readily accessible [content] information becomes increasingly difficult to find the further downstream from a registry you go, which many times does allow for reporting to the wrong party to get abuse resolved. Uncertain incentives of reporting parties. They'll use a scattergun approach. They'll tell everybody at the same time, which makes a lot of

noise in the system. It also means that, by the time one party may look to resolve the abuse, another party has acted upon it, and thus it's a timewaster across the parties. A possible solution of course is to create a single point of context termination, where a reporter can identify the type of abuse and get directed to the appropriate parties.

Next slide, please. So our findings.

Next slide. Basically, a lack of coordination leads to inconsistent approaches to DNS abuse management. We think the opportunity is for a common abuse response facilitator.

Next slide. So the common abuse response facilitator sits in the middle of the ecosystem, where there's ICANN contracted parties, there's ccTLD operators, hosting providers, Internet service providers, CDN or Contacted Delivery Network types, where they can be at the center of the reports and facilitate them going to the right party. You can reverse the arrows in that model.

Next slide, please. So the mission of a common abuse response facilitator would be to scope the problem space, convene some relevant stakeholders, implement best practices models, create the evidentiary standards we spoke about. They don't necessarily have to be created. There's some really good ones out there. They could be adopted. Execute a common abuse handling framework, develop abuse reporting approach that includes elements of SAC 115, establish standardized methodologies to build trust and abuse reports, and to report regularly on the effectiveness of the facilitators programs.

Next slide, please. So our recommendation. The SSAC recommends that the ICANN community continue to work together with the extended DNS infrastructure community in an effort to examine and refine the proposal for a common abuse response facilitator to be created to streamline abuse reporting and minimize abuse victimization and to define the role and scope of work for a common abuse response facilitator using SAC 115 as an input.

Where we're looking here is that the ICANN community is a broad group of stakeholders that can take this on, but also there's a lot of parts of the infrastructure in the stack that are not commonly active within the ICANN community. And we need to outreach. We need to work with those parties as well. Thus, we're not asking ICANN to create this or ICANN to run it. We are asking ICANN to use its role as a community-driven facilitator to get something like this moving.

Next slide, please. And the next slide goes to Jim Galvin, so I don't know if you go on with the deck or have questions first. Your call.

ROD RASMUSSEN: Thank you, Jeff. I just want to add just a bit of color to this as well. We've had several conversations with other groups throughout the ICANN community about this: the Technical Committee, I believe the Registry Stakeholder Group. It's on the docket, yeah. ALAC. Oh, Public Safety Working Group. That's the other one. We've had that. And I believe—yeah, we're on the docket for either CPH or the Registry Stakeholder Group. I don't remember which off the top of my head at this hour in the

morning. So we are reaching out and having discussions with various folks throughout the community on this.

I'll also note—this has come out in several different venues as well—this is one of the documents where we have had some alternative views in the sense raised by some of our members—fairly lengthy bit on that. And it's in the document. For those of you who haven't read it, please take those comments in as well. It's mostly focused on a couple various … One is whether or not the overall effort is worth it and can we actually [inaudible] and actually get something [to go], which is a legitimate question. It's a lot of trying to get a whole bunch of different silos and folks to work together. So it's ambitious in that regard, and we believe that will be part of that conversation: to make something like that happen. There's also some concerns around the way the presentation around response times and how that was derived and put down. The goal of course is everybody agrees to minimize how long it takes to deal with abuse. I think we all agree on that. But those are brought up in those [inaudible]. So it's important to understand that those are in there as well. Please take those in because they're part of the conversation we're trying to have: to have those discussions around the challenges, etc.

Yeah, actually, since we've got this here, if there are any questions, etc., that we want to directly, we'll go ahead and do that because we will have a couple of big topic areas today. So might as well cover those as we go along. So if there are any question … Or do we have anything in the pod, Kathy?

JEFF BEDSER: There was one, Rod, but I answered it. A second one just popped up. The question was probably for you, Rod. It's asking, can you explain the status of alternate views? Does the main report recommendation stand? And you referred to the alternative views as background. Or is the Board somehow supposed to rationalize it, too?

ROD RASMUSSEN: That's a great question. We're actually in that internal discussion as well because this is one of those things where we're trying to provide a way for people … There are a wide range of experts that will have different opinions on things, and then we try to reach full consensus on in the SSAC on topics. If there are other inputs, then those alternative views are also presented.

As far as the official Board recommendations, the SSAC recommendations are what the Board actually has to deal with. The other advice can be taken in as (or the alternate views, etc.) can be taken as helping inform the Board or whoever. After we make a recommendation, do we [inaudible] the GNSO or to some other and have some alternate views presented alongside that? Those are just to be taken as inputs to consider as well. The main advice though that the Board should be responding to is actually the recommendation that's contained within the document itself. Hopefully that clears that up.

I see another one has popped up. I'm not familiar with the CARF. Is anybody in the SSAC familiar with the CARF?

JEFF BEDSER: I think that's Common Abuse Response Facilitator. It's the acronym [inaudible] came up with.

