SINGAPORE – Name Collision Mitigation
Monday, March 24th 2014 – 13:30 to 15:00
ICANN – Singapore, Singapore

FRANCISCO ARIAS: Hello, we're about to start the name collision session. I would appreciate if you take your seats. We'll start in a few minutes.

Hello, everyone. This is Francisco Arias, director of technical services at ICANN. We are going to start with the name collision session. I have here also Jeff Schmidt, CEO of JAS Global Advisors. Together we're going to talk about name collision more detailed update.

So the agenda for today I have to confess is very similar to what we had in the webinar a few days ago. What we did is try to incorporate some of the questions we received in the presentation. So, hopefully, that would make -- that would solve some of the questions that we had before.

So the genesis of this is that the collision occurrence management plan was approved by the ICANN board new gTLD program committee on 7 October. The plan contains a few points. First is the deferred delegation of home and corp indefinitely to commission a study to develop a name collision occurrence management framework, which is what Jeff is going to talk about in a few minutes. And also, following that framework, each new gTLD will receive an assessment based on the framework that would tell that registry what are the mitigation measures that have to be implemented regarding name collision?

Finally, there are two or more elements. The plan also includes what is called an alternate path to the delegation for eligible strings. Most of them were found to be eligible. Only 25 were not.

Those strings that are eligible for the alternate path for delegation, they can proceed as of now to delegation as long as they implement what is called SLD block list. This is a list of names that they have to block for activation in the DNS until the framework is approved. And then they proceed to implement the measures described there. And, finally, the last element in the plan is an outreach campaign to those entities potentially affected by the name collision.

So next I'm going to ask my colleague Nicole to talk about the outreach.

NICOLE DAVENPORT: Good afternoon, everyone.

So in December of 2013, we launched an outreach campaign that's multifaceted. The primary goal was to raise awareness of the name collision issue but, more specifically, to help educate IT professionals on how to safeguard against name collision.

Our first step was to create a resource hub at icann.org/namecollision. Some of the resources that you'll find there will be a video overview, a blog post by our security expert Dave Piscitello explaining at a high level what name collision is. And this is complemented by a very detailed report that Francisco and his team put together called the Guide to Name Collision Mitigation for IT professionals. It has all the details that you'll need.

We also have frequently asked questions and a form allowing Internet users to report if they've had an incidence of name collision that has caused them harm.

So once we created all these resources and put them out there, our next step was to do some outreach and get that information into the hands of people who would be most affected by it. So we targeted key technology and media outlets and influential industry associations. We've issued an initial press release. And, since that point, we tracked nearly 40 pieces of coverage in 14 countries and 6 languages. The response has been really good. The response from our -- we've actually also reached out directly to industry associations. And the response has been very good. We've reached out directly to 100 associations. And they've been willing to either publish the information or get the information out directly to their own members.

We followed up with another press release, which is serving as the point of reference for both associations and media. And it's received positive response from the likes of the European Telecommunications Network operators. After getting information out, we are amplifying through social media. We have everything that we release, all educational materials, all press releases are published via Twitter and amplified. We have about 66,000 followers and 10,000 on Facebook. We also launched a campaign on LinkedIn. We went specifically and we identified groups in the space and went and engaged with them through their forums. Some of the groups include the CIO Forum, the CIO Network, the CIO Exchange, CTO, CIO Leadership Council, CIO Community, and the Forbes CIO Network. Our outreach is ongoing.

We recently created a name collision kit. It's basically a zip file with basic info on including infographics and other materials. If it's something that you're interested in having and sharing, you can contact us at GDD-communications@icann.org.

And Francisco and his team just recently launched a public mailing list. And the idea there is to allow the community to come together and talk about how we can further our outreach efforts. So I believe on the screen we have the information there. It's NC-info@icann.org. Is it Francisco or Jeff? Jeff. Great. Thank you.

