ICANN75 | AGM – NCAP Final Update: Preparation for Public Comment
Tuesday, September 20, 2022 – 13:15 to 14:30 KUL

KATHY SCHNITT: Onsite participants who will use a physical microphone to speak and you should leave the Zoom microphone disconnected.

For the benefit of other participants, please state your name for the record and speak at a reasonable pace. You may access all available features for this session in the Zoom toolbar. With that, I'm happy to hand the floor over to one of our co-chairs, Matthew Thomas.

MATT THOMAS: Thank you, Kathy, and welcome, everyone, to today's Name Collision Analysis Project update. My name is Matt Thomas. I'm co-chair of the NCAP along with my other co-chair, James Galvin, here. We're happy to be here today to give you an update as to where the NCAP is in the Study Two and the development of its Study Two report as it prepares to go out for public comment. Next slide, please.

So, briefly on the agenda here, we're going to give a quick background on NCAP, what the project proposal actually entails, and specifically what we're working on Study Two in this context. We'll then dive into the completed work items that have been achieved so far in Study Two. This consists of several different reports as well as numerous discussion group calls. And then from that completed work discussion, we'll go into describing some of the findings that we have found so far regarding name collisions.

But the meat of this presentation will really then flow into item four, which is the description of a workflow around name collision assessment. The main objective of this workflow is to create a sustainable, repeatable, deterministic process that allows for ICANN to be able to assess name collisions going forward. We'll follow up with some general information on how to participate in NCAP, and then open up the floor to Q&A and hope to have a robust discussion with the community here today about this. Next slide, please.

So getting quickly into the background, ICANN Board asked SSAC to study data material and do some additional research into name collisions and to provide its point of view and some guidance on two main points. First, it was looking for advice on three specific strings from the 2012 round. Those are .home, .corp, and .mail. In addition to that advice for that particular question, there was a set of nine questions that the Board developed and asked for some advice and information on regarding name collisions. So that is the main two objectives coming out of this in addition to the development of this workflow.

Currently, the NCAP project as a whole has been designed to be conducted in a manner that is inclusive of the community. Currently, we have 25 discussion group members that also include 14 SSAC work party members. And currently, we also have 23 community observers. Next slide, please.

This is just more for historical record. These are probably the most important and relevant documents regarding the Name Collision Analysis Project for your reference. There are four links if you look inside the presentation to the Board resolutions, which contain those questions that I just described. The project charter for NCAP, the project proposal, and then of course the community wiki which contains all of our discussion group meeting recordings, notes, presentations. There's probably a hundred meetings in there so there's lots of material for you to go back and look at if you want to. Next slide, please.

So the NCAP was actually originally broken down into three distinct studies. The first one was a gap analysis study that was conducted several years ago. There are two main objectives of Study One Gap Analysis. The first of which was to define a proper definition for name collisions. That definition is what provided the context for what is in scope and what is out of scope for the remaining studies as it relates to name collisions. The rest of Study One was really focused on collecting all material related to name collisions and synthesizing that into a very detailed report and to perform a gap analysis to identify any shortcomings or missing information within NCAP.

The second phase, NCAP Study Two, is really looking at the root cause and impact analysis. This was focused on three main objectives: determining what kind of criteria when an undelegated string could be considered a string that manifests name collisions or what we call a collision string. One other criteria could be given when a collision string should be not delegated, and how can those strings be removed from that list going forward.

Study Three is to be completed and is a future piece of work for the NCAP. This study is mainly focusing on analysis and mitigation options for name collision strings. Next slide, please.

Now, these studies, their original scope was probably a little bit broader than what actually happened in a real context of what happened within the Name Collision Analysis Project. Study One had originally included an item of building a data repository. Upon getting into the work, it didn't become a feasible, scalable, or actually just doable item. So that item was removed from Study One. But one of the recommendations coming out of Study One was that Study Two should not proceed as designed. So accordingly, the NCAP Discussion Group and SSAC evaluated the original goals of the Study Two, taking that information from the Study One report, and reevaluated what was more appropriate for Study Two. As you can see, in the original goals there was that building in the data repository as well as building a test system. Both of those were removed from the scope of Study Two. But there was also an additional item added in Study Two that was looking at an impact analysis. So here we are in the presentation today, taking a look at the work that has gone on Study Two, and we were developing that report and hope to put it out for public consultation in the coming months or quarter, and Study Three will be afterwards. Next slide, please.

