ICANN77 | Prep Week – Name Collision Analysis Project Study 2 Update
Wednesday, May 31 2023 – 18:00 to 19:00 DCA

KATHY SCHNITT: Thank you. Hello and welcome to the Name Collision Analysis Project Study 2 update. My name is Kathy and I am 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 the session, question or comments submitted in chat will only be read aloud if put in the proper format as I will note in chat. I will read questions and comments aloud during the time set by the chair or moderator of the session. If you would like to ask a question or make your comment verbally, please raise your hand. When called upon, please kindly unmute your microphone and take the floor. Please state your name for the record and speak clearly and at a reasonable pace. Mute your microphone when you are done speaking.

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 into Zoom sessions using your full name, for example, a first name, a last name, or surname. You may be removed from the session if you do not sign in using your full name. And with that, I'm happy to hand the floor over to one of the NCAP co-chairs, Matt Thomas.

MATTHEW THOMAS: Thank you, Kathy, and welcome everyone to the Name Collision Analysis Project Update. My name is Matt Thomas. I am one of the co-chairs here along with Suzanne Woolf. And we're happy to give you an update as to where NCAP is progressing in Study 2 today. On the agenda, we'll just kind of similarly like we did at ICANN76, go over the background material of what the NCAP project proposal actually entailed, a little bit of history around Study 1 and Study 2.

Then we'll talk about some of the completed work items that have been accomplished by the discussion group, specifically three reports, the case study looking at Corp Home and Mail, the DNS perspective study looking at name collisions at different perspectives in the DNS hierarchy, and finally, the root cause report that looked at the reports received by ICANN for name collision incidents since 2012.

We'll briefly speak about some of the board questions and then dive into the findings that the discussion group is working on right now, as well as some tentative recommendations, which will lead us into the proposed workflow that the discussion group is working towards. We'll finish with a call to action for any of you that might be interested in participating in the final throws of NCAP here and how to participate, and then close out with some Q&A. So just for some brief background here, this is one of the diagrams that's currently in the Study 2 final report in the background section that really highlights the amount of work and bodies of effort that have been put into name collisions going back to 2012.

As you can clearly see, there is a plethora of different publications and other research that have been done in name collisions over the last decade. Back in June of 2020 is when we started the Study 1 report for NCAP. And Study 1 report was essentially an exhaustive bibliography or reference material of all things NCAP-related or name collisions-related that was produced in a final report that went out in March of 2021. Since then, we've been working on Study 2, which consists of a few deliverables that we'll go over here in a little bit. But I think it's just important to take a look at this historical timeline to really understand the sequencing and the ordering of things of when they happened. It's very rich with context to understand what some of the recommendations and the findings are when they're put in this certain historical context.

So regarding Study 2, the board actually sent a few specific requests to SSAC to conduct studies, to present data analysis, and to solicit points of view and provide advice to the board for how to deal with name collisions going forward in a sustainable, repeatable fashion. There were a few specific actions that the board asked regarding advice around the three non-delegated strings of home, corp, and mail from the 2012 filing, as well as a series of questions to solicit just general advice regarding name collisions going forward and how name collisions should be measured and assessed in terms of risk and potential harms going into subsequent rounds. The studies have been designed to be inclusive manner, open to the community as a whole.

There are 25 discussion group members, including 14 SSAC work party members, as well as 23 community observers at the current time. The

NCAP project proposal here has a few specific links here that will be available in the presentation on the ICANN schedule that will take you directly to the board resolutions, the project charter, the project proposal, and then the community wiki, which is packed with over 100 and some recordings of our NCAP discussion group meetings and discussions and presentations there. So going back and looking at the NCAP studies, NCAP was designed to have three studies, the first of which was the study one that I talked about earlier, in which it also included properly defining name collisions.

As you might discover when you review the background material in the study two report when it's published, name collision definitions have kind of had an evolution over time in the denotation and connotation of name collisions hasn't always been aligned within the community. So one of the action items from the discussion group was to come to consensus on a properly defined name collision. And like I said before, the study one also really was an all-encompassing body of work to review and analyze all previous studies on name collisions. Meanwhile, on study two that we're working on is looking at suggested criteria for how undelegated strings could be considered if something like a name collision string.