JEFF SCHMIDT: Good afternoon, everybody. My name is Jeff Schmidt, and I run a firm JAS Global Advisors -- not consulting, but close -- that ICANN asked to look at this problem.

The -- the scope of our study -- what ICANN asked us to look at was the impact of potential collisions on the end user on the consuming devices and community. And this is an important distinction that I want to be clear of right up front. The initial evaluation component of the entire new gTLD program comprised a DNS stability string review component that was designed to look at potential strings impact to the overall security and stability of the Internet and the Internet DNS. That was from the perspective of the global Internet DNS. Our study was the other way looking at the collision issue from the perspective of devices and software that are consuming the global DNS.

And we found no evidence to indicate that the overall security and stability of the global Internet DNS was at risk as a result of collisions, essentially confirming the results of the study during initial evaluation.

The objectives from -- you know, starting in the beginning of our study, there's a lot of discussion about the frequency of the potential collisions. The very good Interisle paper, studies from SSAC all indicated the potential for name space collisions and gave some indications of their potential frequency. But the consequences or the impact was really the unknown at this point.

And so that was where we focused our efforts was, okay, you know, we understand that there is a probability that collisions will occur. But then what? What are the potential impacts, potential consequences? We started off recognizing that, you know, every time there's a potential for a collision, it could be indicated by a record in the DITL data sets, the day in the life data sets, operated by or consolidated by DNSorg that are quite frequently used in collision analysis. Or some other indication of collision doesn't necessarily mean that a collision occurred. If a collision occurs, that doesn't necessarily mean that there's a serious problem. And so getting an understanding of the potential range of impacts of a whole range of end systems was a key component from our perspective.

Quickly, we also had to set a foundation with a definition of what is a collision? I'm not going to read these to you. But the Interisle paper defined a collision, and SSAC report number 62 also defined a collision. And, you know, there's different words in these. But the emphasis added on the slide is mine where, really, the crux is a confusion aspect or a difference in expectation between the end user or the resolving

device between what they got and what they expected. So that's the -- that's the crux of what makes a collision a collision.

Just briefly summarizing, there's been quite a bit of community interaction around collisions. This is kind of in reverse chronological order. VeriSign hosted a -- actually, a fantastic event at the end of IETF two weeks ago in London. We want to thank them for that. Brought together a number of researchers in the space to talk about collisions.

Our paper was released at the end of February. And then there were a couple of other events prior to Buenos Aires where the issue of collisions were also discussed. And we took part in all of those in varying capacities.

Over the course of our studies, one of the things we committed to do was to be open about where we were, our thinking, and the progress was a very time-compressed effort. And so we really wanted to have a public comment period before the public comment period. And we did that through interacting with various technical mailing lists and a couple of guest blog posts on DomainIncite, which were I think fairly widely read.

We also established bidirectional communications with a number of folks that are actually experiencing collisions in their operational systems. We were able to, you know, identify very specific situations where collisions were occurring and reach out and find those users responsible for those devices and talk about the experiences and in some cases conduct experiments with them, in some cases link them up with vendors of their hardware and software to talk about the issue.

One of the key findings was that collisions are, in fact, not a new occurrence. DNS namespace collisions.

DNS namespace collisions occur throughout the entire hierarchy. There's indication that they have occurred at the top level for all previous introductions of TLDs since at least 2007. That's where the data is available.

And as an indication of that, this chart, the DITL data sets -- again, that's the DNS OARC data sets -- have pre and post TLD delegation information back to 2007 for the TLDs that are listed on this chart. And one of the things that we did was compute a theoretical SLD block list, a second-level domain block list, based on the available DITL data for these TLDs.

There's no requirement for these folks to block anything. We just did it to kind of get a sense of if there was a block list for these folks, how big would it be, and used that as an indication of preexisting collision and collision-like activity in these TLDs.

Because of the way that the SLD block lists are computed, it's basically a union of all available data sets, all available years with DITL and other data sets, and so based on the dates of delegation on these TLDs, not all years are available. The ones in red actually only have one or two years available, and the numbers in a theoretical block list are still, you know, fairly significant.

