ICANN77 | PF – Joint Session: ALAC and SSAC
Thursday, June 15 2023 – 13:45 to 15:00 DCA

YESIM SAGLAM: Hello, and welcome to joint meeting of the ALAC and SSAC. My name is Yesim Saglam, and I'm the remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior. During this session, questions or comments submitted in chat will only be read aloud if put in the proper form as I've noted in the chat. I will read questions and comments aloud during the time set by the chair or moderator of this session.

Interpretation for this session will include English, French, and Spanish. Click on the interpretation ICANN in Zoom and select the language you will listen to during this session. If you wish to speak, please raise your hand in the Zoom room and once the session facilitator calls up on your name, kindly unmute your microphone and take the floor. Before speaking ensure you have selected the language you will speak from the interpretation menu. Please state your name for the record and the language you will speak if speaking a language other than English.

When speaking, be sure to mute all other devices and notifications. Please be clearly and at a reasonable pace to allow for accurate interpretation. This session includes automated real time transcription. Please note this transcript is not official or authoritative. To view the real time transcription, click on the closed caption button

in the zoom toolbar. To ensure transparency of participation in ICANN's multi stakeholder model, we ask that you sign in to Zoom sessions using your full name, for example, first name and last name or surname. You may be removed from the session if you do not sign in using your full name. With that, I will hand the floor over to Joanna Kulesza, ALAC Vice Chair. Thank you.

JOANNA KULESZA: Thank you. Thank you very much, Yesim. Thank you everyone for joining us for this traditional joint session between the ALAC and SSAC. In case you're looking at the agenda and you might be confused, I am not Jonathan Zuck, but I am step-- Right. I know you're surprised. But I'm stepping in very briefly just to introduce our guests, and I am thrilled to be joined by our SSAC liaison. I understand that Andrey Kolesnikov is on the line and if that's indeed the case, I would be glad to share the floor with him directly to give us an intro on what is expected and what are the topics of joint interest that we are looking forward to discussing. Andrey, are you with us and are you able to take the floor?

ANDREY KOLESNIKOV: Yes. I'm here. I hope you can hear me. This is Andrey Kolesnikov.

JOANNA KULESZA: We can hear you well.

ANDREY KOLESNIKOV: Maybe you can even see me because I turn on my camera for a moment, low bandwidth, but let's start our traditional ALAC and SSAC meeting. I don't know how many years it's been since the first meeting started, but it's a good tradition. Let's keep it up and let's jump to the agenda immediately. So, the session will be pretty much structured the same way There'll be opening commands by Jonathan and Rod, and then there'll be a number of set topics, regarding the working groups and working parties, then there'll be elect topics, and then there'll be the closing comments for the session. Well, I won't take much of your time. So, let's give the floor to Jonathan and Rod Thank you very much.

ROD RASMUSSEN: Thank you, Andrey. Thank you to the ALAC for the scheduling the time with us. It's always one of our most interesting and I think useful meetings of the of these ICANN meetings to get together and get caught up on what initiatives are going on because we have a lot of overlap. I know there's some newcomers in the room, etc. So, they just want to let folks know that Security ability advisory committee, advise the ICANN board on a plethora or a lot means a lot of issues and around the security and stability of the Internet.

These often overlap with end user concerns because people like a safe experience when they're using the Internet. So, a lot of the initiatives that the ALAC is working on we are working on as well. And we often find ourselves in alignment on views on things, and sometimes not. And it's good to get together and have those conversations and we can often dig into some of the technical details of what may be questions you have that you may not have the expertise or as many experts as we have

on and there are plenty of experts in the ALAC, though, of course Jonathan excluded. Anyhow, I just wanted to give a little background there for digging in.

So, we're going to go through a series of topics, which I am going to hand off to the various SSAC members who are chairing those efforts within our committee. I will remind our committee members and I'm trying to do it myself is to speak slowly and clearly. So, this is interpreted properly. Technical folks have a tendency of running on fast, so remind yourselves of that. So, can I have the next slide, please? So, I'm actually going to beg indulgence and ask to skip ahead a few slides to the name collisions section because so keep please keep going, and we'll see it here in a moment. Keep going. Asset. Okay. And right there, I'm going to hand it over to Matt Thomas because he has a hard stop and I will go back and get those other topics. So, Matt, over to you.

MATT THOMAS: Thank you, Rod, and hello, everyone. Sorry, I can't be there in person. I have some personal conflicts this week. But my name is Matt Thomas, and I am one of the co-chairs of the name collision analysis project along with Suzanne Wolf. I'm sure many of you are familiar with the name collision work that has been ongoing for years. We've provided several readouts to the ALAC at the last several ICANN events, but wanted to just give another brief update here today. So, in general, the NCAP project is looking to answer a series of questions and to provide some advice from the ICANN board on the topics of named collisions.

There was a specific ask on advice regarding the three non-delegated strings of Dot home, corp and mail, as well as answering a series of 9 questions presented by the board. The group has been working with the community over the last several years in an inclusive manner. There are roughly 25 discussion group members as well as 14 SSAC work party members currently part of the effort. So right now, we are looking at almost finishing up NCAP study 2.

And the NCAP study 2 consists of 5 major deliverables of which one of them was the board questions that I previously mentioned, but it also consists of three independent or three separate reports, the first of which is a case study of collision strings looking at corp, home, mail in addition to internal, LAN and local. And that study was conducted using DNS query data from Verisign's A and JGroup servers going back longitudinally for many years. That allows us to highlight the changes have occurred within the DNS relevant to those strings and their named pollution properties. That report has been completed. It went out for public comment and the discussion group has called consensus on that report.

A second report is a perspective study on DNS queries for non-existent top-level domains. And this report was important because it allows us to look at the various different vantage points in the DNS hierarchy to understand how where and what limitations there are when looking at name collision DNS telemetry data that can be used to assess name collision risks going forward. That report has also gone out for public comment and has been called for consensus within the discussion group.

