ICANN79 | CF – Joint Session: ICANN Board and SSAC
Wednesday, March 6, 2024 – 9:00 to 10:00 SJU

FRANCO CARRASCO: Hello, my name is Franco Carrasco, and I am the participation manager for this session. Welcome to the joint session between the ICANN Board and the Security and Stability Advisory Committee. Please note that this session is being recorded and follows the ICANN expected standards of behavior. Interpretation for the session today will include English, French, Spanish, Chinese, Russian, and Arabic. Please remember to state your name for the record and the language that you would speak if speaking a language other than English. Please also speak at a reasonable pace for our interpreters.

The discussion is between the ICANN Board and the SSAC. Therefore, we will not be taking questions from the audience today. However, everyone is welcome to use the chat. Please note that private chats are only possible among panelists in a Zoom room format. Any message sent by a panelist or a standard attendee to another standard attendee will also be seen by the session's host, co-host, and other panelists. Having said all this, I will now hand the floor over to the ICANN Board Chair, Tripti Sinha.

TRIPTI SINHA: Thank you, Franco. Welcome everyone to the joint meeting between the SSAC and the Board. We get this opportunity twice a year and it's a

moment when we can have a very candid conversation on various and sundry topics. I believe we've exchanged some questions that will be the focus of our discussion today, so with those opening comments, I turn it over to Jim, who will chair this from the Board.

JAMES GALVIN: Thank you, Tripti. Welcome everyone. This is one of my favorite meetings of this week, bringing the Board and the SSAC together to my favorite activities during the ICANN week with so much interesting stuff going on. But for today, our focus will be talking about the Name Collision Analysis Project in general and the aspect of it that is of interest to members in the room. This project has a long history. That history has been well documented in many briefings across the way, so I'm not going to dig into that at all.

For today, what's important is we do have a first draft of the final report has been out. Public comments have opened and closed, and the discussion group is actively reviewing those comments to produce a final report. SSAC has led the NCAP project throughout its lifetime. Alongside the community, of course. That was one of the requirements from the Board when we were first asked. SSAC was first asked to consider this work and to move forward with it. SSAC is actively considering the essential, if not critical, security and stability principles and the risk factors of the discussion group proposal for an updated name collision framework, which, of course, many in the community are very much interested in and its relationship to the new gTLD program coming up.

So, we'll begin with the NCAP co-chairs making a presentation about a summary of some key takeaways from the NCAP discussion group, and then that'll be followed by an SSAC's presentation about their current thinking about the NCAP discussion group proposal. So please ask questions at the end of each proposal at the end of each presentation. We'll do that. A quick comment about logistics, and I'll repeat this later. We will have roaming mics. I know there's not as much room up here on stage as there is down there in the audience, but all SSAC members and all Board members are welcome to speak at any time. Please just stand up or wave, and someone will bring a microphone to you. And with that, let me turn this over to Ram Mohan, Chair of the SSAC.

RAM MOHAN: Thank you, and thanks, Tripti and Danko and all the other friends here on the Board. It's lovely to be here with you. We are looking forward to a candid and direct conversation with you about our observations and findings from the Name Collision Analysis Project. That's our sole topic for discussion. The SSAC, prior to this meeting, if you go to the previous slide, prior to this meeting, the Board and the SSAC exchanged this particular conversation where you had wanted to understand the SSAC's point of view on this topic.

The way we're looking at this, we have two presentations to you. The first part is going to be led by Matt and by Suzanne, who've been the chairs of the SSAC-led community-wide discussion group that has been looking at name collision. So, that's the group that has published a draft final report with public comments that have come in. So, they will provide you, the Board, with what they have found and the key

takeaways and the recommendations. And then we will break for questions and answers and a discussion between all of us.

Once that part is done, we will go to a second part of this meeting, which focuses on the SSAC itself, which has formed a work party to discuss name collisions and what its findings are and kind of where we are in that process. You will find that that part of the conversation perhaps may be more pointed in its conversations and in some of the findings as compared to the other one. So, without further ado, I'll hand it off to you, Suzanne and Matt.

SUZANNE WOOLF: Thank you, Ram. Thank you, Jim. Thank you, Tripti. And we are here. I'm Suzanne Woolf, and I'm here with my co-chair, Matt Thomas, to review the work of the Name Collision Assessment Project Study 2. As Ram noted, the draft report was published a few weeks ago and the public comment period ended earlier this week.

Before we really get underway, I want to note this work was produced by an awful amount of effort by the Cross-Community NCAP Discussion Group, and we wanted to thank those participants for that effort. People really stepped up. We also note that we tried to work by consensus as much as we possibly could. So, the contents here are the group's best assessment of the issues and the best guidance we can offer the Board and the community.