What criteria could those collision strings determine if they should not be delegated into the global route? And also, what criteria or recommendations would there be to have those list of collision strings removed via remediation or mitigation plans? That's really the focus of where the NCAP study two has been focused on for the last several years. Finally, there's a scheduled study three for name collisions that's

looking at the analysis of mitigation options that will happen in the future. But for today, let's focus on our study two items.

So study two, the completed works and study reports. So as I mentioned before, there were three specific studies that were undertaken as part of study two for NCAP. The first one was doing a collision string analysis of Corp Home and Mail. However, that study was also expanded to also include other labels such as internal, LAN, and local due to the high query volume that they were receiving at the time of study using data from VeriSign's A and J route servers. Those three additional strings were receiving more than 100 million queries per day. That case study has been published and went out for public comment. It has been decided that there is consensus within the NCAP discussion group that the report is complete and it highlights changes over time of the properties of the DNS queries and the traffic alterations as the DNS has evolved over the last decade.

Those studies highlighted certain aspects surrounding various DNS protocol evolutions that we might talk about a little bit more here today about things like QNA minimization, the suppression of queries via things like aggressive NSEC, the introduction of public recursive resolvers, and so forth. The second study was a perspective study of DNS queries for non-existent top-level domains. And the goal of that study was really to understand the distribution of name collision traffic throughout the DNS hierarchy. And this is important because it provides some insights into how name collision telemetry can be assessed for risk management purposes going forward. The guardrails need to be known in terms of where you look at data and what that data

tells you to make an informed decision or assessment about potential name collision risks.

So that study looked at the differences between the various root server identities using the day in the life internet from DNS org as well as some data provided by some large recursive resolver operators that highlighted some differences between the traffics at the roots versus recursives. Again, that report went out for public comment. Consensus has been called within the discussion group, and those led to some of the findings in the study two report. The third report, the root cause analysis of new gTLD collisions, was conducted by the technical investigator hired by ICANN, Casey Deccio, to analyze the various aspects of name collisions from 2012 round based on the reports received by ICANN for their name collision incidents.

That report has also been completed. The discussion group has called consensus on it, and it has a number of insightful findings that will go into the study two that we'll talk about a little bit later here. So some of the key takeaways on these three studies would be, again, going back to the case study, is that the data clearly shows that the impact has increased for all of those strings over time. And that there are certain things that the discussion group has come to term critical diagnostic measurements that can help predict the impact of name collisions. We'll probably talk about critical diagnostic measurements a little bit later in this presentation, but essentially you can think of them as a way to quantitatively describe name collisions based off of DNS traffic data observed at some vantage point in the DNS hierarchy.

For this, it was namely the root servers. They follow things like query value, but also look at other dimensions like diversity, where diversity can be in multiple dimensions, like the number of different networks, the different IP addresses, the different number of ASNs, the different number of second label domains, queue types, so forth and so on. And essentially these are very heavily lifted from the fantastic work that Jazz did back in 2012 when they did all of their analysis back then assessing all of the applied-for strings. Now what we also identified in the case study is that two of the major problems for or root cause of name collisions still are DNS service discovery protocols and suffix search lists.

Now switching over to the perspective study of DNS queries, it showed some of the similarities and differences between the root server identities and public recursive resolvers. It gave us some insights into how much assurance one might have by taking measurements solely at a root server versus looking at other places within the DNS. And it also helped provide some foundations, talking through people like ICANN's ITHI or their magnitude service that they have deployed, that there are opportunities to extend existing measurement platforms that can help inform applicants a priori their application to be aware of potential name collision issues.

The final report, the root cause analysis, really again identified that the private use of DNS suffixes is widespread and that the impact of the TLD delegations ranged from basically no impact to somewhat severe impact by a few entities reporting. But overall, all of these case studies highlight one major theme, and that is name collisions are and will

continue to be an increasingly difficult problem. Now I also mentioned that one of the portions of the study two objective is to answer the board questions. The board solicited information specifically about Corp Home and Mail, but they also sent a series of nine questions that they wanted to get some more advice on.

Here we can see those kind of put into a rough taxonomy where the first again was focusing on defining name collisions. As I mentioned before, the evolution of a name collision definition has varied over time, so we're hoping to establish something that can be consistent and go forward as a stable definition. The second question really focuses on what is the role of negative answers? And this is a critical component of understanding the risk of name collisions. These strings that are leaking or intend purposely or unintentionally into the global DNS are currently receiving non-existent domains or negative answers from the root servers, and it was trying to understand what role that answer plays back in terms of risk, harm, and just general understanding of how those negative answers might impact services or software.