Not listed on this page, there is the third report which has been conducted by the technical investigator hired by ICANN to help support the NCAP study 2, and that work was to look at the actual name collision reports received by ICANN since 2012 until now and to do a root cause analysis. That report has also been finalized. It has not gone out for public comment, but will be part of the fifth deliverable that I need to mention, and that is the entire study 2 report, which is essentially a large report that brings together those three independent reports, the board questions and puts forth hopefully a package of information relevant to ICANN board in terms of how to proceed with name collisions going forward.

So, to highlight some of the key findings, the discussion group has Through its research, identified that name collisions still are a large problem that exists in the DNS. And that a large amount of the name collision problems stem from things like DNS service discovery protocols and suffix search lists. The discussion group has also come to term a topic or a set of measurements called critical diagnostic measurements as a way to help quantitatively describe potential risks associated with a string.

We'll talk about those CDMs in just a little bit. But the important thing for you probably to understand here is that as those CDMs increase in both volume and diversity, where diversity can mean multiple different dimensions of things like the number of networks or second label domains or something of that nature, the mitigation and remediation becomes more problematic and the potential for impact or harm also increase. Finally, I will also note that we have identified in the discussion group that there are opportunities for existing measurement

platforms to be extended to help applicants a priority on submitting their application to be potentially more informed about potential name collision risks for certain strings that they might apply for.

So, some of the critical diagnostic measurements that the group has decided, have talked about, also stem from some of the measurements that were used back in 2012 by interline and Jazz. Collectively these critical diagnostic measurements look at things like query volume, query diversity that could span both the directions of IP or ASN distribution and other properties of the DNS queries like q type diversity, the number of labels present, it also needs to include a quantitative component that looks at other good old fashioned investigative discovery of why those name collisions are potentially leaking out into global public DNS and what potential harm that they could pose due to the nature of those queries.

Overall, though, the group or the goal and objective out of study 2 is to ensure that name collisions can be assessed. And in order for that occur, it means making name collisions visible. One of the things that the discussion group has noted throughout its years is that there was a technical gap from 2012 till now. The DNS ecosystem has dramatically changed We've seen the introduction of public recursive resolvers. The deployment of protocol changes like queue name minimization, aggressive NSSAC, NX domain cut, all of which impair previously relied upon observational abilities at the root to do that assessment.

So, in order to do that, a proposed workflow has been established that deviates slightly but still in the spirit of the 2012 round that allows name collisions to be conducted in a more sustainable, repeatable and

deterministic way. And this is what that workflow would look like at a high level. One of the important things that you'll note on this graph is that there are four distinct offramp options. One of the things that the discussion group really spent a lot of time discussing about was that there needs to be the ability for applications and applicants not to get stuck in a limbo situation like corp Coleman Mail.

The proposed workflow now provides four different distinct opportunities for that application to be withdrawn or terminated. The workflow in general starts with the applicant deciding what type of information or what string that they would like to apply for. As I mentioned earlier, we've discussed that there are opportunities for existing platform measurements already developed and run by ICANNs such as their ITHI program as well as their magnitude page to be extended and have that data publicly available to applicants prior to their application submission that they can go there and at least get an initial sniff test to see if they're string might be problematic from the get go. The applicant would then submit their application and the process begins. The process would then go to the purple section there in which the TRT, the technical review team, would conduct its own static assessment.

This technical review team which I'll describe a little bit more on the next slide will again look at the publicly available data that it has or is generally available to do an initial assessment for collisions, risks, and impact for that particular string. If it is a heightened risk and the applicant does not wish to proceed, there's an offramp opportunity, otherwise, the string will be delegated into the Root server system and it will be pointed to an authoritative name server controlled by a trusted

neutral party or the technical review team during a phase called the passive collision assessment in which that's TLD or applied for a string will still essentially return a non-existent response or no air, no data, but it has now allowed the TRT data team to perform a near like root server system wide collection of DNS telemetry data for that applied for a string.

That collection also gets around some of the technical problems that I mentioned before like queue name minimization in which certain aspects of the DNS data are now more rich with information that a better name collision assessment can be conducted. As that time period goes on during PCA, if there is a heightened risk or concern identified by the TRT, there would be an offramp opportunity for the applicant again before going into the final phase of assessment, which is called the active collision assessment.

Active collision assessment is essentially instead of returning a non-existent response for the string, a routable IP address is returned in which the TRT is then able to actually log and monitor connection attempts and better understand what kinds of services and systems and software are potentially using that string to better inform the name collision assessment risks. Again, there is an offramp option here if the assessments identify potential risk or harm, but ultimately the TRT will put together its package of a recommendation, send it to the board, once the board has made its decision, the TLD or string would then be granted to the applicant.

So, I mentioned before that this technical review team in the neutral third party are two main important features of this new workflow. The

TRT is really a body of folks that need to have a deep understanding of DNS and understanding how Internet infrastructure works and how the DNS interfaces and works with applications and services. They're going to need to be able to understand the data collected those critical diagnostic measurements that we mentioned before and how to put that into ultimately a risk assessment kind of model.

We foresee them having four major responsibilities and that is assessing the visibility of name collisions, documenting the data findings and recommendations, potentially assessing any mitigation or remediation plans if they are necessary and then also serving in an emergency response function. Similarly, there would be the neutral service provider. This may be a separate entity from the TRT. It may be the same. Again, this is an implementation detail, but the four responsibilities of this neutral service provider would be to operating that passive collision assessment environment as well as the active collision assessment do any log processing and analysis for the TRT and again serve in an emergency response function.

So, I'll note that the discussion group is in the middle of finalizing its findings and recommendations. We are more than welcome and happy to have new people join the group. So, if you are interested, the information on joining the discussion group is here. You can always reach out to me. I'm happy to take any conversations offline. And with that, I'm also happy to answer any questions.

ROD RASMUSSEN: Thank you, Matt. Are there any questions in the room, or in Zoom for that matter? How many should check there? No questions. Okay. Oh, sorry. Jonathan?