The numbers here aren't designed to give a false sense of precision, just other than the fact that they're large in many cases and they're not zero indicates that this is not a new occurrence.

If we did a survey of the -- in fact, we did a survey of the research and found that namespace collisions and related phenomena have been mentioned by a number of DNS-related researchers since as early as 2003. In a nutshell, there's a significant amount of queries that appear at the root that basically should not appear at the root for a number of reasons. One of the biggest contributors is queries into undelegated space, queries into TLDs that have not been delegated. And this phenomena has been pointed out very widely.

The -- We also note that ICANN conducted two previous pilot rounds as well as the, you know, ongoing delegation of CCs and IDN CCs, and no serious consequences around collisions have been indicated from those events.

We found that the specific failure modalities -- There's a couple of primary drivers for collisions that we'll talk about here in a couple cases, but they basically lead to a couple of common failure modalities. And those modalities are similar no matter where you are in the namespace, whether you're at a large delegated generic, a small CC, an IDN, at the first level, at the second level, or even at the third level. The way that machines tend to fail when presented with a collision tend to be the same at all those levels.

And our conclusion from this was that the currently contemplated namespace expansion at the top level that ICANN is in the process of implementing, you know, it doesn't fundamentally change or increase the risks associated with collisions. Collisions happen. They have basically always happened, as best the data can tell us. They happen throughout the DNS. They will certainly happen in the newly delegated

space just as they happen in the currently delegated space. And the modalities and the consequences will basically be the same.

So why? Why are we seeing so many collisions currently and historically?

So one of the things that we found over the course of our study was a surprising lack of appreciation and understanding of the DNS at a fundamental level.

Many systems interact with the DNS in opaque and unclear ways that are not necessarily known to the administrators and operators of those systems. The most common situation there we found was related to corporate directory infrastructures, LDAP infrastructures. Things like Lotus Notes and Active Directory and such that have basically an overlay component of the global Internet DNS that is not always clear to the administrators at the time that these systems are set up.

DNS search list processing is another significant contributor as to why this is happening. DNS search list processing is basically our machines all being very helpful and causing a structured exploration of the DNS namespace whenever we type anything. For example, your machine may add JSadvisors.com, yourcompanyname.com, or other parts of the namespace to help your short queries work.

Search list processing has been a topic of another -- of a series of papers, including a recent SSAC paper, but I think this is a -- this is a really important topic to understand.

DNS search list processing effectively creates a large number of synthesized requests into the DNS, and many of those collide into undelegated space or fail for various reasons.

It's also interesting to note that there's actually two levels of search list processing. Host operating systems, since the beginning of networked operating systems, have done some sense of DNS search list processing, but many applications, particularly Web browsers and, to a lesser extent, email clients, do some sense of DNS search list processing as well that interacts with the OS level search list processing in sometimes unexpected ways.

So those are both kind of accidental causes of DNS namespace collisions. The first two bullets.

The third bullet are intentional use of namespaces that either are not delegated or not under the control of the person that's using the namespace.

So this is a situation where I intentionally use dot corp or I intentionally use dot company name or dot home or something along those lines just because either I didn't know that I couldn't or maybe I did and I thought it was okay. But for whatever reason, I'm intentionally using a namespace that I don't control.

The -- Another cause of namespace collisions that we found was the retirement of host names and second-level registrations. This is, you know, a machine is retired or a service is retired, but devices are continuing to query that device, the retired name. Sometimes that

name gets recycled. Sometimes that name gets repurposed or, in some cases, picked up by another party, causing a collision by definition.

And finally, you know, we found that colliding namespaces can be purchased. And in fact, they often are purchased. This goes by a lot of different names. Some people call it investing. Some people call it domaining, drop-catching. But those are all collisions by definition if you go back to the -- to the two definitions that talk about a user winding up somewhere that they didn't intend.