So just a quick recap, because as Ram said, we've heard of these things before. If you think of domain names as being like house addresses, the drawbacks of having duplication are very clear. If two houses share the

same address, it's hard to know which one mail or delivery should go to, and furthermore, if the mail or the packages end up in the wrong hands, there are many, many consequences that could be quite serious. In the DNS, name collisions occur when a domain used in the global DNS namespace is also used in a different namespace. Most commonly, this is your corporate IT, but there are also lots of other ways it can happen. And where it's used in a different namespace than expected, users or software or other functions might misinterpret.

And why is it relevant for at this time and for new gTLDs? There's significant risk of unintended consequences if we don't undertake this work. Businesses have used labels as internal TLDs in private namespaces. Those do turn up in the global internet and have side effects. Introduction of new gTLDs increases the probability of name collisions just by having a larger pool of potential names that could lead to collisions. More possibility that a gTLD string might overlap with names already used in private networks or internal naming systems or some of the other ways people work on the internet.

It's also important to note that measuring name collisions is difficult and more difficult than it was in the past due to the evolution of technology and network infrastructure. It's the internet. It keeps changing. It keeps evolving. And so, what we have to do in order to collect the data we need to assess what's going on is different and, in some ways, more complex and more difficult to get. Particularly privacy enhancements in the DNS and alternative naming systems have made the DNS landscape more complex and measurements more difficult. I'll hand over to Matt.

MATT THOMAS: Thank you, Suzanne. This is Matt Thomas for the record. So as stated earlier, the main mission from the NCAP discussion group, which was led by the community as well as SSAC, was commissioned by the ICANN Board to give some advice on how to proceed with the next round of gTLD delegations. And this work was really aimed to create a sustainable, repeatable and deterministic process for how to address name collisions going forward.

One of the questions that the Board asked SSAC to do and present data and analysis on was specific advice around corp, home and mail. We also received a series of questions. I think it was nine different questions to provide general advice around name collisions going forward. And all of this work in the NCAP's discussion group was done in an inclusive manner with the community. Experts in various different fields joined the group and provided their advice and insights informing the Study 2 report that's out for comment, or that just closed for comment.

So, there are three specific studies that were commissioned as part of the Study 2 report. The first study was really looking at collision strings, and that report looked longitudinally at multiple years of root data from ANJ root servers looking at corp, home and mail, as well as lan, local and internal. And that study really provided some interesting insights because it showed us that the impact of name collisions over time has still increased. And that was determined by looking at something called Critical Diagnostic Measurements. And these CDMs are a way to

quantitatively look at name collision telemetry to potentially gauge or measure the impact of name collisions for a particular string.

Now, that report was very insightful because it provided some of the guidance and the recommendations on how to measure and look at name collisions going forward. But the second study, a prospective study of looking at DNS queries for non-existent top-level domains, also provided some insights or guardrails in terms of how name collision data could be assessed and where it could be measured in the DNS ecosystem. So as noted before, the ecosystem for the DNS over the last decade has significantly changed.

We have new technology changes like QNAME minimization, aggressive NSEC caching, deployment of local roots, and the deployment of large public recursive resolvers that have happened since the 2012 round. So, this study really took a look at how the DNS data could be analysed from one or multiple root server instances versus that of what could be collected at, say, recursive resolvers. And this study really tried to provide some insights and guidance in terms of the pros and cons of what can or cannot be measured for name collisions in different perspectives.

Finally, the third report was looking at the root cause analysis. And this really focused on the 47 reports that ICANN received since 2012, stating that name collision issues were actually observed in the wild from the 2012 delegation of gTLD strings. And so, that report really looked into understanding how big or impactful were those particular collisions and what the underlying root causes of those were. So, this

really helped to provide some concrete data in terms of real-life experience with name collisions from the last round.

But all of this together, along with the discussions that the group had over numerous years, really framed the key finding of name collisions going forward needs to be treated in a risk management perspective. And that name collisions are here, and they will persist and continue to go forward. And so, in order to preserve the DNS security and stability, we're recommending that risk management frameworks be applied to name collision assessments.

So, there are several key findings that are coming out of the discussion group Study 2 report. And one of these things, like we mentioned before, was that it needs to be treated as a risk management problem. But in order to do that, that risk management problem needs to have data to be able to decide are particular strings potentially harmful or are vulnerable to some kind of name collision risk.

In order to determine that, the discussion group is recommending the creation of a Technical Review Team , or the TRT, who is a group of experts who are going to be able to look at the name collision telemetry data and do an analysis of those critical diagnostic measurements to make a recommendation on the applied for strings. But in order to do that, this requires data.