So the second revision of Study Two, like I mentioned before, we removed the building of a test harness in the data repository. But the main two goals of the Study Two that we were looking at was identifying the root cause that most name collisions that were observed since the 2012 round and understanding the impact of those name collisions. So to achieve that, this included several different tasks conducting a root cause analysis, which is looking at the ICANN Name Collision reports that were received since the 2012 round in which name collision problems manifested themselves and were reported to ICANN, and also conducting impact analysis relative .corp, .mail, and .home, as well as looking at a data sensitivity analysis that would help inform the discussion group to understand the limitations and guardrails for assessing name collision telemetry within the DNS hierarchy at various vantage points and what appropriate guardrails should be placed on that kind of analysis and usage. Finally, all of that work is then coming to the culmination of what we're trying to prepare right now and present to you and that is a report on Study Two, and then that report, like we mentioned before, it will be going out for public consultation. Next slide, please.

So let's talk about some of the completed work that we have achieved so far here. The first thing is the Case Study of Collision Strings. Now, these case studies really focused on .corp, .home, and .mail, .internal, .lan, and .local. .internal, .lan, and .local were added in addition to the original three because they were receiving at the time more than 100 million queries to A and J root servers. This just gives a little bit more robustness to the report and being able to compare the original three strings to some others within context.

The Perspective Study, like I mentioned before, was also trying to understand the parameters in which DNS telemetry data can be used for name collision assessments and how that can be put into our name collision analysis workflow that we'll be describing here in detail.

Then the last item is the Root Cause Analysis. Again, this is what I mentioned before, taking a look at those 2012 Name Collision reports and understanding what the underlying underpinnings of those reports were and what actual harm or impact there was to the affected parties. Next slide, please.

So I'd like to focus a little bit more here on the key takeaways. First, on the Case Study, we learned a few things by looking at some longitudinal data .corp, .home, and .mail at A and J root servers going back from numerous years, and that is that the impact has increased, that the traffic and the diversity for those particular strings that are leaking into the public root server system has increased. Accordingly, that case study helped identify what we have termed Critical Diagnostic Measurements. These measurements are a type of quantitative measurement that allows us to better assess the impact or potential harm of name collisions. We'll speak about those a little bit in a coming slide. The other things that we have noticed is that, again, some of the leading causes of name collisions are still what we originally identified in the 2012 round are DNS service discovery protocols and suffix search lists.

The second study that the group undertook, the Perspective Study, really kind of gave an understanding of how DNS name collision telemetry can be assessed when looking at the root server system and how accurate and/or complete a portrayal of a particular name collision string is when looking at one or more RSIs as part of the entire root server system.

Finally, we have the Root Cause Analysis that looked at I think 47 reports that ICANN received for name collisions. Some of the key takeaways from that included the use of private use DNS suffixes is still very widespread within the Internet community and that name collisions are strongly supported by the data that has been observed within the report. Most importantly, the delegation of certain TLDs had some impact ranging from severe to minimal. But I would say from all of these, the main takeaway here is that name collisions are and will continue to be an increasingly difficult problem to manage and assess. Next slide, please.

So let's talk a little bit more about these findings. So these are some of the key findings that we've identified out of the culmination of those three reports, as well as the NCAP Discussion Group presentations and discussions. There are definitely more findings that are going to be included in the Study Two report, and we'll encourage you to go take a look at those there. But these are some of the ones that we wanted to highlight. Again, I'll just reiterate one more time that one of the key takeaways is that name collisions are and will continue to be an increasingly difficult problem. That the identification of these things that we have termed the Critical Diagnostic Measurements via the case study are a way analogous to what was used in the 2012 round via that Jazz and Interisle reports on a way to assess DNS telemetry for name collision risk assessment, to better inform the risk assessment in our workflow that we'll be describing here today.