JONATHAN ZUCK: I'm going to ask an expert question if I might, Rod. I'm very fascinated by this risk assessment concept and sort of this offramp concept, and is there different levels of responsibility for an applicant based on the pre-assessed risk or are they sort of separate in that helps them decide whether to apply, but once they do, the set of responsibilities is the same? Is there is there a threshold or concept within your quantitative work that says that you have to do this level of rigor if you're at this level of risk? Does that make sense?

MATT THOMAS: The very good question, and unfortunately, with the quantitative measurements, there is no magic number the discussion group feels that above this level of query volume or above this measurement in some CDM that you're in one bucket or another. Unfortunately, it is more of a somewhat subjective, quantitative qualitative assessment that really will require the expertise of the TRT to kind of make that judgment.

ROD RASMUSSEN: All right. Have we got any other questions? All right, and as Matt said, we're coming down to the end here. It'd be great to have more people having a look at this and providing input it is late in the day, so you're

going to have to catch up fast, but this oh, well, that's before 78 now. These are the same slides, basically, we had last time. Yeah. But the NCAP is very important. So, I want to make sure where everybody's where we're at. we had a great meeting here in DC, and we're looking to see if we can add some an extra face to face as a result of the great progress we made here and replicate that so we can make sure to get this done and launched.

Well, ready for the other people to do whatever launching they want to do as a result. So, thank you, Matt. Thank you. You backed it up. Can we back this up now? We'll go back to the DS key management stuff. Our DS automation. Thank you. And I'm going to hand this over to Peter. And Peter, if we can just hit this a little bit of the highlights. It's the same slides we presented last time, but we have made progress on where we are with potential going forward. So, if you could concentrate on that part.

PETER THOMASON: Yes. Thank you. This is Peter Thomason in a SSAC. Are these slides being presented directly from Google slides? Because it doesn't seem to be the most up to date, but I can also present from those. If it's difficult to reload.

ROD RASMUSSEN: We're going to have to go with what we have in front of us.

PETER THOMASON: Okay. That's fine. All righty. So, as we have discussed multiple times in this forum, when you have a domain that has DNSSEC, and the DNS operator rotates the keys, it is necessary, especially when the key signing key gets changed to update the DS record set at the parent domain and that can be done in an internal operation between the registrar and the registry when the registrar also is the DNS provider for that domain, but when the that is not the case, then somehow one has to convey the DNS key parameters from the DNS operator who can be some third party to the registry. And so far, that is often done manually by the registrant, and the process is quite well, not successful in all cases.

So, there is a paper that shows that about 40% of registrants only succeed in doing that when they attempted and it is perceived as frustrating and difficult. So, what's needed is actually more automation there. There're multiple ways to do that, and so I'll talk about that more in a little bit. I think on the next slide, we have a diagram that shows the manual stuff. Exactly. So, At the bottom there is the child level, and at the top is the parent level, and in the middle, WHOIS neither the child nor the parent is the registrant and perhaps some domain reseller. So, when you have secured your domain, the DNIS operator has put their zone onto the authoritative DNS server, and then the next step is when the DNS operator is not the same as the registrar, that the registrant fetches the DNSSEC details from the DNS operator and somehow passes them on to the registrar through some web interface at the registrar or the reseller or something. And then the data takes its way all the way to the TLD server.

The difficult to realize in this intermediate step where a manual action is involved. Yeah. So, some automation is already there. Some ccTLDs about 10 have automation implemented. The way it works is that The DNS operator at the child publishes so called CDS records, which indicate that the change of the parent site DS records is due. And the parent can discover that by looking for these records, for example, daily. So that's called CDS scanning. And scanning has two problems.

First of all, when you look daily, then the convergence time is about a day. So, when a change happens, the registrant or the DNS separated doesn't know immediately at what time the change will be affected. Instead, it'll take some time up to a day or so depending on what the scanning interval is before the changes are detected and take effect. And second, there's a scaling problem. So, of course, the more frequently the parent looks for the CDS records in the child, let's say hourly, then there would be 24 times as much load as if you do that daily, and also, depending on how many domains are, securely delegated from the TLD, you would have to scan all of those every day, and so the load goes up as that number increases. And that doesn't appear to be problem so far based on the numbers from ccTLD registries who already do that, so we ask them and they say, are not too concerned about scaling, but it could be that in the future expectations get stricter and perhaps people want the scanning to happen more frequently and then it could be a problem.

So, one mitigation that could be introduced to avoid that stuff is kind of improving the opportunistic scanning that happens every day by having a trigger to do the scanning. So that the DNS separator would somehow let the parent know. Okay. Now I put a CDS record in the child. You can

now look for it. You don't have to wait until your timer is due tomorrow night or something. so that's called a NOTIFY. Similar thing exists for replication between name servers, so you might have heard about NOTIFY, and this is a generalization of that. But the generalization hasn't been specified yet, and so that would have to be done, and the according pieces of software would have to be developed.

That's kind of in an early stage, but there is kind of light at the end of the tunnel regarding these scaling and delay problems with the CDS scanning approach. And then in addition in the RRR model, there are two players at the parent's side from the perspective of the DNS operator. one is the registrar, and the other is the registry. And both of them potentially could look for the CDS records in the child and affect the update the registrar could do that by talking to the registry via EPP, and the registry could do that directly and just update the DS records and the parent at the risk that the registrar doesn't learn about it and perhaps the database will be out of sync.

So that's a bit of a challenge and the question is who should be doing the scanning? And it seems that it's more natural in the RRR model when the registrar would be doing the scanning. So, I guess that is what would be the smoothest approach and there's actually one registrar who does that, they're called the domain name shop from Norway. But there's also about 10 ccTLD registries who also do the scanning. So, let's say, for example, under one of these 10 ccTLDs who do the scanning, there's a domain that's managed by a domain name shop who also does the scanning, then you have two parties scanning, and potentially there could be a collision, and it would be good if that were eliminated by those two parties the registry and the registrar agreeing

on who does the scanning. Now that sounds very detailed. What I'm trying to convey is that although the scanning approach sounds nice, there is as usual the devil in the detail of several devils.