The third question really focuses on harm and understanding what potential types of harm there are from name collisions. And this is a difficult question because it's somewhat non-tangible and there's not firm guidelines on understanding all the potential implications of how negative answers can be interpreted from systems or users out there. So the discussion group has heavily debated about between harm and measuring impact or risk. Questions four and five then talk about specific advice from the discussion group to the board about what potential options are there for mitigating that harm. How could there

be ways to notify these leaking parties or to disrupt any kind of impact once a particular string might be delegated? Question six then is on the flip side is what are the remaining risks of after delegation occur?

And then the last three questions solely focus on undelegated strings and collision strings in general. The discussion group has already provided some tentative answers to these board questions. We plan on revisiting those answers as we complete the study two document to cross core or cross reference certain aspects of the larger study two report with the more succinct answers in the board questions to make it more impactful and helpful for the board when they interpret these answers from the discussion group hopefully. So let's move on into our findings. So our current findings really evolve around again defining a consistent definition of name collisions. And like I mentioned before, the theme of that name collisions are and continue to be an increasingly difficult problem.

As I mentioned also before about critical diagnostic measurements, CDMs, the discussion group has come up with some recommendations or proposed ways that can both quantitatively and qualitatively assess some of those potential name collision risks based off of the DNS data that's being observed. Now there's also clearly some limitations with that data that the perspective study highlighted as well as just understanding how the DNS works in general. And how that might also be impacted by potential scenarios for gaming or manipulation of data in certain contexts.

One of the other things is that we've identified that qualitative measurements are also really necessary to help understand the

broader picture of name collisions in which sometimes that qualitative assessment almost has to be done on a string-by-string basis. Suzanne, is there anything else that you would like to chime in here or should we keep on going.

SUZANNE WOOLF: I think this is pretty much where we are. We've been through the findings. We have some measure of agreement, but this is where this is currently where we're working here and in draft recommendations. So if you want to go ahead and also get into what we've got so far.

MATTHEW THOMAS: Sounds good. Thanks, Suzanne. Going towards more of the draft recommendations, one of the problems that we're really trying to solve here within the discussion group is creating a sustainable, repeatable, deterministic methodology for evaluating name collisions going forward. What we're trying to do is be able to identify potential high-risk collision strings, things that should not be delegated.

Now, when we talk about this, the discussion group has talked about this as framing name collisions in general as a risk management problem. How can name collisions be a risk management problem based off of the data that's available and the expertise that is available to the community there to look at those name collisions and then come up with an assessment and provide advice to ICANN board.

Specifically, the discussion group has talked about an entity called the technical review team that we'll talk about a little bit in the proposed

workflow going forward and how we potentially see name collision assessment being conducted. But at a high level, we're trying to look for those high-risk labels. If we can, can we also identify any do not apply labels with the given caveat that the rest of the strings would perceive normally through the normal application and delegation process. In order to do this, name collision data has to be available to some entity to understand it, measure it, quantify it, either and do some qualitative analysis on it. That means making those name collisions visible.

Discussion group has produced a technical analysis gap document between the study one and study two report that really highlighted some of the technology changes since the 2012 round. Things like I mentioned before around the introduction of large public recursive resolvers, things like QNAME minimization that literally suppress the labels to the root, so you're only receiving the label for the top-level domain. Other things like aggressive NSEC that might be suppressing queries that were sent up to the root because they're now being cached or things like NX domain cut, so forth and so on. It's come to the attention of the group that the impairments of what is visible at the root right now is not the same richness as what was available at the roots back in 2012. There needs to be a mechanism to potentially make that data a little bit more visible in order to do the data measurements and do the risk assessments for that.

Then finally, there also needs to be an opportunity for applicants to have the ability to create a mitigation or remediation plan if a particular string is deemed high in terms of name collision risks. Those plans

would need to be tailored probably on a string-by-string basis, but they would be largely informed by the data that is being observed and needs to be made visible.