The third main thing that we've identified as well is that mitigation and remediation is problematic. However, that challenge increases as the volume and diversity of those CDMs also increase. So as more sources are emitting or leaking these types of queries for a variety of different CDMs, the mechanism in which you can remediate those are increasingly difficult.

Finally, we've also identified that there's potential opportunities to extend existing measurement platforms to help inform applicants a priori the application round to inform them a potential name collision risk before they submit their application. Next slide, please.

So as I mentioned before, the Critical Diagnostic Measurements or the CDMs are a set of quantitative measurements that describe the traffic that's seen within the DNS at particular vantage points in the DNS hierarchy. The first of which is query volume. Now, query volume might be a leading indicator but it does not portray the entire picture or risk profile of a particular name collision string. Just because the string has high query volume doesn't necessarily mean it has high impact or risk. Other quantitative measurements such as diversity help portray that risk profile of a name collision string in a more holistic view. Things like looking at the query origin diversity such as the number of IP addresses issuing that query, the number of networks, the /24s or ASNs issuing those queries helps give another dimension to assessing that risk. Other properties also include other diversity things such as the query type or they offer type A, AAAA, MX, so forth and so on. Also, what are the label diversities? Do they all seem to be coming under a particular second level domain? Do they exhibit other known labels that are associated with threat vectors such as WPAD or ISATAP or other known DNS service discovery protocols that are vulnerable to attack and name collision scenarios. Of course, then there's also a qualitative component of name collisions that can only be described by using some other open-source intelligence, googling, searching for data, and better understanding what the root cause and why those queries are being seen within the public DNS. Next slide, please.

So at this point, I'd like to hand it over to Jim who is going to start describing the workflow. Thank you.

JIM GALVIN: Thanks very much, Matt. In some ways, this might be some of the most interesting part of all of this. Everything we've listened to so far has been all of the technical discussion we've been having about trying to understand the problem space in which we're working. And then the rest of this is what would we ultimately be recommending to the Board to do?

So a fair question to start with in all this is what problem are we trying to solve? This is kind of important because we've actually gone around a little bit in our discussion group about exactly what it is, how to interpret the questions that the Board had asked us so that we can provide something that's actually useful. What we've settled on so far and the path that we're headed down is providing a methodology for identifying high risk labels or what the Board is calling collision strings. That's what it put in its resolution that it gave to us. So these would be strings that would be applied for that the Board is more likely to not delegate as opposed to just letting it pass on. Then implicit in that being the problem that we're trying to solve, it suggests that no other string would be blocked as a result of name collisions.

So one of the observations to make is over the last decade, we have delegated almost a thousand labels and, really, most of them, more than 90% of them have had the presence of collisions. And yet, in spite of that, nothing really bad has happened, at least not on an Internet context. Obviously, the fact that something happened at all is quite serious to the recipient, the person who experienced it, and we don't mean to dismiss that or underrate it. But this is important. It's just important to acknowledge that, that we're really only looking for high risk strings and collision strings. There's a lot of people who have suggested that "Gee, 10 years of experience. Why aren't we doing anything at all going forward?" and I think that that misses the point. The point is that there's always a risk, and that's the issue. ICANN has an obligation, as we all do in the community, to really assess that risk before we move forward. It's not an absolute yes or no question. So we just have to pick the risk that we're willing to assume and the risk that we're not.

Clearly, name collision analysis is a risk management problem. We've said that several times. We're going to continue to repeat that. That really is what's going on here. It's important to accept that and to acknowledge it. It's not a black and white situation. It's not an objective situation. There's a certain amount of expertise and subjectivity involved in analyzing all of this, and there really is just no way to escape that. In fact, you still have to accept that you could get it wrong, because in reality, it only takes one collision to be a very harmful event. So, volume all by itself is also not necessarily an indicator that something is really bad and is not harm that can't be mitigated.