So, one of the things, as we've mentioned before, is that the DNS has changed over this last decade due to technology changes. And that the data used in 2012 just simply is not available or is highly impaired now. And that new data needs to be collected in an alternative manner. So, there are some enhancements into how that data is assessed and

collected, which we'll talk about a little bit more going forward, via these multiple different assessment methods in which the discussion group identified four methods for collecting and analysing that data. But overall, there are two main goals in the side of the framework that's being proposed in the name collision report. And the first is to ensure that the name collisions can be assessed. And that is to make sure that the data is available.

And in order to make that data available, we are now proposing from the discussion group that the applied for strings are temporarily delegated into the root zone prior to granting to facilitate data collection to the Technical Review Team. And this will be able to enable the TRT to use that data to define and collect their risk assessments. The second goal is to provide ICANN with access to evaluation of mitigation and remediation plans. And so, while known causes may inform mitigation and remediation plans, there needs to be the ability to create tailored or bespoke mitigation plans for certain strings.

So, it's a relatively straightforward process that the discussion group is putting forward for name collisions. And it starts at stage zero or pre-application for an applicant. What here we're encouraging in the framework is that applicants a priori submitting their application use publicly available data such as ICANN's ITHI or ICANN's magnitude data, for instance, to see if their applied for string already has potential name collision issues. Now, certainly the presence or absence of that string on any list is not a guarantee that it isn't void or has severe name collisions, but it at least gives the applicant some insights early into the application process before proceeding with their applied for string.

Once the application is submitted, this Technical Review Team will also do their assessment of publicly available data to determine the risk of that name collision. It is intended that the Technical Review Team for the most part will always recommend that that string then goes into a temporary delegation into the root zone, but their review of that publicly available data is there to assess high-impact strings to see before it goes into the root to ensure the security of the root server system. Now, if the string is high risk, the TRT is going to submit a recommendation to ICANN Board, or the applicant may need to provide their own mitigation and remediation plans in conjunction with the TRT.

Now, proceeding on to stage two, once the applied for string is actually delegated into the root, this enables the data collection that we mentioned earlier. This allows the TRT to collect the DNS telemetry data and do their assessment. This can be done via one of the four methods that were described earlier in the discussion group. The first of which is called No Interruption, and I won't go deep into the technical details here. They're in the Study 2 report, and for your information, but no interruption you can simply think is a wildcard empty zone that will preserve a no error, no data response, similar to what the non-existent domain response was being received at the root. So, the impact to the string should hopefully be minimal, but the data collected is very useful in this.

The second of which methods could be employed is Controlled Interruption, similar to what was used in 2012. And here controlled interruption again facilitates the collection of that DNS data, but it also performs the function of providing some notification mechanism to the

impacted parties so that they could detect that they're experiencing potentially name collision issues.

Now, the third and the fourth methods being proposed are called Visible Interruption and Visible Interruption and Notification. And these methods are a more advanced way in which additional, more granular data can be gathered for the name collisions. Because in both of these scenarios, the zone for the temporary delegation now contains a routable IP address that will direct all queries to a sinkhole server that allows now not the collection of DNS data from just the recursive resolvers, but actual connection attempts to the sinkhole server from the impacted parties. So, you actually get information from the devices themselves that are causing the name collisions. All of that together, then the TRT will take that data and they will perform their risk management assessment and then provide a recommendation to the Board to proceed or to not.

So clearly the Technical Review Team here has a large role in the framework going forward. And to that end, the TRT needs to have some qualifications. They clearly need to have deep experience in the DNS in terms of both management, as well as how it works, and how to analyze that data, which supports they need to have a strong understanding of the general internet infrastructure, as well as taking that data that's collected and providing a trend analysis and predictive modeling in terms of the risk assessment based off of that.

The TRT basically has four roles, the first of which is assessing the visibility of name collisions, then documenting the data that's collected, their findings, and the recommendations. They must also

provide ability to assess mitigation and remediation plans, if needed, for particular strings, and then also advise in any type of emergency response function, if needed. So, Suzanne, do you want to talk on this one?

SUZANNE WOOLF: Sure. Except I think you already went through it. No, I haven't. I haven't. Thanks. So, we have had many conversations about data collection, and not only the techniques, but the principles. And one of the things that we need to make sure is really clear, as I said a couple minutes ago, you can't simply reuse the collision detection methods used in 2012.

It is unfortunate, but true, that there's really very-- well, there is less data available in the sense that, for instance, QNAME minimization is a technical term for a change to a way of structuring DNS queries where you don't send the full domain name. You don't send the fully qualified domain name to the root servers, and the data available to the root server does not include the complete domain name. That makes it challenging to find out certain things about how people are using those names, or what the unexpected consequences might be.

There's also, well, probably a smaller factor, but still relevant. There's an awful lot more of IPv6 in use now than there was in 2012. So, it's probably relevant that controlled interruption, as implemented in the last round doesn't work as well for IPv6. It's more complicated, and the data is fuzzier. There's also many, many changes on the regulatory side having to do with operations, privacy, and so on and so forth.