ROD RASMUSSEN: Oh, okay. We got a new acronym.

JEFF BEDSER: Seems so.

ROD RASMUSSEN: I guess so. Maybe we should make a facilitator with a Ph. Then it could be the CARP. It's a phishing joke. Sorry. Jeff, Dad jokes are your [province].

JEFF BEDSER: Still, well done, Rod.

ROD RASMUSSEN: Yeah. For six in the morning, I'll take that "well done."

The GDPR of WHOIS data. Well, I think that's kind of a different topic area, but certainly the obfuscation of WHOIS data or RDS data has had an impact and has been discussed elsewhere on some aspects. If you think about being able to do outreach, which this document covers, and how to be able to notify domain holders of issues that may be affecting their domains, whether it's something like a hijacking or it's been

compromised or is being used for spam or something like that, that does make things more difficult as has been discussed in other areas. So having access to that information for instant response would be really helpful, for example.

Yeah, depending on what kind of entity you end up with, there could be a lot of work with data or it may be something as light as best practices. That's the going-forward conversation to have: what kind of entities might exist.

And Donna asked a question. "I was [directed] to the community"—yes. We've had the conversations. There was not specific recommendation to the Board directly. We went back and forth on this. This is a bigger-than-ICANN-type initiative in that some of the parties that need to be dealing with these issues and having interoperable handling of abuse, etc., are at the table at ICANN. But many are not. So the hope here is that we're raising awareness within the community and a desire to do outreach to other groups that are looking at this through their own lens. Some of the groups were mentioned above and were in the presentation and conversation of the Board Technical Committee about this a week or two ago. There is potentially a role for ICANN to be a mustering point for that.

Right now, the point of the discussion is around getting ideas on how to approach that and move it forward and who and where we might be able to convene that table of people to bring together. So we all have a role in that. It's more a matter of being able to plan our next steps and have enough interest in people that have the want to move this

forward. We'll see where that takes us. If we can get to a point where we can come up with a concrete proposal for coming together—some sort of meeting/summit/what-have-you on that …

And then we have another one. Common abuse response includes [inaudible] by design various actors and [inaudible] abuse mitigation process would be streamlined. So I think that's more of a comment, but hopefully the idea is to put together the people who have the ability to do things, whether it's report things or respond to things, to all have the common vernacular and the common set of baseline principles and processes they all use to streamline things. That is the idea.

Okay. And Jeff, did I miss anything that you want to chime in on?

JEFF BEDSER: Nope. It looks like the questions are either answered in the pod by typing or answered by you verbally. So I think we're good.

ROD RASMUSSEN: Okay. Well, thank you very much, Jeff. I'm going to turn it over now to Jim Galvin to talk about the NCAP.

JIM GALVIN: Thanks, Rod, and thanks to my Co-Chairs, Patrik Faltstrom and Matthew Thomas. This is the Name Collision Analysis Project. As a reminder, folks will remember that this project has been around for a couple of years. It resulted from a couple of Board resolutions back three years ago. ICANN to the Board has been supporting this project

and all of the analysis work that's been going on. We'll see more about that in a moment.

Next slide, please. So, as we engaged in Study 2 this past year and got our funding to move forward, as part of getting organized, we have broken our work up into five major categories of tasks: root cause analysis, data collection, Board questions, a case study that we're doing, a case study designed to speak directly to one of the Board's questions, and then of course bringing all of those elements together to produce the real final work product that we're looking for in all of this. So I'm going to take some time here to talk about each of one of these, so we'll move to root cause analysis. Next slide, please.

Along the way over the last nine years since the last round of new gTLDs, there have been about 40+ reports that ICANN has received. Folks may recall that ICANN had set up a website and a place for people to go if you encounter the name collision problem and you could say something about that and report it. There was a mechanism that was put in place as a response to the controlled interruption that was part of new gTLDs at the time.

We have been supported now by a technical investigator. [Casey Deckio] is going to be doing this for us, who ICANN has hired. He's going to be reaching out from ICANN to those 40+ reporters. And we're looking to do some root cause analysis on all of those things that happened. We're trying to understand what we can learn about name collisions that have occurred and what remediation was done at the time—so how it was detected and what they ultimately did about it downstream

and protected themselves with. So we're hoping that this will feed—well, we're expecting rather than hoping—our final work product and also contribute to a Study 3, which is a study of mitigation methods of name collision activities.

Next slide, please. The next big work product is some additional data collection. For those who've been following the work this year so far, Matt Thomas, one of our Co-Chairs, and also with Verisign, who's an A and J root service operator, has been doing quite a yeoman's job of analysis on that root service data and looking at the presence of the non-existent domain [entries] that root server operators do get.

And what we're looking to do here is to reach out to both of the root server operators so that we can do a comparison of whether or not the data that we got from Verisign is actually representative of root servers in general. That's an important characteristic to understand in all of this space. One of the things that has changed in the infrastructure today on the Internet as compared to the infrastructure of 2012 and the analysis that was done prior to that round of new gTLDs is the presence of global resolvers. So we're also hoping to get some data collection from global resolvers, a similar kind of analysis we're getting from the root server operators, so we can look at all of that and we can see what name collisions really look like, what's present, what the behavior of those non-existent domain name queries look like. Again, we're hoping that this work product … It will be feeding into our overall analysis of name collisions, so we'll need it for our Task 5.