All right. So, this is a diagram that shows how the situation is simplified when you have that kind of scanning, so as you can see, the intermediate layer between the parent and the child, where the registrant had to do something manual is gone. This diagram shows how the registry can do the scanning directly so they would discover the CDS records from the authority name server from the child DNS operator directly, and then update the DS record in the TLD server and also let the registrar know through some EPP back channel that a change has happened, but that is not supported by all registrars. And as I said before, it can happen that the registry database is left in the jag and it gets out of sync. The next slide shows same thing, but with the registrar scanning, which is the more natural approach I just mentioned.

So, the registrar is already the party who usually gets new DS information, so far manually usually from the registrant and so the scanning, which is arrow number 2 here, would just be another way for the registrar to get the same information, directly from the DNS operator, and then all the remaining steps that saw the registrar forwarding that stuff to the registry and so on and so forth would remain the same. So that's the that's that. So, we have some preliminary positions in the SSAC about this whole topic and so, this stuff as I said is preliminary. We will, I guess, settle some of the details in the upcoming weeks and the goal is to publish a report on this ideally before the next

ICANN meeting, I think, October it is in Hamburg. It'll be very close to my home.

So, the preliminary positions are that the DNSSEC is really hard to set up for people, lots of people don't manage to do that successfully, and DNS tech adoption has remained pretty low and for these reasons, it would be good if DS automation, specifically the automation of updates where required functionality and registries and registrars would support that. And there's various things that you can try to make that happen, one thing would be if ICANN org would hasten the DS Automation implementations, and, for example, support registries and registrars to implement that and also supports participation in the ITF processes that are necessary, for example, to Settle the issues related to the scanning delays and scaling in combination with this notify approach that I talked about earlier, which isn't specified yet.

So, it'd be good if these things were pushed a little bit, and then there's also domains which are not under gTLD or ccTLD authority, but under RIR authority, like the reverse DNS stuff. In principle, you can also do DS automation there, and in fact, RIPE does it, but there is four other RRRs too. And so similar advice should be provided to them. As I said earlier, if the scanning is done by the registrar, that's the most natural approach. So, we call that the ideal arrangement here. But still, some registries, especially in the ccTLD space already do the scanning, and that should not be discouraged because we'd be a step back. Instead it would be great if the registrars would include support for the DS automation in their services and once a parent who does the scanning so far knows that a registrar is now supporting it also, they can cede

control over the automation stuff to the registrar incrementally as registrars and the necessary capability.

So that's pretty much what is to be said about CDS updates. Then, interestingly, the notified part also involves the DNS operators to actually send them. So, this is an additional group of entities that would be affected by this advice. Ideally, DNS operators would, once it has been specified, support these notify things and actually send those messages whenever there is a key change that requires the DS records to be updated. But as I said before, this is kind of very early to make that an actual recommendation because the basis for acting on it is not yet there, so this is kind of very preliminary. I believe that's all the slides we have, and if there's any questions, we can do them now, I guess.

ROD RASMUSSEN: Any questions on this?

MATT THOMASON: And if there is no question

ROD RASMUSSEN: No questions in Zoom? Oh, James Jim.

JIM GALVIN: So, Jim Galvin, for the record, I just wanted to call out. I don't know if ALAC has a standing position per se about things that are DNS related but I did want to call out and ask that ALAC consider this obviously, all of this and this issue with DS records and such directly affects users and

registrants and their ability to use and deploy DNS Act. And I just wanted to call that out in case it wasn't obvious. I know that sometimes these things can seem overly technical and you don't immediately realize what's going on. But would like to encourage ALAC to take some time to think about that. We certainly have folks such as Peter and others in the work party who I'm sure would be delighted to come and talk more about it with some folks. But give some thought to how you feel about this work and taking a position on it perhaps considering whether or not you want to have a role in advancing it, which I think we would very much welcome. But far be it for me to speak on behalf of SSAC, but Rod can add to that.

ROD RASMUSSEN: Thanks. You get to do that to the board. You're used to it. Oh, go ahead. I was going to add something that what Jim just said there and it might address what you're about to ask. We had the DNSSEC workshop, which I know many of you are aware of, but still plug it again. We do that every ICANN, meeting. we have a workshop. The policy form, it's a half a day. And this time, we dedicated completely to this issue but also putting this issue in context. It's like, why are we spending so much time and effort trying to get this weird esoteric but important operational issue dealt with.

And this is about kind of completing the foundation for DNSSEC and making it a truly solid operational service that other things can be depended upon. And this is one of those things that hangs up the adaptation of DNSSEC which in turn affects businesses and other folks who want to create new services that could use the benefits of a

universally trusted cryptography and basically a whole system that you can hang digital trust off of and interoperate between islands of trust throughout the world that maybe do using different systems, but the only thing out there that is ubiquitous worldwide accepted both from an operational and governmental perspective and a technical perspective is DNSSEC. And it has everybody's like, where are we spending almost time on energy working on it?

Finishing off those bits is really important to be able to then to hang dependable services that governments can use it to interoperate businesses, can trust it so it'll be there and be able to make money selling things that are based on those services. A lot of interesting things there So I encourage you if you in your copious spare time, we do record those sessions. And this year's we'd actually put this all-in kind of perspective.

This might be a good one and to take a look at, and also feel free to talk to SSAC members about that, especially like Jacques over here, wave your hand, Josh. He's doing some really cool stuff in Canada with digital identities and being able to get different provinces to interact with each other on their digital driver's licenses and things like that, that it requires something like DNSSEC to make work. So that's why we're spending so much time on this. Because at the end of the day, it's for the users of the internet to be able to use it for all kinds of really cool new stuff. So, Jonathan, you had a question.

JONATHAN ZUCK: Thanks. and I agree completely. That wasn't going to be my question. Like, why are you spending so much time on this? in fact, the I think the

At-large community, and it was before my time, but took a very early supportive stance on the name collision work that you did way back before the last round, and been very supportive, I think, along the way on that work here. So, we haven't yet taken to my knowledge and official stand on DNSSEC in particular, but I think we've been largely supportive. And I feel like the form of that support has mostly been in form public comments when you put things out for public comment. But if there's something more creative we should be doing to evangelize this, I am personally and increasingly more folks are interested in how we can make use of the our largest, so to speak, to evangelize some things out to broader audiences as well.