So one of the things that we did as a part of our study was look at other important namespaces that have changed over the years. Obviously the DNS is a very important namespace globally, but it's not the only very important namespace to have changed.

The examples that we found were around phone numbers and postal codes.

Between the '40s, when phone numbers were -- or when phones were not dialed by number, you know, you called your operator and made a connection, to three-number dialing, four-number dialing, in some parts of the world five-number dialing, six-number dialing and now in the U.S. we're all the way up to 11-number dialing in some parts, every one of those changes has necessitated a change in behavior from a user perspective as well as a change in behavior from a systems perspective.

One of the concerns whenever the namespaces in the phone number system has changed has been around PBXs and in other sorts embedded devices that interact with the phone number namespace.

Also in the U.S., and in other parts of the world, the area code and the exchanges have actually been changed in a number of places. So you have namespace expansion. You have additional numbers, but then in the U.S. you also have, you know -- I was born in the 216 area code, which served the west side of Cleveland, Ohio, and that namespace became the 440 area code in my lifetime. All those phone numbers that everyone had actually changed.

That has happened a number of times all over the world as additional demands for phone numbers have emerged.

Postal codes have also changed. Large buildings, growing cities, shrinking cities all get new postal codes, and sometimes postal codes change, sometimes postal codes are decommissioned. So these are important namespaces that have changed over time.

One of the things that we found when this has happened, that it's helped make these changes nor effective and more seamless, are the use of advanced notification. So basically a marketing campaign. If I'm going to change your area code or change your postal code, I have a notification campaign that says, hey, you're postal code is changing. And then some kind of a grace period or what we termed a NACK or a negative acknowledgment period. This is where if I call my old 216 phone number, instead of it going to some other random party, the phone company actually returned an error saying, whoops, by the way, you need to be dialing 440 now. In postal codes you get a nice little stamp on your envelope and it's returned to you saying it could not be delivered.

But some kind of a confirm nation that you're doing something wrong is an important component of a change to an important namespace like this.

What we found was that those negative acknowledgment periods and those transition periods typically are relatively short in the overall scheme of things. 30 to 90 days were typical.

And, you know, one of the other things we found that was actually kind of amusing, in looking into the phone system, we found in the '50s, the advent of what was called the Anti-Digit Dialing League, which apparently had chapters all over the U.S. and they were very concerned about the change from named exchanges to numeric exchanges. And the concerns that they were raising were both would cultural as well as operational and technical. They were concerned that this was an important namespace and it was changing. And they got quite a following to try to convince the phone companies to not make a change to the namespace.

I told a story at the London IETF that my dad would probably have been a member of the Anti-Digit Dialing League when our phone number changed from 216 to 440 back in the Cleveland area.

But the message here was there's always resistance to change, but history has shown that important namespaces can be changed, and -- and negative acknowledgment periods tend to help that.

So a quick summary of our report. The -- you know, we're recommending that home, corp and mail are not delegated. In the IP space, there's what's called RFC 1918 space which comprises the IP

addresses that may be familiar with you, 192, 168, ten dot, et cetera. These are IP addresses that are known to be okay to use locally. They're not supposed to be routed on the Internet.

No equivalent of RFC 1918 space exists in the DNS, but it's fairly obvious that particularly home and corp have been basically appropriated for private use. The use of these namespaces is very broad and hard-coded into a number of installations, a number of scripts, a number of devices, and it would be fairly difficult to undo at this point.

Also, it's pretty clear that the Internet has shown a need for RFC 1918-like namespaces in the DNS. So it seems natural to use the ones that have, again, basically been appropriated for this purpose to -- for that.

For the remaining proposed TLDs, we're recommending a negative acknowledgment period we called controlled interruption that buffer the potential legacy use of a TLD from the new use of a TLD.