Next slide. The next thing that we're looking at is … One of the Board resolutions had a list of nine bullet points. They weren't really framed as questions, per se. We call them "Answering the Board's Questions." But there were nine bullet-pointed areas in which the Board was asking for this NCAP to speak to some specific concerns that it had about name collisions. So we are working through the process of answering each of those bullet points. We're reframing them as questions to ourselves. And we'll be collecting that input. I'm not exactly sure what that's going to look like in the final work product, but we do want to make sure that we speak to all of the concerns and questions that the Board had when they gave this project to SSAC.

Next slide, please. And the last one of the four prep works before we get into the real analysis is … The second Board resolution was explicitly to ask NCAP to look at .corp .home, and .mail. Folks who were tracking back in 2012 … There were applications for .corp, .home, and there was a realization that .corp, .home, and .mail kind of fall into a special category of some sort that we really want to try to support with data and provide some well-defined characteristics so that we can identify these kinds of things in the future. Those were three names in particular that had just a huge amount of activity in the root server operators even prior to there being any motion of a TLD at that level.

We added to our case study analysis and we've certainly gotten a lot of data from the A &J analysis that Matt Thomas has done. We added land, local, and internal to that case study report that we're producing here. The reason for that is that the six of those names were the six names, at least in the A & J root, which were getting more than 100 million queries

a day. So that's why those six were chosen and where they came from. That was just where we drew the line in looking at the data that they had. So we have the ICANN research fellow, John Kristoff, and Steve Sheng that are working together to create that first draft of that, which will then become input to the NCAP group as a whole so that we can consider that data in the full analysis that we're doing.

Next slide. This is Task 5. It really is the final work product. This is really the essential deliverable that we're looking for from NCAP. We're going to be taking the other four elements together. What we really are trying to get as a principal objective here is, what process would we propose to the Board—what decision tree would be offer to the Board to consider—as it looks at future applications and the presence of name collisions? Our starting assumption here is that name collisions do exist and they will always exist. They're never going to go away. So now what do you do with them and how do you manage them and how do you respond to them? So we're hoping that, as part of all our analysis, we'll be able to come to some guidance, some conclusions, about how those name collisions can be examined.

And then obvious the next step would be to look at the mitigation of those things but that'll really be a Study 3 effort. And we're hoping that we will be able to begin looking at mitigation efforts as part of this full-on name collision analysis. As we take together all the data that we have here, what can we say about name collisions today and going forward and how they might be evaluated by the Board and looking at future applications?

I'll call out one last key thing here: the data sensitivity analysis. One of the things we're thinking a great deal about is we've realized, as we've gotten into Study 2, that one of the things that's very different today than was true in 2012 when they're doing this analysis is there's different data available today. In fact, there's actually less data available today than there was then. This is in large part because of the change in the Internet infrastructure. It's important to keep that in mind.

So one of the things that makes evaluating name collisions hard is actually being able to look at them. We are fortunate in this project to have some friends in root server operators and in global resolver operators who are going to provide some analysis data to us. But one of the things that we have to think about going forward is, what do we do in the future—what does the Board do in the future—given that it will not likely have access to that data in the future when it's trying to evaluate name collisions.

So that's a really important part of the analysis that we have do as we look through all of this, not just understanding what we can about what we have, but what are we going to do in the future? So that'll be one of the more challenging discussions we'll have as the project goes forward.

We are currently targeting to be done early next year, sometime during the second quarter, so that we can be having reports out for public comment because the timeline that we had set out for ourselves when we got funding for this particular study was to be able to deliver the

final work product to the Board in June of 2022. [inaudible] this is '21. And so far, we're on target for that. And, again, we haven't hit any bumps in the road yet because we haven't gotten any real data to look at. So we will see as we go.

And that's it. Thanks.

I guess next slide –oh, actually, I actually forget there was the next steps thing. I kind of already talked about that, so we're done here. Thanks.

KATHY SCHNITT: There's a question in the Q&A pod. It says, "Jim, of the 40+ reports, when was the last received?"

KIM GALVIN: I don't know specifically, Donna. My recollection is that we were told it was within the last couple of years, the last few years. Most of the reports came in shortly after 2012 and then there were a couple that trickled in over the next six years. So it hasn't been recent. That's for sure. Thanks.

KATHY SCHNITT: And we have from Susan Payne, "This is June 2022 for Study 2, not for all the work on Study 3 as well?"

JIM GALVIN: That's correct. We will be examining the role of Study 3 as we get to the end of Study 2, but we have no commitment yet on Study 3 work.

KATHY SCHNITT: Thank you, Jim. And David Conrad has his hand raised.

JIM GALVIN: Oh, he might be able to speak to the question of when the last report came in. If we could let him speak, that would be helpful.