But it's also true that to seriously analyze name collisions, you need a variety of sources. It turns out, to our disappointment, that it's impossible to build a generalized case for root causes. There was a Study 3 that was originally scoped to provide further analysis on a model or a generalized description of root causes, and we discovered that there is a little too much individuality. There's no one-size-fits-all. There's not even a one-size-fits-most.

And we'll note that impact assessment could require large amounts of data over significant periods of time. ICANN Org has expressed concerns about risks to privacy and confidentiality with some of the proposed data collection methods. These concerns do need to be thoroughly understood and addressed as the DG recommendations move towards further discussion and implementation. Actually, this is a good time to point out that we got a number of comments as part of the public comment, including the one from Org. We greatly appreciate all of them. They all were substantive, and the discussion group is working our way through them. We can turn it over to-- Do you want questions?

JAMES GALVIN: Go back to the previous slide, and I will take it from here. So, we'll take this opportunity to take questions about the key takeaways from Matt and Suzanne about the presentation from anyone. As a reminder, there is an opportunity for all SSAC members and all Board members to speak. There are two roaming mics. They'll stand up over here and just watch. If you want to speak, please just raise your hand, and we'll call

on you. I'm going to use moderator's prerogative and ask the first question.

There's one detail in the study, the perspective study and the corp, home and mail study, that I think is interesting to call out. As we all know, corp, home and mail are currently in a deferred indefinitely status with respect to ICANN. Something interesting turned up with respect to .mail , though, in all of this, the CDM values for .mail . And I wondered if you could speak out, not the actual values, but the relative values of .mail versus there were seven, I believe, other TLDs over the course of the 10 years that have been delegated. And I wonder if you could just speak a bit to that about just the technical details. Thanks.

MATT THOMAS: Yeah, so I'll refer everyone to the actual report for the details. It's in the appendix of the Study 2 report. But what Jim's asking about is the details of the CDMs, or those critical diagnostic measurements that we talked about earlier for .mail . Looking at multiple recent years of route traffic for corp, home and mail, mail exhibited a very different, lower set of critical diagnostic measurements relative to corp and home. But when compared to some of the previous strings that were delegated in 2012 round, there were actually seven strings with higher CDMs than .mail that were actually delegated. And some of those strings were the recipient of some of the root cause reports, or the name collision reports that were sent to ICANN in the 47 reports that we talked about earlier.

JAMES GALVIN: Thank you. I know we have a question from the floor from Sajid. And a reminder to our remote SSAC members, please just raise your hand if you'd like to speak or put a note in the chat. Thanks.

SAJID RAHMAN: Thanks, Jim. I was wondering, has there been any impact assessment of including TRT on new gTLD in terms of cost and timeline?

SUZANNE WOOLF: I didn't quite get the question.

RAM MOHAN: The question was, has there been any impact assessment of all of this in terms of cost and timeline?

SUZANNE WOOLF: I believe that's part of the next step in the process now that the discussion group has presented their findings and recommendations, and that's to be undertaken next, I believe.

RAM MOHAN: And just briefly, Sajid, ICANN Org has, in its public comments, provided some input on that, and the discussion group has just started to look at the cost and the timeline feedback that has come from ICANN Org.

MATT THOMAS: I'd just like to add one more comment on that and make a comparison to 2012 round when name collisions were essentially bolted on at the end of the process because name collision issues kind of were raised after the whole application period had closed and finished. Going into this next round, clearly, we have a green field opportunity to consider when the timing and name collision risk assessments are done. They don't need to be a bolt-on at the end.

So, there's opportunities to put that review risks assessment process earlier into the application process and run it in parallel with other processes of the application. So, name collision risk assessment is not a gating factor going forward. It can be done in a more streamlined manner with the rest of the overall process.

JAMES GALVIN: Thank you, Matt. Sarah?

SARAH DEUTSCH: Yes, hi. I was working in Verizon in 2012 when the first collision analysis was done, and we had the ability, I guess, to see which of our customers at that time could be impacted. Tried to call them, and it was an impossible situation, A, to find the right person to explain what this issue was, and to get them to do something. So, I'm just wondering, this time around it sounds even harder. What is the impact on people inside businesses who may not know that this collision in their system exists?

SUZANNE WOOLF: Yeah, we actually spent quite a lot of time discussing what kind of notification and outreach could be done within reasonable constraints of time and cost and operational and legal and all of those concerns. And it is quite difficult. We do understand the situation a lot better now than back then, so I do expect that that will make it a little bit easier to figure out what to do with these findings when we have them. But this is why this is a serious problem that we're spending all this time on, because finding out who is impacted.

The other thing that is in the report is a recommendation that, since it's very difficult to get the relevant data out of-- It's very difficult to get the origin or the root cause of a collision out of the data that is ever going to be available to us, so we also said that the education and outreach that ICANN Org has done needs to be continued, just because there's never going to be an in-band way of making people aware of the problem.