On the next slide here, this is really the workflow that the discussion group is trying to flesh out the details both technically and procedurally as to how name collision risk assessments could be done going forward. One of the key tenants that we wanted to ensure here was that there is also capabilities for applicants in the future to be able to have, "off-ramp options."

We don't want TLD applicants to be stuck in a state of limbo, like Corp Home and Mail has been for such a long time. The workflow was really designed to also have a few strategic opportunities for the applicant to potentially decide to abort their application. The process starts at the application submission, and actually it starts right before it. What we call here is the static assessment.

We mentioned before that we identified a few opportunities that there are existing measurement platforms out there like ICANN's ITHI platform or their magnitude page that can provide some insights into name collision risks for applicants prior to them submitting their application. This at least gives the applicant an opportunity to already know before they submit their application that this particular string is likely to have some kinds of name collisions if it's on there.

Now, of course, if it's not on that list, it does not mean that that string does not have name collision risks. It's just that it wasn't in that particular measurement and processing view for that tool. Once the

application is then submitted, it'll go through a series of technical review team assessments, the TRT. We envision a independent, trusted, neutral third party to provide this technical review team functionality in which they will be in charge of collecting and assessing the name collision data, and ultimately providing an assessment and a recommendation up to the board. Here again at the purple step in the workflow, the TRT is going to perform their own static assessment. They'll likely look at the same measurement platforms that I mentioned before. Maybe they have some extra ones that aren't publicly available at the time.

But that'll give them some assurance that if the application is to proceed, that the name collision risks are non-material that delegating the TLD into the root to make name collisions more apparent, is okay. If that's the case, then that's when we entered PCA, the passive collision assessment. This is where the TLD is now delegated into the root, and the technical review team will be operating an authoritative top-level domain server for that particular string, which essentially allows them to do the equivalent of a root server system-wide collection of all the name collision data for that particular string. That assessment will run for some duration, and after that, there's a potential off-ramp option for the applicant if something sparks or jumps out as an aberrant anomaly or something, that the technical review team identifies a heightened situation in which maybe the applicant decides to no longer pursue that based off the name collision risks.

Barring that, the technical review team will then go into the proposed next stage of collection or assessment around something called active

collision assessment. This is technically still being fleshed out within the discussion group of exactly how this works with scope and whatnot. But essentially, this is a more detailed-oriented view for the technical review team to capture potential actual connection attempts outside of the DNS for systems that are relying on that string or that TLD. That would allow for a more detailed and thorough assessment to understand potentially systems, applications, software processes that are using that string and to provide a better qualitative and quantitative assessment of risk for that.

Once that assessment is complete, the technical review team will package together all of their findings and ultimately come up with a recommendation that is submitted to the board for the board's approval. If the board then approves, the TLD would then be granted to the applicant. Suzanne, is there anything else I missed on this one or you want to highlight?

SUZANNE WOOLF: Only that this is still very much work in progress. But there is the outline here. I think the thing to emphasize at this point is just that there's the meta problem. In many ways, the biggest thing we have to consider here is we have more time than we did in 2012 to think these things through ahead of actually opening a new gTLD process. But at the same time, we have to be careful because in some ways, we're going to be able to do better than we could before. In some ways, things have changed. In some ways, what we were able to do in 2012 will work just fine for now also, and we have to avoid adding complexity that we don't

need. So there's a balancing act there, and I think you can see it in this draft workflow.

MATTHEW THOMAS: Excellent. Thanks, Suzanne. Moving on, so just some of the responsibilities of the technical review team. I mentioned this before. We envision them being an independent and neutral set of experts. They would have to have a large understanding of how the DNS works, clearly how then also to measure and quantify that, as well as perform the risk assessments. But we really see that they have four responsibilities. One would be assessing the visibility of name collisions. Two, documenting their data, their findings, and ultimately their recommendation to the board. Three, assessing any mitigation or remediation plans that might be submitted by an applicant. And then four, acting in emergency response situations in case a string during the assessment period post-delegation needs to be addressed in a technical manner.

We also envision that there could be a separate party, the neutral service provider, who could be the same as the technical review team operator or another one, but they would be more focused on the operation of the servers that will actually collect the critical diagnostic measurements. As we mentioned before, certain aspects of that data collection have data privacy concerns, and that is still under discussion with the NCAP team.