DAVID CONRAD: Actually, we received a name collision report relatively recently—I think two or three months ago. I don't recall. It was anonymous, but it was submitted via the standard mechanisms by which the name collision reports were supposed to be submitted. Prior to that, I think it was on the order of a year or so ago. So they're still trickling in. You'll still occasionally see hints of name collision reports being made. But it has gone down quite a bit over time.

KATHY SCHNITT: Jim, there's another question in the Q&A pod from Mason Cole. "The RrSG referred this morning to SAC 023 regarding WHOIS as a source for spam origination. Any plans to update 023 or otherwise address WHOIS as a source of spam?"

ROD RASMUSSEN: Yeah, that's more a question of—

JIM GALVIN: Yeah, I would say that's not for this project. That's a question for Rod in the current work party section.

ROD RASMUSSEN: Right. Thanks. And thanks, Mason. We're going to have an open Q&A for unrelated questions. But I'll take that on right now. There are no plans at the moment. But that's an interesting topic area. I know we've certainly had some discussions of late in other quarters I've been present in. So we'll put that down as an input for the community, as a potential idea. [inaudible]. 023—that's a while ago. So things have changed a little bit since then. So thanks for the input. Let's make sure we capture that, Kathy.

Yeah, no problems, Mason. It's okay. You've got your piece in while you can, and maybe you have to go somewhere else, too.

So we've got those questions [handled], so let's move on to the next section. Current work parties. Just a heads up to work party chairs on make the call. I know at least that this is not a friendly hour for. So I'm going to run through these. We've got slides on the ones that we haven't already discussed. I will ping work party chairs to go through and comment as well.

NCAP—Jim just covered. DNS abuse—the SAC 115 we covered. As I mentioned, we're doing awareness building and outreach and we're looking to affording a strategy for moving that forward once we have a chance to talk to people and gauge interest and the like. So that's where we're at with that.

We have work parties on all of these topic areas: routing security, root server warning system, the ongoing EPDP, work around domain registration data. Registration transfer [inaudible]—we don't have a work party but we do have a fine interest in that and we talked about that and involvement there. The two-plus-year-old-now, I guess, effort on [inaudible]. We are reviewing community feedback on SAC 114. That was our comments and recommendations on subsequent procedures and related activities. Gotten a lot of feedback on that. We're actually actively reviewing that and engaging with parties on the input we got on that. We may have some updates on that in the near future based on that input, etc.

And of course, ongoing things we do. I mentioned earlier we have a formal process for tracking our advice [as] monthly updates on that going back and forth with ICANN staff on making sure those things are moving forward. We're all understanding where things stand. And then we have the DNSSEC and security workshops at all the ICANN meetings. There's an abbreviated session this time. We're looking forward to a longer session at the meeting in October. And then we're going to talk more about membership.

Could we go to the next slide, please. So routing security. Actually, Russ, do you want to chime in on this one? I don't want to steal your thunder if you want to run through that.

RUSS MUNDY: Sure. Thanks, Rod. I'm doing the report today. The Co-Chair of the work party, Tim April, is also on, as well as a number of the participants in the work party.

So, as Rod says, this is ongoing work. The slide lays out our objectives. This work party has been going for a number of months. The focus of the work party is to really touch on the four bullets that are listed for the initial publication. We're trying to have it be as relevant and informative as possible for [inaudible] incidents related to DNS and DNS operations.

We have a number of examples that we've looked at already and we're trying to, if you will, get our collective brains around what's the best way to identify, describe, and explain both the problem set as it affects things such as DNS and other applications and how we can best inform and present this material. We've talked about within our work party a number of different ways to do it, including the possibility of using other media besides [inaudible],a document like SSAC usually publishes.

So no decisions on any of these things yet, but we continue to work and try to develop this material so it will be useful for the community.

Could we have the next slide, please. Oh, I thought we had our second slide in there. Okay. Well, just then let me add verbally that one of the things that we would like to get from the public meeting is any thoughts from any participants in the public meeting about what they actually would like to see come out of a work party of this nature, trying to provide useful information about routing security and the impact that routing problems have on applications.

So, if you have any ideas or thoughts about what you would like to see, we would very much like to hear them. The e-mail you can send to is listed on the slide. So that will be in the archive and available. If there's anybody that has Q&A, if we have time, we'll do them now. But I suspect we're quite short on time, so we'll just go on to the next party, unless somebody has something urgent.

ROD RASMUSSEN: We're actually doing just fine on time, Russ.

RUSS MUNDY: Oh, good.

ROD RASMUSSEN: Yeah, we got another 35 minutes in our session. If there are questions—I don't see any right now, but please add questions or thoughts either to the chat or the Q&A pod. We would love to get your input there.

All right. Thanks, Russ. Let's move on to the next slide then. Root servers early warning system. Unfortunately, for Geoff Huston, it's not a good hour for him in Australia right now, so I will take his stead on this one.