MATT THOMAS: I'll also just note, as Suzanne mentioned earlier, that each name collision string is a little bit unique in itself, that there's no general remediation plan that can be applied for everything. We've seen that certain ones can be remediated through outreach and notification more easily than others. We have blog posts and documents inside the Study 2 report showing how numerous different strings have been remediated in the last several years, some of them with Warren over there at Google, but some of the other ones take a significant more amount of time and effort, because they might be, the root cause might actually be embedded in hardware that's home routers or

something, so you have a very long tail of trying to fix that kind of problem. So, everyone's a little bit bespoke on that.

JAMES GALVIN: So, I've got two queues going here. Questions from Danko and Edmund, but I have two fingers from Warren and then Barry. So, please go ahead, Warren.

WARREN KUMARI: Yeah. Just to expand slightly on what Matt said, it depends kind of on the type of the collision. In some cases, it's a single organization that is using it within their network, and other times it's built into hardware. For some of the ones where it's within a single network, Wes Hardaker, actually no, Matt has done some good outreach on that. And Wes did some good outreach when a colliding string was built into a piece of software on Android, and that fixed some things fairly well.

The fact that it is, as you say, hard to reach people is why we have suggested visible interruptions this next time. So instead of people having silent breakage or seeing an IP address and they don't understand what that means, we're suggesting it should be a web page or similar, so people actually see it and understand the impact to themselves. That obviously only happens once the network breaks, though.

BARRY LEIBA: This is Barry Leiba. I'll also add that the report is going to be telling what needs to be done, what sequencing needs to happen, and all of that sort

of stuff. The costing out of this thing and deciding how to communicate with applicants and all of that stuff is part of the implementation that's out of scope for this report.

JAMES GALVIN: Thank you. One more two-finger, and then I do want to move on to other questions. Go ahead, Ram.

RAM MOHAN: Thank you. Sarah, the other thing is that there might be a perception that this is relevant for open strings, open generic names, but it's equally applicable for brands as well. So, the collusion is not something that is unique only to generic strings.

JAMES GALVIN: Okay, thank you. Danko, you have a new question?

DANKO JEVTOVIC: Well, if it's okay with you, I have a couple of clarifying questions so if we can do it maybe interactively. So, first of all, many thanks to SSAC for doing this really important studies and bringing it to the Board and of course, the importance of stability for the next round is huge. But I was trying to put my mind around all of that and try to simplify it in my-- Well, I used to be techie, but my brain is not as it was with all this governance stuff.

So, first of all, I understand that there are kind of two problems. One is how to get the data that is usable, and that data also provides

risks because that can be privately identifiable data it. And the other thing is how to make a risk assessment based on that data. And the proposal for a risk assessment is this technical expert, Technical Review Team. So how do you envision that? That's my first question.

Because I see those that can be two or three roles that that team might fit or someone else might do. One is to kind of create a set of criteria that will be applied. The other role is to look into data and apply this criteria to data. And the third role would be actually to make a decision should that delegation be proceed because I understand you want and we want to put it, front load it into the process. So how do you see that Technical Review Team and how you think it should be created and what's maybe the role for the ICANN Org there? Bit of clarity on that. Thank you.

SUZANNE WOOLF: Thank you. I like that question because it sort of gets at why we have looked at the set of issues as risk-based rather than as an engineering problem because there are unknowns and there are decision points in what you just described. We do believe from experience and from discussion that it will be relatively straightforward for Org to figure out how to populate the Technical Review Team. There are other reviews that are part of the application process and various other things that the community does. And we didn't want to go beyond our remit to over-specify, but we do know in this community how to do expert groups or expert panels and like that.

As far as how to use the data that's obtainable, this is again a combination of the quantitative, the CDMs that we believe tell a fair

amount about any name collision situation. And there's some discussion in there of the more, frankly, qualitative. And one of the things that does come out as kind of a meta-observation, if you will, is that the quantitative was never going to be enough by itself.

DANKO JEVTOVIC: Thank you. That really helps. But kind of my thinking there was, obviously, we know how to create community-based groups and the expertise needed is clear. But my thinking was, is there a part there that can be actually kind of looking at more technical roles or that can, or do it or outsource it or whatever? And the other part that is like to create criteria and to make decisions to be part of the group. I'm a bit worried that if we task this Technical Review Team to actually collect data, look at the data, it's too much.

MATT THOMAS: So, within the report, there are some opportunities to describe that some of those roles could be operated by different entities. So, the data collection could be performed by a different entity who operates the environment and maybe the Technical Review Team only does the analysis and recommendation processes. So, you could have a separation of duties in some of those there.