So, we should probably have more conversations, not only about the substance, but some of the mechanisms of who you think needs to hear, because we might be in a position to help those people hear the message.

ROD RASMUSSEN: Let me say that really aligns well to some of the things we're in the process of doing, we have our annual workshop in September. This topic, the kind of the grander scheme of where DNSSEC and security in general pain thinking ahead to the ICANN strategic plan. As on our docket, we also have work that we're starting to get actual some traction with ICANN's communications department. I'm about creating messages and all that. So, we'd love to, I think, collaborate on that because you are a great method or channel, if you will.

JONATHAN ZUCK: Amplification.

ROD RASMUSSEN: Amplification, yes, for those things. So, love to do that. Okay. Oh, I'm sorry. Hadia has a question? Go ahead, please. That was a fail. Can we try again?

HADIA ELMINIAWI: Okay, this is Hadia for the record. So, I think the obvious part for At-large would be promoting the adoption of DNSSEC. But being nontechnical, well, not all of us of course, are nontechnical, but the At-large being basically nontechnical people How do you expect us to promote the adoption of DNSSEC? Thank you.

ROD RASMUSSEN: That's a great question. And so, some of the things we need to dig into, I think that at the end of the day, it's not the technology or itself, it is, but it's what you can do with it and so if you think about, I want to use cases, so I want to jock my favorite use case I've seen so far as the one Josh uses all the time. Somebody has a digital driver's license or Canadian they go to Tokyo, they want to go into a bar, they want to get a drink, how do they prove they've got a digital driver's and what's that? But the technologies there, you could actually do that. But it requires that underneath.

So, forget the advocates 20-year-old Canadian girls who want to go to bars in Tokyo. And that's the example. And think of use cases like that where it's end-users who want to do something and the technology can

enable that. And that's something we can probably dream up so a lot more things like that. That's what you could then advocate on. I want this to happen, and I know that in order for that to happen, we need this technology enabled.

MATT THOMASON: And perhaps I can add to that a little bit. So, the basic promise of DNSSEC is not only those additional drinking use cases, but also, of course, to prevent DNS spoofing. And so how do you evangelize that? So, like 10 years ago or something, HTTPS really took off and let's came along, certificates became free, and adoption has really gone up.

And it was pretty clear to everyone, encrypting HTTP traffic, does make sense even if you don't know exactly one of Cypher supers. And so, I think the same applies to DNSSEC. If you have a domain name, you should be able to expect that it is secured with the whatever the state-of-the-art technology is just like you can expect today that your web host is going to use TLS for your website. And if you buy a car, it has an airbag and you don't have to do like any airbag app stuff, it just works, and that's what it should be for domains.

ROD RASMUSSEN: And I had a request from Barry.

BARRY LEIBA: Hi. This is Barry Leiba. What Rod and Peter answered is why you should be promoting this, but you ask what you can do to promote this as a nontechnical person. And I think the answer is that as the At-large

community, you reach out to the operators to the service providers that you're all using, and tell them you need DNSSEC, and here's why, here's the why that they've told you about, and you go out and promote that. And say, I'm paying for your services and I need this service.

AMRITA CHOUDHURY: Thank you. And I'm Amrita for the record. And just to take a cue from what Barry was saying, what we can do and what-- as in we can be the amplifying channels, but what we would need is more communication specifically made for different targets. For example, if you're talking to policymakers, we need a one pager as to why it is important. If we are going to an ISP or a service provider, we need something which we can show to them as a benefit, not only for me as a user, but also for them to implement it.

So, when you are having those communication materials being made, if those kinds of communication material can come, it can help us to pass it on to our RALOs who in their respective regions, countries can promote it because when technical people are speaking or business is speaking, if it is backed by end users, it amplifies the information. So that's the help he would require. Not how we do it, but that communication. That's what we need. Thank you.

ROD RASMUSSEN: And that's perfect now. That's great input for us in our having our discussions with the ICANN communications team. Thank you for that. In the interest of time, I want to move along. Did you have something really quick you wanted to?

HADIA ELMINIAWI: Okay. really quickly. This is Hadia for the record. So, I would say all operators, actually, registered operators know about DNSSEC, and those who are actually not implementing DNSSEC. They're not implementing it for reasons that they have. So, in order to be able actually to promote this, we need to, I think, work together on having answers to the reasons for which they're not implementing DNSSEC. That's one and two, I think also we need to be working together on creating the use cases that you were just talking about. Thank you.

ROD RASMUSSEN: Great. Thanks. Yeah, why is the saying, is why is the juice worth the squeeze, right? Okay. And this was a little bit broader than the topic, but this is exactly what we're talking about in DNSSEC right now. So, I'm glad we had that conversation because it was some really good input from you guys on how we need to think about our own conversation. So excellent. We're going to do this one. Can we go to the next slide, please? That oh, and what do we have? An online question? I'm sorry. That's a comment. And thank you, John, that's a great comment. Yes, web developers and resellers, yes. Oh, yes, for the record, the ISPs and the registrars are the easiest to educate on DNSSEC. The difficult thing is getting the web developers and resellers to adopt it. That's around 25% of the gTLD markets. That's the end of comment. That's an excellent observation. Totally agree. Okay. Let's move on to the next one. It's a little confusing because they're two different things.

We got to move on. So, this is alignment on next June gTLDs. I just want to point on a couple of this. This is one of the outstanding questions. We have NCAP, obviously coming to an end. From the SSAC perspective. The other thing that is of interest, I think, in some newer developments is abuse in new GTLDs, which we pointed out in SSAC 114, we'd like to kind of have that problem addressed. So, two things there. One is we knew how new contract amendments that everybody I know is looking at, which are a big step forward, and I'm sure there's some comments about things that we'll all be providing. But in the in general, we are moving forward at least that's on the response side.