This one we're very close to publication on. We're taking a look at OCTO 15, which took a look at an early warning system for root zone scaling and what that would look like. We've done an extensive review of all this material that we could dig up on this topic and how the term actually came to be, which was kind of interesting, which will be covered in the paper. We wanted to take a look at what was going on in

OCTO 15 and elsewhere on this topic space and comment on what that means, what it would look like, what kind of information is out there, how useful something like that would be in the current environment and the current state of knowledge of what all you can look at to get some sort of early indication of problems with the root service operations, etc., and how the root server system evolved over time and adapted to various challenges and the various things that have been added to the root itself as far as … It's not just the number of TLDs but there's also various DNSSEC and the various types of DNSSEC that have gone along and a whole bunch of things that have changed the way it has worked over the years.

As I said, we're very close to publication on that. I'm not going to steal the thunder of what that's going to look like, but I don't think there's going to be anything really controversial out of it. Let's put it that way. We tried to do a very thorough view and tie it all back together to both OCTO 15 and other work that's been done either by SSAC or RSSAC or elsewhere and do hopefully a fairly definitive review of all that.

So if you have any questions about that, feel free to ask them. As I said, I don't think there's going to be anything controversial here. But it is really interesting, especially to tie it together and to look at the history all in one place.

I'm not seeing anything on that. Feel free to chime in and we can catch up on that at the end.

Next slide, please. Steve or Tara, we'll let you give an update on the … This is one slide, so … Whichever one of you has the duty. Tara, it looks like you do. Go ahead.

TARA WHALEN: Yeah, I do. Thanks, Rod. So on behalf of my Co-Chair, Steve Crocker, I'll give the update of what's going in the EPDP Phase 2A work.

So these are the elements of this EPDP that SSAC was most involved in—or elements that we contributed. To give folks a little bit of context, this was some questions that were remaining about differentiation of legal and natural persons' data and also the feasibility of unique contact to have a uniform anonymized e-mail address. The legal/natural persons was a question that SSAC was most involved in.

As laid out here, really what we're talking about is how to make the determination about public disclosure of data. The framing that has been used is using legal and natural persons as an approximate proxy for this as to whether you would, for example, the data of legal persons and have it treated differently from natural persons' data, which would have personal data inside it.

And what we have been promoting is the use of more explicit declarations, more clarification, on these elements and also to ensure there is clear and explicit guidance because, of course, folks are going to need to know what they're signing up for. If you are declarations about how you would like things to be disclosed, you need to know what is likely to happen to your data.

There's also this, I guess, issue to deal with, which is that we have a bunch of existing registrations that are in one state as we may move to a different ecosystem. So if we aren't sure about the status of those folks—for example, they haven't yet gone through the process of having signified whether they are a legal or natural person—then we need to have a state for that. So, in this case, we recommend an unknown state that we have proposed. So then we can put that forward and then, as this information is collected, we can update that and then gradually reduce the number of registrants who are in this unknown phase.

We were also careful to note that whatever model we developed, whatever framework and data elements we put in, we want there to be some extensibility. Right now, we're operating in an environment where we're focused on legal and natural persons, but there may be different elements of a registrant that we may need to take into account about what their desires, what their requirements are, what their regulations are. So we want to make sure we don't nail our feet to the floor a bit, so we don't constrain ourselves unduly, and as things change, that we'll be able to respond accordingly and publish things in the right way.

And we're not being overly prescriptive, of course, on how the registrars and registries engage in collecting this information. They have ways that work best for them as to how they provide guidance, the stages at which these registrants are asked for this information, how they obtain consent. We want to ensure the registrant is well-informed and given the right set of courses.

And of course all of these things we put forward are consistent with our general stance, which is that we want to have maximum disclosure in the sense of availability for information that is useful for things like investigations and abuse in the system and also that there'd be an expected use of differentiated access. So if there is there is security resource or different uses for this data for people that had different perhaps credentials or different requirements, those things will be able to be accommodated in a system.

Right now, we're in the public comment period. The draft initial report has gone out and the public comment period is open until the 19th of July. So folks are welcome to have a look at that and weigh in.

And that's it from me. So we're open for questions or any follow-ups from Steve.

STEVE CROCKER: Thank you, Tara. Let me just add three very specific points. Your first point that legal versus natural is an approximate proxy for whether the data should be publicly disclosed—let's put an emphasis on that "approximate" because there's certainly cases where the decision would go the opposite direction, where legal persons would want their data protected, where natural persons would choose to have their data exposed. So it's very important not to absolutely tie these things together.

Next is that, in the discussions that have taken place in the working group since we put this slide together, I think there's been a preference

for the word “unspecified” as opposed to “unknown.” So if anybody is tracking carefully and you see a discrepancy, it’s because there’s been evolution in that.

Further, the idea of “unspecified” is not the same as not having answered the questions. So it’s more of “I’m not going to tell you what the answer is. So whatever you’re going to do with my data, do it without the knowledge that you’re asking for.”