So we're left with these questions more specifically is whether or not it's possible to identify a high risk label. We do have a methodology. We're going to walk through a little bit of that here which we think at least is based on what we know today how we can seek to identify some high risk strings. One of the questions that's interesting and hanging out there is there are some separate policy recommendations, most notably in the Subsequent Procedures recommendations, where it was suggested by the community that they would like for there to be a list of do not apply labels. At least currently, the place that we're at is it's really not possible to do that up front. You can have a list of do not apply labels. Basically, that list is the set of things that you're not going to delegate and you will discover them as you go along. So somebody will apply for a string that you ultimately decide not to delegate, and then it gets added to that list. From the Board's point of view, they had labeled these collision strings, the idea that they had in their resolution was what strings would get on this collision string list and thus would not be delegated? Then they also asked for a framework for managing that list. Meaning, is it possible to get off that list once you're on it? And that's a question that we'll have to get to that's really a Study Three question here along the way, but we'll probably say a little bit about that in Study Two once we get to the end. Next slide, please.

So the goals of our workflow are just to ensure that we can assess name collisions. It's just to create a methodology, a mechanism by which we can cause name collisions to be visible, the existence of them to be visible. And thus, you then have an opportunity to consider whether or not you want or need a mitigation or remediation plan. We do want to observe that it's probable in many cases that you don't have to have one of those plans explicitly. If the CDMs are on the low side in both volume and diversity, more than likely a mitigation or remediation plan would not be needed. You would have to allow for that not to happen and you would wait for something to come along some harm to really be experienced to become visible, and then you would have to reexamine the need for mitigation or remediation plan. It's important to note that that's essentially the system that's in place today. That was the decision that was made in 2012. There was no mitigation or remediation plan requirement ultimately in the 2012 round in the end. They just decided that there was controlled interruption. If they got that far, unless some real harm was identified, you just went forward. And there was no concern for trying to do something in response to the collisions that you had. They were only looking for knowledge about harm to human life at the time.

So the way to achieve this goal, the workflow, where you want to make them visible so you can assess them is we've come around to this idea that you need two operating roles in order to conduct measurements. This is expressly intended to be a functional description. It's not intended to represent how this would be implemented. This is important in the ICANN community because the obvious two choices for implementation—and there are other options but two choices obviously are in RSEP kind of model where there's some kind of standing panel of experts that are out there and they draw from when they need them in order to assess any given application. Or perhaps ICANN has to somehow acquire and cause to come into existence a team of experts that might be employees, it might be outsourced. I mean, those are sort of the two examples of how it might be done. Our role here is just to describe the job that needs to happen, and then the implementation of it will have to be figured out later. Next slide, please.

This is a quick look at the workflow. It's kind of a teaser at this point in the presentation I just wants you to see. This workflow really has not changed in the large since like last January or February in the discussion group, although we have absolutely made some changes to some of the details on the inside in order to make all of this fit together. You're going to get a good hard look at this as we go through the rest of this presentation. This was just intended to give you a quick picture of what's coming. So next slide, please.

The first of the two roles that has to exist in order to assess name collisions is what we're calling a Technical Review Team. This Technical Review Team, they need to be a set of independent and neutral experts. That's just a functional characteristic that has to exist. Whatever that means in terms of implementation, that's the goal that you're trying to achieve here.

The technical experts really do have to be experts in a number of specific areas. They do have to understand the DNS and the DNS infrastructure. They do have to understand the collected data, the

CDMs don't have to represent and should not represent just DNS queries. We are proposing that there'll be additional data collection about other protocols. So you need people to understand those other protocols and how they're used and be able to assess the connection data that's collected about the use of those protocols. Of course, they need to be able to understand and assess risk.

They have four responsibilities in those particular technical expertise areas. The obvious first one is being able to assess the visibility of name collisions. They really do have to look at the data and make a decision, make an assessment of whether or not that represents all the name collisions that could be there, the quality of the visibility of those name collisions. One of the concerns that we have is whether or not the Internet infrastructure and the DNS infrastructure and those kinds of things will change with time. Because a very significant observation between 2012 and now is the DNS infrastructure looks radically different than it did 10 years ago. Most of that is invisible to users. People wouldn't see that. But technologists are very aware that the DNS infrastructure is different. So the ability to assess name collisions is very different now than it was then. So there has to be an opportunity to evolve.