There are studies being done, both of the DNS Abuse Institute, but more importantly, I think, by ICANN itself, ICANN is not doing the study. It is funding the study, which is really important. It's a big distinction. Well, it's better in a way to in some regard, and that an independent study and is going to be a little maybe less political. However, that provides us with little to none look to know opportunity to kind of comment on how things are going during the process since it's being done with the University of Genoa, I believe. Having talked to the CTO about this a bit, there may be some opportunity to provide kind of external advice to the to the researchers.

So, we're going to take a look at doing that as SSAC, is what or some parameters and things and some ideas about where to look for or data to look for in trying to understand the causes and, and prevalence of abuse in different TLDs and registrars, etc. And that might be something the ALAC might want to take a look at as well. I'm sure you have some ideas too, and we'll be happy to share anything we come up with on that front. But that's an opportunity if we can get, then the

question is, when is the study's going to be done vis à vis lining up of when the process kicks off? And the guide book is completed, etc.

So, would like to see the studies be done and time for it to actually have an impact on how things are evaluated or people put together programs for anti-abuse, etc., and their new TLDs. Any Jonathan, did you want to add anything more on this topic area before we move on to other stuff?

JONATHAN ZUCK: No. I think we're very aligned on that and we even had a conversation earlier today that that the real milestone there is when those contracts are put together probably after the application period, but make sure that if there's a PDP that needs to happen or something like that for to address some of the discoveries that were made in the paper that there's an opportunity to do that before things are fully delegated.

ROD RASMUSSEN: At all external noise there. Okay. Could we move down to slide 23, I believe, in the presentation? I mean We're having some somebody's mic is on. I was going to turn this over to Gautam, who is, by the way, has been accepted for membership into the SSAC just recently. He has to be approved by the board, but he's as a result, he is now and officially an invited guest until the board has its approval process. But he's been done some great work in leading this work party, so we've invited him to join us as permanently as a member. Gautam, can you want to give us a real quick update on progress. Thank you.

GAUTAM AKIWATE: Yes. I'll try to keep my update short. So, the registrar name server management work party, the goal of this work party is to study the unintended part products of undocumented registrar operational practices that deal with domain expiration. And at a high level, this is because of EPP constraints at the registry on how case line where a domain provides name service to other domains in the same top-level domain is handled and because of those constraints, registrars have come up with a way a bunch of operational workarounds, some of which first domains to resolution hijacking risk. And we found as part of some of my PhD work, which is documented in the paper, which is linked in the next slide, we found that roughly half a million domains have been exposed over the last 9 years. And this includes a lot of sensitive domains including good houses, law enforcement, and some government domains.

So, if you're interested, there is a paper that I worked on during my HD work that sort of goes into the more technical details of how this exactly happens, and the work parties. So that's the technical aspect. But as part of this full party, we have invested guests from registrars and registries to sort of discuss what we can do. And the work party is focusing on two main things. The first being how do we mitigate the exist in domain resolution hijacks. So, domains that have already been hijacked. So not only were half a million domains hijacked, but nearly one third of those domains were actually hijacked by actors trying to use those domains for various purposes, primarily search engine optimization.

And then the second question that the work part is trying to address is, like, what are the options to prevent new exposures? So what practices can registrar actually adapt in order to not let this thing happen in the future. And the high-level idea that we're trying to adapt is to analyze each of these options using the framework of what are the benefits of the approach, what are the burdens that it imposes and what are the residual risks to all of the stakeholders. And, primarily, the residual risk in each of these cases is to the registrants themselves. And that said, a lot of the prior thought has been we should let registrants. The registrants should be responsible for monitoring their own domains, but it looks like in addition to registrants monitoring their own domains, it would be helpful if we sort of put in operational checks so that this doesn't happen in the future. And the hope of this framework is that it allows registrars and registries to evaluate the different options and choose an option that works well for them. And finally, the work party is also inquiring about ICANN org monitoring the zones for any new any new practices that arise that also put domains at risk. So that's where we are at.

ROD RASMUSSEN: Thank you. Yes. And there's been some really good progress on this, and we have included in this work party operators that are not normal SSAC members, but their work party members, which we do when we have to bring in some expertise, especially in the operational side of things when we're trying to dig into an issue like this. So, I think will be helpful in those crafting the messages to put in front of the audience, registries, and registrars around being mindful of the approaches to this problem and the various solutions they can choose from. Do we have

any questions on this topic? It doesn't look like it at this point. And we're hoping to progress this work and get this done.

Well, we'll see if we can get it done by Hamburg or not, but would love to be sitting here giving a final presentation with what our actual recommendations were. Not to put any pressure on you, but thank you, Gautam for that. So, if we could go on to the next slide, I believe there was a topic, yes. So, I'm going to go ahead and skip this that hasn't changed since last time. We just think that when you transfer domains, don't forget about DNSSEC. Basically, that's the bottom line. Next yeah. No. The transfer stuff, they weren't they considered it out of scope. It's like, oh, really? Okay. All right. I believe that takes us to the last slide, which was a topic you guys brought up. And I would love to hear more about your concerns and we can share some of the thoughts that some of our members have and at least in kicking around the topic in the SSAC over the last week or so.

JONATHAN ZUCK: This is a topic obviously because even though this was delegated many years ago, Can I help you? Oh, all right. Even though these domains zip and move were delegated back in 2012, they've only sort of recently released them. And I think they had some kind of imaginative idea for them and gave up on that and then just released them as kind of generic domains, and that sort of lit up cybersecurity researchers and others that saw some of the potential problems associated with zip and move because they are what you might consider trusted file extensions. And things that people feel comfortable opening as a general rule, and so people have done several videos that have shown how to create a fake

URL that makes it look like, it's a link to an attachment or something like that and there's some risk associated with it.

And so, I see sort of three components to this. one is what can be done now, if anything, and that that may be nothing. Maybe it's just going to be education since we're already involved to some extent in education. We add this to it. But I don't know the answer to that. I did have a productive meeting with Crystal Ono from Google, who's talking about how heavily They're monitoring them as they release them to see what level of abuse is taking place. And ironically, what's happening right now is that white hats are registering interesting names like attachment dot zip in order to demonstrate the problem, not to actually exploit it.