Finally, with this reference to differentiated access in a future system, there’s a much longer discussion to have about the status of that work and the impact that the current state of development is having in these discussions. That is, (not to be coy about it) there is a credibility issue with respect to whether differentiated access is ever going to come into existence, whether it will be efficient, cost-effective, etc., etc., and that is overhanging these discussions within the working group, where some of the people say, “Well, better get as much as I can now because I’m not sure I’m ever going to get the data through any other path.” So in a sense, that looks at the whole process. There’s a certain tilt to the whole thing.

TARA WHALEN: Very helpful. Thanks, Steve.

ROD RASMUSSEN: Okay. Any questions or comments on this topic before we move on?

I'm not seeing anything. Okay. Let's move on to the next slide. Steve, I think this is for you. A quick update on that.

STEVE CROCKER: Thank you. So there is a new GNSO working group looking at registration transfer policy. One of the peculiarities or features of this particular effort is that it's on a superfast track which has raised a process question. SSAC, as other groups, have been invited to participate, but our processes inside of SSAC aren't fast enough to respond at the speed that this working group is trying to respond.

So the resolution of all that is that I'm serving as an invited subject matter expert as opposed to a formal SSAC representative. From my point of view, I tried to do roughly the same thing and keep SSAC colleagues informed.

The main focus of the working group is on auth-code and loss of access to contact details following GDPR and trying to sort out whether the process should be refined at all, whether 60 days is the right amount of delay, whether there's any security issues with respect to using auth-code, and so forth.

However, one of the things that I've raised and that Jim Galvin, who's also a member through the Registry Stakeholder Group, has been also active on is that there is a somewhat awkward relationship between the registrar and the DNS operator. That is, it's very smooth. It's not awkward at all in most cases in that the registrar provides DNS service, but when there is a transfer from one registrar to another, if that

registrar is providing the DNS service, then they stop providing that DNs service. So there's a forced transfer of DNS service along with the transfer of registration.

That's fine in a lot of cases, but in the case that the customer—the registrant—wants continuity of service, and in particular if they want continuity for a signed zone—continuity means the references both resolve and validate—then there is a rather tricky sequence of steps that have to be taken and it requires cooperation on both the losing side and the gaining side. That's been outside of the purview of the ICANN contracts because DNS operations are not part of the contracts, whereas registration is. And it has also been outside of the tension and focus of the registrars primarily.

So there's a discussion. I have to thank Roger Carney, the Chair of the group. They've been very thoughtful and considerate in allowing that discussion to take place. That will play out over some period of time. We will probably come up with an SSAC recommendation aimed at registrants in their dealings with their DNS operators, some of whom of course will be registrars. But they may be third-party DNS registrars.

So that's the primary state of affairs. There's some additional discussion about whether the auth-code is the equivalent of a one-time password and exactly how strong it should be and what the rules should be about its lifetime, etc., etc. Thank you.

Jim, do you want to add anything to that? You're as much in there as I am.

ROD RASMUSSEN: Shaking his head no.

All right. Any quick questions or comments on that one?

No? Okay. Well, if you think of something, throw it in there. We'll catch you at the end.

Next slide, please. The organizational review implementation. That's basically done. So I don't know if anybody in the review implementation team would like to say anything on that one. We didn't really put much on that. You can see the slide there. But we are what I would call done. There's some clean-up work that is being implemented. We're waiting for some update, which I believe is the only thing we're waiting for there.

Anybody else want to chime in on that? Who's involved in the review work party?

Okay, nope. That's fine.

Let's move on to the next slide. Threats to Internet naming and addressing. This is our threat scan, our ongoing internal work that we kicked off in order to identify gaps and areas that we wanted to potentially form work parties and do reports on. The Routing Work Party is an example of one of the things that came out of that. We're also using that to identify skills, etc., which Julie is going to talk about here in a little bit, for future membership outreach so that, if we have gaps or areas where we have coverage, we could likely have some more

people with expertise in that are to help share the load for those who are [inaudible]. We can recruit for those positions and recruit members to help us with those areas.

An update on this is we had a really good call with the Board Technical Committee in the last two weeks to go over this—where we stand with that—and there has been a lot of work done on really categorizing the threats that we see out there. And then we overlap that with some of the things that the Board Risk Committee has been looking at as far as risks to ICANN and the identifier system, etc. So we've got into depth on that, and I think we've done a pretty good job of having those discussions with the Board to make sure that they're fully aware of the things that we've identified and vice versa.

Currently, that stands as an internal work. We're discussing with the BTC and potentially the whole Board on whether or not to try and move forward with anything beyond that. As far as publication goes, one of the biggest concerns we have right now that that is a list of risks without a mitigations, which is a dangerous thing to publish as it stands. So something like that would require some fairly heavy lifting to flesh out, which may be a challenge for us to do, given all the other work and the volunteer nature of the SSAD.

So we're in discussions and consultations (I guess you could call it) on how to use that going forward beyond this internal tool for the SSAC, which we will continue to keep updated for ourselves to be aware of things and to share with the Board Technical Committee and Risk Committee on an ongoing basis. So that is being done. As I said, we'll

see if there's a more public version of this that we could somehow work out towards publication at somehow work out towards publication at some point in the future. There are no current plans for that, but as I said, we are still figuring out if there might be a way of doing that going forward. I know there's been some interest in the community in having a look at something like that.