They are also going to have to document all their data findings and recommendations. I mean, this is fairly obvious. It's an administrative job but it really does have to be there. There's some pretty important decisions that are being made along the way here. So keeping track of what's going on. Another reason for that creating that documentation is for longitudinal review. This Technical Review Team is going to learn as it goes along as well as what it knows up front. And it needs to be able to have a body of reference material so that it can evaluate changes and see how things are better or worse going along.

It does ultimately have to assess mitigation or remediation plans when those are created. They're not going to be required at every step but they will be something that'll be necessary if something ends up being categorized as a high risk string or a collision string, as the Board would call it, then there'll be a need for a mitigation or remediation plan, and they will need to be able to assess that.

There's emergency response. Emergency Response is something that was talked about in the 2012 round but there was never any documentation about what to do in that situation. What emergency response in this situation means is if you delegate the label, which you have to do as part of the assessment process, and even in 2012, that's what happened in controlled interruption, the string was delegated, you really do have to have somebody who's going to make the call that this has to be removed from the root zone, this has to be undelegated. That really bad stuff is happening right now and it has to be undone. So there needs to be documentation about emergency response and all of the appropriate authority structures have to be put in place so that that can happen in a very short amount of time really measured in hours. So there has to be the ability to collect all of the right people together to make that happen. Next slide, please.

The other functional role that has to exist is what we're calling the Neutral Service Provider. In the 2012 round, controlled interruption was done individually by each registry operator. We're actually suggesting that given the nature of the expertise required to do everything that has to happen here, that that really should be centralized, that from a technical point of view, the best solution is for that to happen at a single technical location. So this would be some kind of operation where they would be responsible for operating all the servers that collect the CDMs in all of the Passive Collision Assessment and Active Collision Assessment, which we'll get to explaining in a minute here. So that's just something that has to be addressed. There obviously will be some data privacy concerns. We're not actually going to be able in our group to solve that problem because that's really more of a legal question than it is a technical question. Our responsibility is limited in scope to the technical part of this. So we will simply make our best recommendation from a technical point of view about what has to happen, and then obviously that'll have to be considered and applied in a broader context.

The four responsibilities that the Neutral Service Provider have, they do have to operate each of the collision assessment environments. So there's two of them: the passive collision and Active Collision Assessments. They will have to do some initial log processing and analysis in preparation for the TRT. The TRT is not going to want the raw log files. So we're presuming here that this Neutral Service Provider will do something in order to make the data presentable to the TRT. We're all leaving open the possibility, in fact, that this might address some of the data privacy concerns. It might be that only the raw data will be at the Neutral Service Provider. They'll have to do some massaging to it to make sure that they remove any sensitive information that might be in the raw data to give to the TRT to do its analysis. But those kinds of details all have to be examined and a decision made about them.

Then emergency response, they obviously have a role here. They might be the first ones that are actually monitoring what's going on. Since they're collecting the logs, it might be that they're the ones that are going to see first that something is going on, and they're going to have to alert the TRT. And then whatever other escalation path that needs to exist will have to come into existence. So they have a role in emergency response if it's needed. Next slide, please.

One of the big questions that we have gotten, especially in the last couple of months, it's kind of hit home, is this question of how does the TRT do its assessment? How does it decide that something is high risk? What is it going to look like? What kinds of questions does it got an answer? So this is just a collection. A small set, high level collection of the kinds of things that the Technical Review Team is expected to do along the way to deciding if something is a collision string or high risk labels. What it means to examine the CDMs is to look not just at the volume of those CDMs but the diversity of the source of those CDMs. So you want to look at things like which networks do all of those queries come from? Which ASNs are they coming from? To the extent that you can see the second level labels, you want to be examining those second level labels. You want to be looking at the source IP addresses. You want to be looking inside a bit of this connection data, especially like the second level labels, and see if you can learn anything from what's in the query string itself, especially on the DNS side. Or if you're looking at other protocols under Active

Collision Assessment, you want to look at some of the rest of the data that's available to you. You really do have to ask what's going on there. What is the impact? The impact is not just the volume of the queries. It's also about the diversity of where they're coming from. If you're getting a whole bunch of queries but they're all coming from one location, then you have to consider is that something which is mitigatable? Could I reach out to that one source and reduce all of those queries? And of course, if that one source is a recursive resolver, that presents an additional set of issues because that just means that they're hiding what's going on behind it.