So, they're still Yeah. So, there it's all fake DNS abuse right now that's taking place, not the real stuff. But an as yet, they haven't monitored. They haven't seen extensive abuse in the doing, but it's only been a month. But so, and they're in a pretty unique position to do monitoring and so. So, we got that reassurance. There was a there was a blog from Google that sort of said, well, lots of file formats match up to current domain names and as I read it, it made me feel old because I had used many of these, but very few of them sort of fall in this category of, like, general recognition, as safe, etc., as a sort of commonly used name and so what immediately struck me is the parallel to the work you're doing on name collisions and risk assessment, because when all is said and done, all a URL is a file location. And so, it seems very similar in the standpoint of I think I'm going one place when I'm going someplace else and that use of risk assessment framework could be applied here, and that's why we thought to bring it up with you.

Barring doing anything about what's gone before, the final question is, should we be trying prior to the new round? To think about other popular file extensions, JPG, GIF, and others that, again, people have come to accept as clickable items in their emails and is that something that we the best chance of actually doing something about perhaps by applying some sort of a risk assessment framework to that. So, I guess that's the background of why we thought we'd bring it up with you guys, because you've been think You guys are hard thinkers about these things, and you had done a project that seems analogous in a lot of ways.

ROD RASMUSSEN: I know, and those are some really interesting questions and thoughts about this. In a way you have the same problem you do with somebody registering a well-known brand and phishing with it. It's a bit of same thing, it's a human psychological thing versus something that necessarily shows up in the DNS because when we're looking at name collisions as things leaking into the DNS, and if stopped leaking into the DNS, it doesn't look like it's a problem from that perspective, right, which that's what we're looking at in the name collision study, but this is a this is a different kind of very similar in that, you have some interesting interactions that can happen. We had some discussions around that. And the sorry. My cable TV provider just decided to call me right now to sell me some more services, I'm sure.

We had some discussions around what are kind of what's the background and what's going on here? Why are these issues? And you have kind of a couple problems. one is the operating system, and we

know who created that, multiple people did, but does interpretations of things, and there might be ways of fooling things to do stuff if you assume that an extension is going to be able to do something. So, there's some vectors there potentially. But there's also the human the user interface problem, and that's I think the more challenging one because you've crammed a whole bunch of functionalities into one at least in a browser context into one space. So, it does many things, so your browser has to be able to properly interpret what it's supposed to be doing with something.

So, if you make lazy assumptions about what a TLD will be, and you're going to try and create a picture instead of going to a website or vice versa, that becomes interesting. As far as dot zip is concerned, it's water under the bridge. It's already delegated. I don't know of any process at this point, especially since nobody's screaming and jumping up and down that the bridges are collapsing or what have you because of dot zip is out there. There's not really a practical thing to do. Looking backwards.

JONATHAN ZUCK: Question, maybe there isn't. But people are jumping up and down. It made the news and things like that. It sorts of one of the few times, something that we work on is actually being beamed into people's living rooms, frankly. So that felt like something worth discussing.

ROD RASMUSSEN: Yeah. And that's what I mean is the I don't see people from wearing government regulatory hats or law enforcement badges jumping up

and down. But there certainly can be plenty of hype about that, which does bring it to an awareness perspective. And I think potentially a challenge for ICANN org to answer questions, if they're getting them, and for us as a community answer as well, and thinking about these things. So that's about the extent of my knowledge.

So, I'm the guy in charge. Now, we do have some experts that have done some things on that. So, I would encourage and I told you ahead of time I was going to do this, folks who know things on the SSAC to say a couple words about potentialities and even if you work for a company that may have one of those things right now, it might be great to approach a microphone. Okay. And then there's oh, and we have…okay. Go ahead. Yes. A direct question was even easier.

BILL JOURIS: I'm not sure it counts as a question more as a comment. It sounds like what you're saying is we can't do anything until we have a disaster to motivate us.

ROD RASMUSSEN: Well, that's some dot zip's case, because we went through a whole process. It passed all the checks. It's live in the water. So, it's nobody's "dead" yet, which is a terrible, but that's literally I think it was the harm to human life. It was the bar that was set when it was approved. So, well, and I would call this analogous to name collisions. But, what mechanism, what level do you pull right now? we know technically how to do it. You pull it out of the root zone, but yeah, okay, great. Then

there's a whole bunch of lawyers that have something to say about that, right? So that's that part.

But going forward, down that that might be an interesting topic and if nothing else maybe an awareness piece. But having a better kind of idea of what where the concerns are and what we might want to say about it is kind of we're sitting here going, well, we kind of talked about this 10, 15 years ago, set that out there. But now that people are seeing it, maybe we need to take a look at what people are saying, which is great to have input from you guys because there's a lot more observation points, especially since we're the technical guys who know how this stuff works, in theory, what else is being said, so did you have some more on the bill?

BILL JOURIS: Maybe what we need to do is go have a little chat with ICANN's General Council and say, do you really need to have something in the contracts to say we can revoke at this delegation if this kind of problem arises. We realize you can't put that in now, but if you drafted the language, at least you'd be ready when, not if somebody dies or whatever, and they and you have to do something at least get yourself in position so you can move quickly when the disaster hits instead of going, oh, dear, we've got to wait until some outside law enforcement by or government agency comes to us and says, this changes whether you like it or not do it. We've been through that with [Audio Break 01:08:07].

JONATHAN ZUCK: -- technical, how this might be applied going forward for this next round? Peter, you had your hand up at one point.

ROD RASMUSSEN: Peter, you had one to do on some.

PETER THOMASON: So, Peter, I would like to put the added risk that that arises from the dot zip domain into perspective a little bit. So, I guess it actually is a problem if you click on a link and it ends in dot zip, and then you don't get what you expect. But I think that is already the case without dot zip domain, because you could have something dot com slash whatever else dot zip, and then that could be redirect to something else that gives you an executable file, for example. So that is not something that is new with the dot zip domain and in fact, with the dot zip domain, it would only be problematic if you are on the main page, because you if you have https something dot zip slash and then something else dot JPEG, then it plays no role at all that there is the zip thing.