But again, here, their four main responsibilities would be operating the passive collision assessment environment that we described in the

workflow, the active collision assessment environment, as well as doing any kind of log processing and analysis to get that data over to the technical review team, as well as assisting or performing any kind of emergency response functions that are needed.

So some of the considerations of what the TRT would actually look at, because there's clearly a large amount of onus on the technical expertise of the TRT there to do all of this assessment and analysis. These are a few of the considerations that we thought would be helpful for them. Try not to be overly prescriptive here, but provide enough guardrails for them to be successful going forward. Some of the questions that would probably, based off of previous name collision risk assessments and outreach and mitigation efforts that JAZ, VeriSign, Google, and other parties have done, really highlight some key questions to understand the risk and harm, such as where are these queries coming from? Do they come from a common set of networks or particular ASNs?

Do they have some kind of common second label domain or some other kind of label within the domain name that can help identify why or what that collision string is associated with? Do the queries come from a diverse set of networks or ASNs, or are they globally dispersed around the world? Those things all play into the success of the assessment. They all play into the success rate or feasibility of potential mitigation and remediation options available for potential collision strings. And then ultimately, two of the biggest questions there at the bottom was for making sure that the TRT has no reason to believe that either PCA or

ACA would be impactful or harmful to the general internet and the stability of the global DNS in general.

So with that, I think we're closing in on how to participate. Here's a link to the discussion group. As we mentioned, the Study 2 report is nearing completion. We are hopefully on the final throws and hope to get that out for public comment very soon. The target was public comment before ICANN77. That is definitely not going to happen, but we hope that it will be shortly thereafter. And with that, Suzanne and I are happy to take any questions from the group.

KATHY SCHNITT: Matt, we do have one in the Q&A pod from Steve DelBianco. I realize that a Do Not Apply list would never cover all collisions, but isn't that a useful step to take?

MATTHEW THOMAS: That's a great question there. Yes, like you said, it's probably never going to be an exhaustive list. Without the TRT potentially doing the assessments prior to the next round, it's unlikely that that Do Not Apply list would be populated without additional effort and time and money spent on assessing current top named collision strings observed at the route. I think the feeling within the discussion group right now is that those Do Not Apply strings would be populated on an ongoing basis going forward as they're uncovered due to applications being submitted and then proceeding through the workflow and the risk being identified in either ACA or PCA.

KATHY SCHNITT: Matt, we do have another question from Philip. Given the substantial possibility that once the string applications for the next ICANN gTLD round are published, blockchain domain creators may mint and launch new blockchain domain name BDN systems with identical string labels. Is ICANN evaluated or is it planning to evaluate the extent to which Web3 BDNs may collide with existing or new Web2 DNS TLDs?

MATTHEW THOMAS: Thanks, Philip. It's an excellent question. I'll just say at a high order bit on this that name collisions within the scope of alternate naming and blockchains doesn't directly fall within the remit of what NCAP is looking at. We have made several comments in the document talking about it tangentially as it clearly has some kind of impact, but the scope and function here of the NCAP is really not focused on that aspect of it. That would need to be undertaken somewhere else. Suzanne, do you have anything else you want to add on that.

SUZANNE WOOLF: Only that, as Matt points out, not only the naming, the process of coining names, but the process of resolving them for blockchain as opposed to DNS is going to be different. We can't really speak for ICANN. If the question was about org or board maybe somebody else that's here can comment. But as far as the work we're doing with NCAP, I think as Matt says, we're not in a position to address that directly. At the same time, I think that there should be some further work on that

and that some of what we're talking about just as far as how some of the issues arise and how some of the challenges work. I think what we are doing is applicable to some later work along those lines. Anybody else. I'm not seeing. So for NCAP the answer is I think we're doing some work that's applicable to that question, but it's not directly in scope for us.

KATHY SCHNITT: And Matt and Suzanne, I don't see any further questions at this time.

MATTHEW THOMAS: Great. Well, thank you, everyone, for attending. And if you have any other additional follow-up questions, feel free to email Suzanne or I. Happy to discuss. And we'd encourage you all to join the NCAP discussion group. We're always looking for new members. And thank you again for the opportunity to talk with you all today. Appreciate it.

KATHY SCHNITT: Thank you for joining. This session has ended. Please stop the recording.

SUZANNE WOOLF: Thanks.

[END OF TRANSCRIPTION]