The technical measure we're recommending there is the use of the loop back or the local net subnet. 127/8 is intended to never leave the interface that it originated from, and it -- for a bunch of technical reasons that are described in the paper, it makes sense to use that for this negative acknowledgment period to communicate with folks that then need to make a change.

Nondelegated TLDs, it's okay or we recommend, rather, to implement the controlled interruption by using a wildcard record that will basically for the 120 day period cause any inquiry into that TLD to return the 127 IP address.

For a delegated TLD that has elected the alternate path, these TLDs are in production. They have registrant information in them. They have active names in them. And for all of the reasons that the ICANN community came up with in the late '90s to not put a wildcard in a production TLD, we are recommending that we do not put a wildcard in a production TLD but, rather, implement the controlled interruption by individual resource records assigned to the second-level strings on the block list.

We recommend that in order to ensure uniformity, ICANN monitor the implementation. One of the things we thought a lot about was the response mechanisms. The SSAC report and a couple of other communications, when collisions has been brought up, talked about potential de-delegation of a TLD if it's causing problems or a response mechanism to make sure that ICANN is, you know, adequately equipped to deal with a scenario where a collision is causing a real problem.

And so, you know, we -- we thought quite a bit about that issue, and for obvious reasons it's really fraught with peril.

The threshold issue is very problematic. Basically, how do you decide between, you know, a couple of parties. One party could be using a second-level registration, following all the rules, doing all the right things, and depending on the proper functioning of that second-level registration.

Another party may be harmed by that second-level registration because of a collision issue and because of legacy usage, be it intentional or unintentional. So how do we decide between those interests?

Because all of those interests wind up being economic at the end of the day and the scope here is global in nature we decided basically that you can't, and that the threshold had to be extremely high for someone to take action regarding a problematic collision.

So the threshold that we're recommending is a very high one, admittedly, but we believe it's the only threshold that can be implemented, again, on a global basis, and that is if there is a collision that, for some reason, is causing a clear and present danger to human life, so this would be an industrial control system, a medical device, something along those lines where inaction is not acceptable, then it would be appropriate in those situations to take action. What does that action look like? Well, obviously, the first step would be to contact the registry to try to get the offending second-level delegations remedied, removed, suspended, adjusted, whatever the potential situation may be.

However, again, given that this is a very serious situation, a situation where there is a clear and present danger to human life, there has to be a backup plan. And the backup plan, in the event that the registry is unable or unwilling to comply, we recommend instead of a root-level de-delegation, which is extremely ugly and fraught with peril and unintended consequences and a very gross remedy to this issue, we recommend using EBERO and making surgical changes to the -- to the zone as opposed to, you know, gross changes to the root zone.

So in this case, if there was a second-level registration causing problems and the registry was unable or unwilling to make a change, that registry could be transitioned to an EBERO. The EBERO would then make the

surgical change required to remedy the clear and present danger to human life. Even though that's a very -- there's a lot of moving parts to that, it's actually far superior at our belief to a root-level de-delegation which again could effect hundreds if not thousands of parties in untold numbers of ways.

We were consumers of the various DNS OARC data sets, particularly the oft talked about DITL, day in the life, data sets. DNS OARC did a fantastic job of making the DITL data sets available to the researchers, a number of researchers that worked on this during this process. But I actually want to acknowledge DNS OARC. They did a great job stepping up to the challenge. They all of a sudden became very popular over the course of the collisions effort, and they really stepped up and did a great job.

I also actually want to acknowledge Sim Machines that was our partner in analyzing that data as well. They helped us and did a great job contributing to our work.

But the DITL data in a number of ways can be improved on. The attention to the queries at the root, who is asking the root what and why and when, there's more work to be done there. There are better ways to collect data to facilitate researchers in the future, and so our report did mention a number of suggestions for additional data collection at the root level to help future generations understand more about what's going on at the root.

A couple of responses to questions and discussions that we've been in since the Webinar.

The -- I talked a little bit about the clear and present danger to human life. I thought it would be fun to introduce another acronym because ICANN is short of acronyms, so there's another one.