DANKO JEVTOVIC: Okay. So, I wasn't totally off. And sorry for taking so much time. I have a second question that's shorter. So, to get to the data, I understand you want to put something into the root zone and then based on that to

get it. So, in the previous round, if I'm right, that task was for the registry because it was at the end of the process. And now in the beginning, there is no contract on anything. So, I understand and envisioned that should be actually performed by the ICANN Org. And that has a specific implication also on privacy and risks and finance. Is that right?

RAM MOHAN: Danko, yes, that is correct. And in the second part of this conversation, I think we will narrow down on that piece and our views on that.

JAMES GALVIN: Okay. I do want to give Edmund up there. Wes, if you don't mind, if we can wait and hold your question till the next time. So, Edmund, close the queue so we can then move on to the next presentation.

EDMUND CHUNG: Yeah. Edmund here. So, I find it interesting, I guess, Jim, you started off with the .mail example and then we went on to a little bit of a discussion about the remedial processes and what next. But the answer was that I find interesting is the, if you're saying that .mail , the CDM value is kind of lower than some of those that were actually delegated. I guess, not to say that not to pre-empt any TRT review on it, but what does that mean? Can we walk through some of the potential remedial things? What does that mean? Does it mean that at some point, .mail can be released or what are the things that might need to be looked at, those kind of things?

JAMES GALVIN: Looks like Warren wants to take that, if you don't mind.

MATT THOMAS: So let me qualify the previous statement around .mail . We talked about the CDMs being lower. That's the quantitative component. There is also a qualitative component to actually understand how .mail is being used and understanding how it is potentially a name collision. So, while the values might be lower, without understanding how that string is used and what potential security threats or risk there are with it from a qualitative view, just having lower numbers doesn't necessarily mean it's good to go. So, you have to paint the whole picture from both sides on that to actually get a real picture, which requires the bespoke analysis on a string-by-string basis that you need both the quantitative as well as the qualitative components.

WARREN KUMARI: And yeah, just to follow on from that. When .mail was indefinitely deferred, there was even a note that this has different semantic meaning or it is not being deferred because it has different meanings and different usage. And I think it actually makes a really good example of why we can't just look at the big list on the magnitude data and say, anything above this is where we're cutting it off. Different strings have different uses. Sadly, semantic meaning is sort of built into names. And so, you can't just look at numbers and say, this will be fine, or this won't. You have to have a larger holistic view.

JAMES GALVIN: Okay. Thank you very much. Excellent discussion Let's move on to the next slide and I'll turn this over to you, Ram.

RAM MOHAN: Thank you. If you could please move to the next slide. So, as I mentioned earlier, the SSAC, as a follow through from the discussion groups report, has created its own internal work party with its internal experts to look at the name collision report and to provide specific guidance and advice to the Board. So far in our analysis, we've found strong agreement with the recommendations and the observations of the discussion group but there are a couple of things for you the Board and also folks in the community to understand.

The first thing is that measurements at this scale are very hard, very difficult. And in the analysis that has gone through so far, as Suzanne and Matt had mentioned, the data is coming from the A and J root service only. And that used to be quite a lot before in 2012. It is actually not enough right now. And as a result, what we are finding is that the advice that we give you cannot be based on just a quantitative measure.

So, you will hear us say there is a persistent risk, the risks are higher. But if you ask us how much higher? Is it 10% more 15% more? We can't give you that answer at this point because the total amount of data that is coming through is far lower than what was available before. So, it's not as quantitative as we would like it to be. So, our advice to you is based on some qualitative measures, perhaps more on a qualitative basis than purely quantitative.

And if you take it all the way to the end, you could even go so far as to say that the advice that we're giving you is going to be based on the expertise and the experience of the people who are gathered in the room to discuss this. And Some could sit there and say there is quite an element of guesswork that is involved. And they would not be completely wrong in saying that but you should know that there's qualitative information and strong beliefs that drive our recommendations.

So, so far inside of the SSAC working group, we have some conclusions. We have not finalized them, so this is subject to change, but we wanted to share them as we are developing them. First name collisions continue to pose a persistent threat to the security and stability of the DNS. We have no doubt about that. But as has been mentioned before, we think to actually get a better handle on it, it's crucial to collect the data and it's also crucial for you the Board to determine when to collect the data and where in the evaluation process to collect it.

So, in 2012, the ICANN framework for name collisions, it tagged the assessment and the mitigation at the end of the process. So, the name was granted to the applicant and at that point, name collision happened, and effectively the responsibilities and the risks were transferred to the TLD registries at that time. But as Suzanne has mentioned, from then to now, there have been quite a lot of changes in the amount of data that comes through to root zone resolvers, to recursive resolvers, and even controlled interruption the way it was implemented in 2012 is not fit for purpose right now. It won't work as it was implemented because perhaps about 40% of traffic now is IPv6.