So any questions or thoughts on that that people want to share before we go on?

Merike, was there anything you wanted to add on that from a Board liaison perspective?

MERIKE KAEO: I can just add that the SSAC has been collaborating very closely with BTC just so that we're all informed in terms of the work that's very similar, [where] the SSAC took a much broader view of the overall threats to the Internet addressing and naming.

But nothing really more to add. I think you've covered everything. So thank you for that, Rod.

ROD RASMUSSEN: Thanks, Merike.

Okay. Next slide, please. Okay. So this is where we start looking forward. We're getting toward the end of our topics here, so prepare any other questions that you may have. This is [inaudible] is a new work, work parties, etc., that we may chime in on here in the near

future, where we have various levels of interest within the SSAC on doing work here in some of the things that are evolving around the Internet, and input from the community have put these onto our radar.

You can see the list here. We've got the, how is the DNS resolution process changing over time and what are the things that we're seeing that will be impacting that and what are the implications of that. And there's a bunch of things that are all kind of lumped under that that are related but not necessarily the same causation or even areas of the same technical area. However, they are all related, and that's why we lumped that together into one potential piece of work.

Related to what Steve Crocker was talking about earlier is some of the areas around management of DS keys, etc., especially around transfers. But there's also ongoing operations. The HTTPS is a bit related to the evolution of DNS evolution. It's also related to the DoH/DoT paper we put out a year or two ago—I think two years ago now. There's this whole blank year of COVID where everything seems like it was last years but it was really two years ago.

Then we have been talking about taking a look at the various bits of data that are flying around and what may be missing, etc., from being able to do SSR work and analysis with as a potential for us to work on. I know there's some various sets of data that are being collected in various places that is not always understood or catalogued in a way that makes it useful, but it could be potentially.

Then there's a very specific issue around being able to hijack domains based on [inaudible] delegations or main domain servers that are kind

of legacy, based on a domain that may have expired or moved on. And there's an area that initially has come up recently. I believe there's some other work going on in this in ccTLDs that would be of interest as well.

There are other topic areas, I'm sure. This is one of the ones where we're always looking for input. We just had one earlier for taking a look at updating SAC 023. So that might make this list. Next time we're on this, we'll have a look at that internally.

Next slide, please. Julie, I know there was something else to important to talk about [inaudible]. Over to you.

JULIE HAMMER: Thanks, Rod. Next slide, please. What Rod mentioned earlier was that we have a range of categories of skills that we look for in our members. We've, within the last couple of years, updated our skills survey and added quite a bit more detail to it, which has been informed by the threat scan that we did. We categorize the skills that we're seeking in members—obviously not every skills in every member, but we're looking for these nine categories. Each year, we update the survey for every SSAC members. And for potential new members, we ask them to fill out our skill survey so that we can identify the desired skillset they might be bringing into the SSAC.

We also use the skill survey as a means of seeing which members have skills for the various tasks and work parties or public comments that we

want to produce documents about. So it's very helpful for Rod and I as to who we should be enticing to come and do work on particular topics.

Next slide, please. What we have recognized is that there a number of areas where we are looking for skills where we are either lacking or only have a very limited number of members with skills in these areas. We've got them listed here: in ISP operations, large-scale measurement, registrar operations, browser development and testing, mobile apps development and testing, low bandwidth resource constraint, Internet connectivity, Red Team experience, risk management, and law enforcement experience.

So we've been trying to do some outreach into the community, which in the current environment is not as easy, as many of you would know, as it is when we can actually engage face-to-face. But thankfully we have had some success in our outreach, and we've just recently brought on board three new members who will become members once we've submitted their names to the Board for approval. So that doesn't mean, though, that we're not still looking for a number of new members with some of these skills or with other skills in areas that we do quite a lot of work on.

We're particularly interested in getting greater diversity in some of the backgrounds of our members, particularly their regional backgrounds, because we recognize that they bring very different perspectives on technical issues. So we're particularly interested in seeking applicants from African, Latin America, and Asia-Pacific.

Next slide, please. Anyone who is interested in approaching us with a view to membership or who know someone who might be a very useful candidate, please do put them in touch with either Rod or myself directly or any member of our wonderful support staff by ending an e-mail to the e-mail address you can see on your screen now.

Thanks, Rod.

ROD RASMUSSEN: All right. Thanks, Julie. Yes, please, this is our time to outreach. As mentioned earlier, we have a few new members that will be coming on. It'll be announced shortly. But are definitely still looking to expand our capabilities and our geographical footprint and our systems capability footprint, if you will. People who work in constrained environments, etc.—really, it's important to get their inputs on how things that may be solved in one areas are still not solved in others and how that actually may—things that we're talking about implementing in one space—have a huge impact in others.

JULIE HAMMER: Rod, if I may, I might just answer Donna's question from the pod.