The rationale for choosing such a high bar was really making sure that a serious situation could be addressed, but obviously I think we all don't think it's appropriate for ICANN to be involved in effectively commercial disputes between parties that have an interest in a particular string one way or another.

So that's one of the main drivers to setting that exceedingly high bar. The clear and present danger to human life is guidance to two sets of people. To the people that are reporting the problems. So, making it clear that, you know, ICANN will only request action be taken under this very specific circle cause a certain self-selection of reports into ICANN about potential harms so that, hopefully, the harms that -- hopefully, there are no harms being reported. But, if a harm is reported, it will be a sufficiently serious harm and we won't get a lot of noise being reported. But it's also guidance then to ICANN as they evaluate the request for potential action to potentially contact the registry or maybe even invoke the emergency response mechanisms at ICANN, including EBERO to respond.

A couple of quick comments about 127.0.53.53. For those of you that are technical in the audience, you will note that's a very odd address. It's odd by design. It's designed to be an indicator to somebody reviewing logs trying to figure out what the problem is. It's designed to stick out so that somebody will, hopefully, use the search engine of their choice to try to find why in the heck is this IP address appearing in my

logs. And then, hopefully, the information about this issue and the ways to remedy it will have sufficiently high search rankings to be visible to the folks that are looking into this issue.

We evaluated a couple of other mechanisms again for implementing some kind of a negative acknowledgment period. The two main alternatives were, instead of using a local host. So a 127/8 address, using an RFC 1819 address such as 10.0.53.53 or something along those lines. as well as using an Internet honeypot. So this would be an Internet routable IP address run by ICANN or some trusted third party that would provide information to help folks remedy the issue.

These -- all of these approaches have various positives and negatives, I would encourage you to take a look at our report where we talk about the plusses and minuses of each of these.

One of the things I wanted to call out is we thought a lot about the audience for the 127.0.53.53 messaging. Who do we want to reach? Who is going to be seeing this in a log, and what do we want them to do?

So we, basically, broke down the audience into a sophisticated category and an unsophisticated category, for lack of a better term. It's not meant to be insulting. The unsophisticated actors, we felt that we needed to protect them, cause no harm while, you know, they were figuring out this issue. So one of the things that we really liked about the 127/8 space was this won't cause any traffic to be transmitted outside of their network that wasn't already being transmitted. We thought that that was a very, you know, good feature. We could break

their system so that they could, hopefully, notice without exposing them to potentially new security issues in the process.

But one of the problems with 127/8 approach is, if you're trying to troubleshoot this issue, your logs are basically localized to the specific machines. So, if I have a thousand machines under management, my logs indicating the failed connections to the 127 space will be on a thousand different machines.

That was actually one of the advantages to the 1918 space in the honeypots were more of a centralized logging facility, so that somebody could get a better sense of what was going on kind of from a macro perspective.

But then we said, well, sophisticated actors actually already have a number of ways to potentially, you know, figure out that this is happening beyond a host level at a network level. And we experimented with a couple in particular using response policy zones and an intrusion detection system to detect at a network level that these 127/8 responses could be identified, if you're a large enterprise or an ISP, and wanted to get a sense at a macro level what was going on. And we know that can be done.

So sophisticated operators have additional tools available to them. Unsophisticated operators are protected from additional harms during the negative acknowledgment period.

Why 120 days? So this was another question that we got during the webinar and also at the London event. The -- this is a hard issue. The expired domain, you know, reacquisition policy, one of the things that

the -- one of the ICANN policies that we looked at that that was an inspiration to controlled interruption, actually only requires the -- the registration from the second level domain to be interrupted at a minimum of eight days. So, obviously, you know, there's a light year between 8 days and 120 days. So that's a benchmark that's out there.