So, as a result, the SSAC is likely to come to you and say you should make name collisions visible. You should collect the data right at the beginning and provide that data to the Technical Review Team for them to review. The points are valid that you have to figure out are there criteria created ahead of time as a generalized framework for evaluation that the Technical Review Team will apply in a standard way, or is it done in a bespoke manner. Those are questions to be answered down the road, but they are valid questions. But we think that you have a great opportunity to not take the risks after the fact after granting the TLD to the applicant.

One other thing that we want to point out-- And we acknowledge also that there are privacy risks, there are privacy implications that attach to our proposed model of visible interruption or a visible interruption with notification. Well, we want to point out two things. If you were to do Visible interruption or even controlled interruption at the end of the process and at the time and it is granted, even in controlled interruption IP addresses come through and, in some jurisdictions, IP addresses are considered PII. So, you don't really walk away from the privacy implications problem. So, our thought is look at it early on, address it early on.

If you apply an interruption framework at the end of the evaluation process and and have the TLD registry to actually do that work effectively, you're transferring the costs and the risks to the successful applicant. And that's a decision you can make. Our concern and our discussion so far is not about that. Our concern is that if you do that, you're now creating a disparate evaluation mechanism. Some registries may be better equipped to perform such tests than others.

The privacy implications in some ways you will end up amplifying it because you're spreading that to a plethora of TLD registries rather than centralizing that at the start of the process.

And we think that given that there are privacy implications and given that there are risks, our assessment so far is that on balance, perhaps there is lesser risk for the assessment to be done by ICANN Org where you can potentially contain or at least have the impact of those privacy risks more understood in a more centralized manner rather than distributing it to all of the TLD applicants. So, that's kind of where we're going. That's likely to be the thrust of our final recommendations to you.

So really, what we're saying to you is when you are choosing mitigation and notification methods, we suggest to you quite strongly your primary focus should remain on ensuring a secure and stable DNS infrastructure. Of course, it aligns with your mission, but what we're also saying is unless there are compelling privacy concerns that just cannot be mitigated, you ought to prioritize finding solutions for the privacy concerns and not compromise on what we believe are the critical security and stability aspects of name collision.

So, that's the direction that SSAC is likely to provide. As I mentioned earlier, this is not where the discussion group landed at, but inside the SSAC, we think we have a responsibility to provide you very clear security and stability-oriented advice. And we think that you have a big role ahead of you, but we need to share with you that just looking at it from privacy alone is not going to be sufficient. We know you already understand that, but our exhortation to you is if there are ways

to mitigate the privacy issues, you ought to do that, especially because we don't actually see any mechanism by which you can walk away from the privacy implications.

Final slide. We're still on track. You asked us to deliver a report and our recommendations to you by the 1st of May. We're on track to do that, assuming that nothing blows up inside of the work party. But we're on track and we look forward to giving you kind of formal final advice at that time.

JAMES GALVIN: Okay. Thank you, Ram. As a reminder to those in the room, we'll take questions, but this is just a discussion between the SSAC and the Board. And again, anyone from the audience in the SSAC and Board who want to speak, please just get my attention and I'll add you to the list. Wes, we'll go to you first.

WES HARDKER: Thank you, Jim, and thank you, Ram. And thank you, Suzanne. Actually, the question that I had earlier is relevant to both of these two discussions. You've both talked about the changing environment and the fact that things are changing and it's getting harder to actually get qualitative data that can be used in any sort of fashion. In fact, it's changed since 2012 and my questions really centered around, I'm sure it's going to change by 2036.

And so, in fact, Ram, you said you should make name collisions visible. But even the changing infrastructure of the Internet may make that

harder and harder over time. There's more and more traffic that goes machine to machine. A lot of things fail silently, and you actually don't notice. So, I guess in the end there is no perfect technical solution to any of the problems that you've just listed. And what I'm wondering is, did the study take into account the fact that you may be asking a TRT team to do something as information gets better or possibly worse over time? And did you consider how are you going to deal with this when there's ten years out or something like that? And then information and the environment has significantly changed.

JAMES GALVIN: Warren, did you want to answer or two fingers? Okay, please go ahead.

WARREN KUMARI: I mean, there is definitely a tension here between privacy and the ability to collect data, kind of by definition, right? As we've seen, the DNS is becoming more privacy-centric over time. So, a lot of the information simply isn't visible. This means that as time goes by, it might be more and more important that you have some sort of way to clearly notify somebody when the system, when the name is actually delegated. We might not be able to predict as well in the future that a specific string is colliding, and we might just have to try it and provide a way to let them know that this is happening, for example, with visible interruptions.

JAMES GALVIN: Okay. Thank you. I've got a queue going here and I do want to do a time check with folks. I have Maarten and John. Tripti, you want a two-finger? Please go ahead.