ROD RASMUSSEN: Oh.

JULIE HAMMER: Donna has asked, "As new members roll onto the SSAC, do other members roll off, so to speak?" Yes and no, Donna. In this case, no. We're not, if you like, replacing anyone who has recently retired from the SSAC. And there is no fixed number of SSAC members that we aim for. It's more a range. We have been somewhat low on members for a while now, which makes it difficult to get work done.

So, in general, we feel that the ideal range for the number of SSAC members is somewhere between, say, 35 and 40. But it is flexible. So the quick answer is, no, that's not the case. Thanks.

ROD RASMUSSEN: Great. Just to add, unfortunately we did lose a member, literally, this past spring. Ben Butler passed away. But we don't have a quota or anything that we replace. As Julie says, we're a little bit low on our folks. I believe we added maybe one person last year and we had a couple of people move on from the SSAC. So we're roughly about where we've been over the last three or four years with adding the three that we're talking about right now, getting it back up to that. Still, there's plenty of work to be done, and having more hands to do it would be very helpful.

I think the next slide is Q&A. And I see a couple more questions are on here. If I could get the next slide, just to make sure I'm right on that.

KATHY SCHNITT: This is the next slide, Rod.

ROD RASMUSSEN: Oh, okay. Yeah. Oh, questions to the community. Okay. Thank you. All right. Caffeine is still waiting to fully kick in here. So, yes, this is the time to bring things like that up.

So I see the first one is … Jeff, do you want to take that one on? Mr. Bedser? If you're available.

JEFF BEDSER: Sure. So the question in the pod is back to SAC 115. "Perhaps SSAC recommends otherwise, but looking at Page 21 of the slides, shouldn't the big orange oval in the center be ICANN? This entity needs to take up all of the tasks for the benefit of the industry groups and the Internet community. I think that ICANN should be the perfect non-profit organization acting as the common abuse response facilitator. Wouldn't it be too time-consuming to form and look for another organization that the Internet community trusts taking up this role?"

We, as a recommendation, believe that ICANN should not at the center of it because they don't represent the full ecosystem, but it doesn't mean they couldn't be at the center of a coordinated effort for the whole ecosystem.

So I hope that's an adequate answer. I think that's a matter of opinion and debate amongst those that form it about where it will reside once created. And there's a couple of entities that are already on the road toward a potential of being a common abuse response facilitator, such

as the DNS Abuse Institute, that could take on that role and could be supported by ICANN.

ROD RASMUSSEN: All right. There's been an ongoing back and forth on the zombie stuff. Thanks, John McCormack, for that. K.C. Claffy is the lead on that one within the SSAC, and any data on that that you want to share on that is awesome. We love data. And that might be really helpful in getting that one kicked off as well.

And then the next question I have here is, "What's the status of SSAC's review of community feedback [on] SAC 114?" Just to get a little more in-depth of what I mentioned earlier on this, we have a work party that has gone through all the questions, both formal and informal, that we've received on 114. We grouped those things together and are taking a look at that and reviewing what the commonalities were. There's been a couple things that were quite evident that people had questions about and either wanted more information or citations, etc., about where we came up with the recommendations or the findings, etc., on that. So we're looking at sources and things like that to help flesh out that information.

There were some questions about a couple of the recommendations that we had. We're taking a look at the language there and how people have understood that versus the intention of it that we were trying to do and may have some updates there.

We have a formal request from the Registry Stakeholder Group, which we are looking towards having a meeting with a couple weeks after this ICANN session gets done to just have a good conversation and working meeting there to go over some of the things we've reviewed and get any further inputs there.

From there, that work party will put their thoughts together on how to either—there are many options there—potentially respond to the questions with some further documentation, potentially do some updates to the actual 114 recommendation itself. We've updated things in the past based on community input and review of our own work, especially as new items have come to life or it's become clear that the intent of what we wrote was not exactly how the community may have been perceiving it.

So that will really probably be later in northern-hemisphere summer/southern-hemisphere winter to come back with that, but as I said, we're having an open discussion there with folks that have brought things to our attention on a more formal basis so that we're all on the same page on what the intent of the SSAC was with that document and what we'd really like to see being considered as part of the process of moving forward with a potential new round of TLDs.

Hopefully that answers far more in depth than what I answered before.

Anybody else from the SSAC want to add anything else to that?

No? Okay. I think we've got that one.

Well, we're right at the top, but we've had some other questions. Any other great ideas or questions from an SSR perspective that you want to bring to our attention or have the SSAC take a look at at some point?

Okay. Well, I think that we are then done with our public meeting. And we've got done about 30 seconds before out time ran out. So usually we're over, so good job, everybody. Thank you for coming today and listening and contributing. Some good stuff came in that we'll take on as well. So thank you very much for that. We're leading to maybe some potential work you'll see in the future.

Thanks, and everybody have a good last session if you're going to the last sessions here. We'll see you in some form in October. Thanks, all.

KATHY SCHNITT: Thank you. This concludes today's session. Please stop the recording.

[END OF TRANSCRIPTION]