So there's another level of investigation that has to happen. These are some of the technical details that have to be examined by the TRT. We're going to have some discussion about this set of things as a way to kick off the TRT. We do expect in general the TRT has to be an expert, knowledge in a lot of protocols, and exactly what they do and how they do it will evolve with time. As they learn more about what's happening, they'll certainly change up and get better at assessing name collisions, and that's to be expected. That's what you would like. Next slide, please.

Okay. So now we're back to this slide. What I'm going to do here is walk through each of these steps. All right. There are essentially five steps on the inside of the process and then one step each on the outside. So there's a beginning spot which is the applicant. Really, all that's going on here is the applicant has to build to get application, and then submit it.

We make the observation that there's an opportunity for what we're calling a static assessment. ICANN Org already publishes a set of data. It's the DNS magnitude dataset that it has on its website that it's already doing. That it's not a definitive thing because it's not all of the data that one would like to see, but at least it's a step in that direction. All it really says is, if the string that you want to use happens to be on that list and happens to be high on that list which changes on a daily basis, then you need to take that into account and you need to think to yourself, "Well, that means I'm going to get additional scrutiny." That's really all it means. If your name is on that list, you're going to get additional scrutiny, and you just have to expect that. Which means your process through, your time through this process might be a little slower because the TRT is really going to be looking at your data, and it's important to understand that. Next slide, please.

So the next thing that happens is at some point application processing begins and the TRT then gets an opportunity to make a decision about the first step in collision assessment. The TRT is going to make three name collision assessments. That's what we're proposing. The first is the TRT is going to look at the same static data that the applicant did. They're at least going to document what's there and what they found. If anything looks significant to them and is concerning to them, then that's what that little number two is up at the top, that's that offramp at the top. If they're thinking that this data, this string falls into the special case category, then they're going to engage with the applicant. Because they're going to have to ask the applicant, "Well, gee this looks concerning here." They're going to have to do some level of investigation and make a decision about whether they think it's a high risk string or not. It's probably a little early to make that assessment, quite honestly. But you never know, they might be able to come to that conclusion. So the applicant should always have the opportunity to decide, "Gee, I don't want to go forward." If you're already going to call me a special case and you're going to give me all of the scrutiny and you really think something's going on here, the applicant has to make a decision about whether they want to go forward or not. And if they do, they might even, at this point, be obligated to work with the TRT to do some additional investigation because it may be they have to start thinking about a mitigation and remediation plan. That's what the option is there at that little number two at the top is whether or not they want to continue and if they want to be responsible for a mitigation/remediation plan. Next slide, please.

So assuming all of that goes forward and they decide they're going to keep it or it has not yet been determined to be high risk, we go into what we call Passive Collision Assessment. Passive Collision Assessment is kind of interesting. It amounts to delegating the string but with an empty zone. So this has the feature that the most clients will continue to see the same NXDOMAIN response that they would have gotten if it wasn't delegated. I mean, in fairness, we do have to acknowledge that that's not a perfect solution because we can't understand everything. You don't know exactly how everybody does everything on the Internet. But in most cases, a very large percentage of cases, we expect that clients will still get their NXDOMAIN query, the NXDOMAIN response, it's just that it will take an extra query cycle for that to happen. So we've just extended the query cycle. Because now, in addition to querying the root, they'll query the authoritative server that the TLD has, and then they'll get their response.