So, what I'm saying is it only plays a role if the URL ends in dot zip, and that has been possible even before the dot zip domain because URL is not always ending domain names, they sometimes end in paths like directories and files. I think the more interesting abuse case is when you craft the URL in another way so that it is additionally misleading so that you also get the impression you are at a different domain name. So, before the domain name, can prepend an at sign, and then before that, you can prepend a username for logging in somewhere, and you can craft that to say git hub dot com or something, and then things

become more complicated. But that also has been possible before the dot zip domain.

You could do HTPS something get up dot com username at some other thing dot com slash something dot zip, and I think the risk of that, which is actually a risk, is the same as with the dot zip domain. I don't think it's any worse and so I'm not saying that that shouldn't be fixed, but I'm saying that I think we don't need to be too concerned about this particular TLD and instead it's a user interface issue where the various parameters that Rod mentioned are getting intertwined. And whether that can be resolved, I don't know. But I don't think the TLD particularly causes any harm, and I'm just speaking for myself.

ROD RASMUSSEN: Ask John Levine to pipe in a bit, and he's done some thinking on this. Right. Okay. All right. Okay. Justine okay. That would so, John, why don't you go ahead and then Yeah.

JOHN LEVINE: I'll be quick. I mean, I need to push back a fair amount on this. I mean, the basic mistake is Microsoft, who made a foolish decision decade ago to use the name of a file as a security indicator and you can create an executable file on any Windows machine that ends in dot M0V, and click on it in bad things will happen. And I'm with Peter, I mean, honestly, I think all the foo for all about zip and move, Mart tells us that it was a slow news day in the security business than anything particularly serious is going on. So, the issue of basically trying to do security by name was never worked, that horse left the barn decades ago. And

while I think it is reasonable to tell people like, hey, if you have particularly badly design software designed in Redmond, Washington.

This is yet another way that it might break. I don't think it's constructed to try to make a list of every possible filename because any string can be a file name and once that aren't file names today, it might be file names tomorrow or it might turn out to be file names in some part of the world that we don't happen to live in. So, I would like, note this and move along because I think we have way more important things to worry about.

JUSTINE CHEW: [Audio Break 01:12:03] I both actually. So just a short comment on the zip and then move thing. I think I agree with gentleman here that it's the shipping sail kind of thing. We should just monitor as a possible DNS Abuse channels and do it that way.

GREG SHATAN: Greg Shatan for the record. I think most of the cases that Peter mentioned are not the case that we're concerned about to the extent that anybody's concerned about anything, it's the fact that as a general business user, you see a handful of different types of file types that are coming as attachments in emails. You see dot doc and dot PPT. You see dot zip. And that's probably 99% of what you see. So thankfully, nobody's applied for dot PPT yet but that's the issue is whether you no one could exploit the that confusion and get somebody to click on what they think is a zip file of compressed documents and instead get sent to a dark space on the in the dot zip domain. All the other examples I think

are irrelevant because they're not the concern. So, I agree with Peter that those are irrelevant, but they never were even on the table.

PATRIK FELTZ: Patrick Feltz from SSAC. To be able to have a constructive discussion, we need to separate the type of file and the name of file, the two completely different things. If it is the case that you happen to use an application that cannot separate the type of file from the name of file, in that case, the problem is definitely not on the TLD side. It is the application that you have decided to choose. So, this is back to procurement and the choice of applications and requirements on application developers. Thank you.

GREG SHATAN: You're just never going to stop people from sending zip files.

WARREN KUMARI: Warren Kamari SSAC. I mean, sort of following on from what Patrik said, what shows up in your mail reader saying, this is blah blah blah dot zip. The fact that that shows up in the mail reader doesn't actually mean that it's a zip file. You can say, here is a picture of my cat and that can be what shows up as the link. But the TOD that it's actually sending it to could be I like cheese dot net. So, having anytime people have any sort of assumption that what shows up in the link or in the sort of body of the message, is the same type as the file is going to be a problem. The same thing happens, you can get an email saying, please go to foo dot net and when you click on it, it takes you to evil dot com. We've been

teaching users for a long time. These are different things. What shows up here is not the same thing that works there.

JONATHAN ZUCK: Right. People can hover their mouse over the URL and see the real URL. But I think we were also saying it's possible to construct something that looks like a real URL well.

WARREN KUMARI: Well, yes and no because the examples that have been shown with URLs with dot zip they actually aren't ascii text. It is a homonym attack. It's got a Unicode slash.

JONATHAN ZUCK: Right, but URI, it's going to look like a URL even when I do what I've been trained to do.

WARREN KUMARI: But that's exactly the same thing for Microsoft dot com. With the o being a Cyrillic O. It's this is a Unicode confusion issue, and that's why, pretty much all of the mail clients and web browsers when you hover over the text show you, it in puny code or similar. Thank you.

ROD RASMUSSEN: We are out of time, so we need to get to a rap here. But as you could see, interesting stuff, and we had similar discussions internally, we're moving towards conclusion. We had one more, do we have time for that, but for number one, not really. Apologies. I'd be happy to talk offline.

JUSTINE CHEW: Could you point me to the IDN experts in SSAC? Maybe I can have a conversation.

ROD RASMUSSEN: IDN experts in SSAC who are in the room. Please wave your hands around. Yep. You got two right next to you. And Patrick and Yeah.

JUSTINE CHEW: So, If I could just borrow you guys for a short 5 minutes or so, please. Thank you.

ROD RASMUSSEN: Thank you. Great. We worked that out.

JOANNA KULESZA: Let Jonathan jump in for something:

JONATHAN ZUCK: Thanks for coming. We always appreciate it, and we look forward to working with you on plotting out an informational campaign on DNSSEC.

ROD RASMUSSEN: Awesome. Thanks for having us.

**[END OF TRANSCRIPTION]**