There's another benchmark that's out there. And that's the CA revocation period, which, you know, we're probably all aware. And that 120-day period is designed -- to serve a similar purpose. To effectively buffer between potential legacy usage of a certificate and potential new usage of a certificate in a newly delegated TLD name space. That period is set at 120 days. So we kind of have this range out there. We looked a lot at how controlled interruption impacts different systems. Some systems have a very bright, very obvious response to the 127/8 response. They fail in a bright and obvious way. And that's great. That's actually what we want to have happen. We want the red stamp from the postal service on the envelope saying here's your problem. Go fix it. But, unfortunately, not all systems fail in a bright and obvious way. Some of the failures are a little bit more subtle, will take a little bit more time to figure out. And so that argues for a longer negative acknowledgment period while folks figure out the potential issues.

One of the other kind of factors here is the collisions issue is a serious issue. We found that real problems can occur when you have a collision anywhere in the name space. Delegated, non-delegated, first level, second level -- there are real problems that can and, in fact, have occurred. We worked with a number of end users, a number of vendors that there are issues.

And so that actually argues, you know, for the longer end of this spectrum while these vendors get additional guidance out to their users, while some of these vendors make adjustments to their software and hardware to actually notice the controlled interruption period. One of the things that we're particularly excited about is we're talking to a number of vendors now about building in detection of the odd 127.0.53.53 IP address. So, instead of an administrator having to go figure out what this strange IP address means, they'll actually get a log from their application saying, you know, we got this interruption, this instrumented IP back. You have a problem and go read this article or go take this action.

So these things tend to argue for the longer end of the notification period. We like the conservative approach, and so our report does recommend the 120-day period.

One of the other factors, finally, we're conscious that, you know, not all interaction with machines or potentially impacted machines is realtime. There are a number of scheduled jobs -- daily, weekly, monthly, quarterly -- a lot of jobs are scheduled quarterly, particularly in accounting and banking and finance. And so we want to make sure that any potential issues and jobs that only run quarterly are -- have an opportunity to be noticed as well.

That's it.

FRANCISCO ARIAS: Thank you. Now I'm going to talk about the interactions between the name collision mitigation measures and other provisions in the ICANN contract.

But, first, I would like to interject one question we received before is whether do we anticipate an effect on the right of the new TLD delegations by implementation of these measures that are being proposed? And the answer is no, we do not anticipate any effect on these.

Regarding the activation of names under the new TLD, it is currently a requirement to not activate names under the TLD for 120 days counted from the contracting from the effective date of the agreement. That has not been changed. That is not being proposed to be changed.

The proposal will be to add a second period of not activation of names accounted from delegation. And these two periods could overlap. In other words, you don't have to wait until the first period from contracting ends to start a count from the other period. In the previous section, my colleague Russ explained that, if all goes well for a TLD, they can go from contracting to delegation in approximately 60 days. So suppose a TLD that is going in this fashion, then you will start the clock for the contracting 60 days. You reach delegation, and then you will start the other clock for the 120 days. And at the end of those 120 days after delegation, you will have an effective period of 180 days in which delegation -- sorry -- activation of names was not allowed for the TLD. So it's not the sum of 120 and 120. It's not 240. It could be less depending on the how fast the TLD proceeds to delegation.

The only exception to the rule remains to be nic.tld. The reason I explained before is that this is required to offer such as like WHOIS. And we made a conscious decision to allow this service to be there because it provides other benefits to the community. And we also have the alternative to use the name collision report mechanism in case there was an issue with this.

Regarding the registration or allocation of names under the new TLD, they will be allowed. This has not been affected. It will remain as it is now that registration will happen subject to the RPM and other requirements in the registry agreement.

So names that are registered will be subject to, for example, sunrise if the TLD is during the sunrise process or claims, if they are in the claims period.

There was a question before whether there would be a requirement for, if a registry is not allowing registration of names in the SLD block list during sunrise. And where there will be a requirement to do another second sunrise period, once they release these names suppose they go for the 120 days in which they publish these WHOIS records in the DNS for the names in SLD block list. And the answer is no, there is no such requirement to do a second sunrise for those names. The only requirement in the RPM requirements is to have those names in the claims period.