The benefit of the Passive Collision Assessment is it really does give more DNS data to the TRT team and ultimately to the applicant if they need it. Because one of the most important things about PCA that it does is it pulls data out of the rest of the infrastructure. Without getting into the technical details, the reality is the static assessment doesn't see everything that's happening on the Internet. It only sees what ultimately gets to the root servers. And not everything gets to the root servers. What PCA does is it really forces to pull data in today's terms out of global resolvers. It forces data at a global resolvers into the root servers and then into the authoritative servers. So in principle, what you should see is greater volume. You automatically expect volume to increase. If you are a troublesome string, you would expect for the volume to increase. Whether or not diversity increases, that's just something you have to look at and you have to see and have to assess.

There's a little 90-day timer down there that's indicated. One of the things that's interesting is controlled interruption in the 2012 round was decided to be 90 days. We're not actually trying to change that. We have not in any of our discussions or analysis found any reasons to change that 90 days. So we're not actually having an opinion about it. It may be that others will have some kind of opinion or some data or evidence to offer about it. But for now, that'll be a 90-day period to do PCA.

One important thing to say before I move on to the next thing is keep in mind that although this looks long, because you'll see the 90 days coming up on the next one, it's also important to keep in mind that all of this can happen while the application is otherwise being processed, and otherwise, all the due diligence is happening. So you can do this technical name collision assessment in parallel with everything else that has to happen when it comes to analyzing a TLD in the application. It also turns out it really is an independent activity so it doesn't have to be delayed. You can do these just as fast as you can do them and you don't have to wait for the processing that's happening on the due diligence side. So that's something to keep in mind when you think about whether this 90 days here and the next 90 days is too long a period of time. Next slide, please.

Active Collision Assessment is distinguished from Passive Collision Assessment in two ways. One is that it actually covers both IPv4 and IPv6 addresses. The controlled interruption mode that was done in 2012 only covered IPv4, it did not cover IPv6. So that's a significant feature. Now appreciate that IPv6 doesn't have broad deployment but it certainly does seem like a real gap that we were not doing that assessment in 2012. I mean, I appreciate why they did it at the time. But going forward, we really do need to find a way to incorporate all of the technologies that we're aware of and that we know we can manage. So that's one key feature.

The second thing that Active Collision Assessment is going to do that we're proposing in all of this is we want the assessment to include more than just DNS queries. We need to learn something about how a proposed TLD string is currently being used to the extent that it's visible to us. So you want to be able to look at the second level label strings and you want to be able to see, to the best that you can, what are the protocols are being queried. It's not just about a DNS query. What happens next? Is it an HTTP query of some sort? Is it a web query of some sort? Is it some other common service on the Internet, some streaming level service that's being queried a lot? Whatever it is, even if we can't identify the service, is there a consistent query for something that follows from the DNS query? That's valuable information to have. It's a starting point for investigations when you're trying to figure out what's going on. You do want to look. Remember from the 2012 round, we had, for example, the WPAD investigation that happened. It was an interesting discovery on the part of Jazz when they were doing some early studies about what was going on with name collisions originally. Something like that could happen again, and you need the ability to see that data and see how the name is being used, what the next step is, what the next protocol is. It's just a starting point for investigation. And that's the purpose of Active Collision Assessment. It allows us to see how the name is being used and it's a starting point for investigation. So it's more information about the CDMs. In the same way that you use CDMs for DNS queries, you use them for these other queries. You get to see the volume of queries at these other protocol places, you want to look at the diversity of those queries in this other protocol, all of this informs your ability to identify high risk strings. This is the best that we have today. One of the things that we will allow for is the option that these kinds of things can change with time. It might be that what you collect, what you use changes. As the Internet changes, as usage changes, you simply have the opportunity to do different things.