TRIPTI SINHA: Yeah, the other thing, as a partial answer to Wes, we did take into account that changes would continue. We were not at a static point, that none of this is static. For example, we were asked about the feasibility and possible mechanisms for creating a do-not apply list. Is this feasible? Is this something that should be done? And we said no, because it would change. It's more dynamic than people think and it's more dynamic than is really practical with the other constraints.

JAMES GALVIN: Thank you. Maarten? Okay, John, go ahead.

JOHN LEVINE: Yeah. I'm John Levine. I'd like to put a slightly different spin on what Ram said. I mean, because the end of this process, we're going to do all the technical stuff and then it's going to be dropped in the Board's lap. It's going to have to make policy decisions. And it occurs to me that the issue fundamentally, every time you delegate a new TLD, is that you are privatizing the benefits and you are socializing the costs. And the privatizing can be quite narrow. I just counted the median number of names in a new TLD is under 400. Then when you delegate, you were probably delegating to an extraordinarily small group of people and as

we've just heard, the potential people that might affect could be anywhere.

And this is, to me, this calls out for a cost-benefit analysis. And when the SSAC last met, we talked to the economist we recently hired, and he seemed really good. So, I'm thinking that to inform this and to inform your decision, this really cries out for a cost-benefit analysis, trying to figure out what are the overall benefits that people are going to receive? What are the overall costs that you're going to inflict both on this group and everybody else and allow that to make an informed decision of is this actually in the overall benefit of the internet that ICANN presumably is supposed to work for? Or is it just, are the expenses out of line for what you would get?

JAMES GALVIN: So, I'm going to take that as a comment on behalf of the Board for the Board to consider and take back rather than answering that now and go to Maarten.

MAARTEN BOTTERMAN: Maarten Botterman for the record. A slightly different question. I really, really appreciate all the work SSAC does and has done over the years to help us get where we are. The NCAP work is important in that. As you know, we're currently embarking on the next strategic plan. RAM was on the Board when we developed the current one. My question is, seeing all the new emerging technologies coming, do you think the SSAC membership is able to deal with that for the next five years to

come or do you see and do you already talk and consider how the membership should evolve on that crypto, AI, etc?

RAM MOHAN: Thanks, Maarten. Great question. And I would recommend folks to look at, the SSAC has published a report on this just recently, the evolution of the DNS namespace, where we consider these new technologies that are evolving. Barry led that working group, so chat with him offline. At the end of the day, the summary of the recommendations are track the space, watch the space. We don't actually see that any of these necessarily pose an existential threat to the DNS. The namespace is evolving, and innovation is a good thing. So, that's the general outcome of that.

JAMES GALVIN: Okay, thank you. I've got Tripti and then Becky, and I'm probably going to have to draw a line for time. But go ahead, Tripti, please.

TRIPTI SINHA: Thank you, Jim, and thank you to the SSAC for all the work on this. Very good work. I think, Jim, during your opening comments, you said something about a deterministic problem. Was it you or was it Suzanne? Matt, did you call it? Okay. You labeled this as potentially looking into a deterministic solution. My takeaway is that the sample size of the study was based on data from A and J and there is no sense of the scope and scale when you look at this globally. So, the cost-benefit analysis will remain, is being done mostly on a non-

deterministic landscape. Correct? Right. And so, our risk analysis is complicated here. All right. Thank you.

JAMES GALVIN: Okay. Thank you. And Becky, please.

BECKY BURR: So, I usually don't say anything in these meetings because I usually don't know anything useful to say, but I do know about the privacy stuff, and it's a real issue. And in order to determine what to do, there's going to need to be a data protection impact assessment. I completely agree with you that the centralized model is a much more privacy-friendly design than the distributed model because whatever mitigation is needed, you can actually control that.

Having said that, what's going to be critical in the impact assessment is whether you really need all of the data elements. Now, you may need to collect it in a string, but if you guys can think about what are the elements that you're likely to pick up in a string, what you need from that, what you can get rid of, and how fast you can get rid of it, that's going to be the kind of mitigation that's required. Everybody around the table knows that there are lots of ways to go through a data set and overwrite nine-digit numbers, strings that might be Social Security numbers, or that kind of thing. So, whatever happens, it's going to have to be collaborative, because we're going to look to you to understand what data elements really need to be looked at, kept, maintained.

RAM MOHAN: Thank you, Becky. You can count on us to partner with you on thinking through that. We don't know the answers, but we know that together we'll probably fumble our way to something that is not very unsafe.

JAMES GALVIN: And I think that's just a fine place to thank everyone for your participation. Unfortunately, we are at the top of the hour. We're a minute over, and we do need to end. So, thank you everyone here for your questions and comments. Please, Board members, SSAC members, we're all here together now. Go out for your coffee break together and continue the discussion. So, thank you very much. We're adjourned.

[END OF TRANSCRIPTION]