You'll notice again here, number four, little blue four up there at the top, there's another offramp opportunity. If the data that's collected suddenly causes you to be identified as a high risk string, that is another opportunity for the TRT to reach out to the applicant and talk to them about "You have this high risk situation. This is going to have to be looked at and studied. Do you want to do that investigation? What kind of data can we provide you with and help you to do that investigation? We're going to need you to provide a mitigation and remediation plan so that that can be evaluated with your application as part of deciding whether or not it's going to be granted to you." So that's the opportunity for that, or it's an offramp. The applicant could decide that this is too complicated, too complex, it's not worth going forward, and so they would withdraw their application at that point. But they otherwise would be allowed to provide a mitigation/remediation plan which the TRT would then evaluate. All of that would be packaged up. Next slide, please.

Oh, never mind. Go back a slide. I thought I had one more with one more arrow about once the TRT is done and all that assessment is there, that package then gets submitted to the Board as part of the rest of the due diligence that the application processes doing. And yes, an important distinction from this workflow versus the 2012 round is, from a technical point of view, part of what we're asserting is that you really do need to assess name collisions. If you really want to respond to potential harm, you want to respond to the risk of a name collision, then you need to do that assessment as part of your ordinary due diligence. That is just a principle. A risk management security principle is the way that we're presenting that, you should do that before you grant the TLD. So that's a fundamental change from the 2012 round. I'm sure that the Board and the community may have comments about that, ultimately. And that's why all of this will go to public comment and we shall see. In the 2012 round, all of the due diligence, except for name collisions, except for controlled interruption, was done in advance. And then they granted the TLD, it was provided to the applicant. Then they had to do the collision assessment at that time. You did controlled interruption after the fact, after it was granted. We're just suggesting here that as part of a whole risk management process, that should all be done in advance prior to granting the TLD. That way, when you grant it and you decide to go forward, it's all clean and you know that you have what you have.

So that's the process overall. I kind of put a lot of words to this picture. There have been some wordy slides in the past that went through it. I hope that wasn't too much to listen to. I apologize for just going on at length about it. Next slide, please.

That just leaves us with the next slide, which is just a reminder that you can join in. You do have to apply to be a member. All applying means is you just have to answer a few extra questions about your relationship to potential new gTLDs in general. But otherwise, it's open to anybody. There's still an opportunity to get there.

We are in the discussion group. As Matt had said, we've been imminent to producing our work product here since June. Part of the problem is, the more you write, the more we realize we have to write. So the study report is getting a bit lengthy at this point. But our goal is to be complete, not to be hurried about getting a work product out. We've gone quite far here. We're finally getting our findings all well articulated and we're just about ready to move to the recommendations. As soon as we can get all of that documented, this thing will come out for public comment. So we're targeting as early in quarter four as we can get it together and get it done. The discussion group, of course, will have to review the full document once it's all there. Hopefully, that won't take too many weeks, and then it'll be out. The usual 40-day public comment period, and then it just gets submitted back up to the Board, and that's where that will be. Then we will begin a discussion about Study Three in the discussion group.

The next slide I think is just for Q&A, and I think that's it. That's where we are. Certainly open for questions if anybody has any comments or questions. I'm looking in the Zoom Room here and not seeing any hands. No one's raising hands in the room. Okay. Do we get to pat each other on the back or something? We must have done a good job, right? Either that or you're all sleeping. Okay. Danny nodded his head, he's asleep. That's all good. I see a comment in the Zoom Room. Oh, okay. No questions in there? Okay.

Well, if there are no questions, thank you for being here. Matt and I were sort of wondering about this earlier. We're not expecting to do another presentation about this. Hopefully, we really will get the document done in this early, in this quarter four. It'll be out the door. In principle, this is it. At best, maybe we might think about—I'm thinking at the moment—unless somebody wants to suggest something different, we might do a webinar of some sort as prep for ICANN76 as part of Prep Week, just to do a presentation. I think that they often do that with PDP results and stuff like that and ICANN and this might be something that that's worth doing that for. So have another detailed discussion about what the workflow looks like and some more details about the findings and recommendations once they've been articulated carefully. It would probably be in order. But I don't expect to be up here in front of the room with anyone in the future unless you want to ask for it.

I guess with that, any closing comments from you? No? All right, then thank you very much for being here. We are adjourned. I give you back five minutes of your day. Twenty? I thought we were only an hour.

[END OF TRANSCRIPTION]