The -- regarding the 100 names that are specified in a specification 5 for the promotion of the TLD, they will not be affected by this proposal. They will still be allowed. The only requirement that will apply is that they cannot be activated until the end of the activation period.

And, of course, they are subject to other requirements in the registry agreement.

Regarding the alternate path to delegation, we believe this new approach, the control interruption measure is superior and supersedes what we have in the SLD block list. And, therefore, once and if this proposal is approved, the alternate path to delegation would not be available any more for new TLDs that are delegated after this happens.

As Jeff mentioned before, TLDs that are already delegated, they don't have to do or they will not be required to introduce a wildcard in their TLD. We don't think that's a good idea. And the requirement will be to introduce these records for the SLD block list.

Finally, the name collision report mechanism, this is something that we have already working since last year. By this mechanism any party can report a significant harm caused by name collision to ICANN, and then ICANN will request action for the registry in regards to this name.

So far we have received no valid requests regarding name collision by this mechanism.

So the only effect on this provision is that there is a -- there will be a clarification on the threshold for this harm that could be clear and present danger to human life as explained by Jeff.

And with this, let's go to the Q&A session.

Eleeza, do we have a question in the chat? Would you mind?

REMOTE INTERVENTION: Yes, we do. The first one is from Reg from Minds + Machines. During the presentation, JAS said we are recommending that we do not put a wildcard in a production TLD but rather implement the controlled interruption by individual resource records assigned to the second level strings on the block list but did not discuss why the wildcarding option could not continue to be available as a choice to TLDs. Clearly, it was sufficient since TLDs have gone live using wildcarding and the Internet was not destroyed. Could an explanation be offered?

JEFF SCHMIDT: Thank you. I don't remember the SSAC report number. Somebody in the audience, I'm sure, does. But it's one of the early once, SSAC 3 or so, that discusses the issue of wildcarding and production TLDs or registry level or Internet level registries, or whatever they called it back then. There's a lot of history and a lot of science and a lot of work has been done that, basically, says that wildcards and production TLDs are a bad idea. And so we didn't want to go against that. And so we drew the line in our recommendation saying that a TLD that does not contain registrant data is not in production. And, thus, a wildcard would be okay and, in fact, preferable in those situations. Once a TLD contains registrant data and is in production, all of the reasons that are described in that SSAC report and others for not putting wildcards in production TLDs become applicable.

JEFF NEUMAN: Sorry. How do you turn this thing on?

Okay. It's on, great. Cool. Jeff Neuman with NeuStar. Thanks, Jeff, for the report. I think it was fantastic. It was really well done. You had a very short amount of time to do it. And I just want to comment you for the report you did. And the outreach that you did also as well, I think was fantastic in reaching out. I think it's the first time I've ever seen a report drafted by an independent group where you solicited comments during the actual drafting of the report. I think that is a fantastic way to go forward with other independent reports and hope that ICANN actually uses that kind of model going forward for other types of papers that it does.

I do want to, obviously, make a comment on the 120 days. As Francisco was alluding to, it really means 180 days from the day you actually sign your contract. I think that period is way too long. I mean, I know that you were looking for a day. And you even showed up on the screen something I was going to comment on, which is a typical NACK period is anywhere from 30 to 90. We operate the U.S. common short code system, which is a registry of 5 and 6 digit codes that you register -- that U.S. carriers register to do programs like -- for people like me that watch American Idol and you vote -- yes, I vote -- they have a NACK period of 60 days between when a user lets their code go to when a new user can register that code, simply because they don't want these types of collisions.

I would proffer that that's a much more reasonable time period than 120 days. I know that the CAs use 120 days, but that's from the date of signing the contract to give them enough notice. It has nothing to do with any kind of collision. And so I also -- you know, there was a thing on there that the clear and present danger to human life. I think any

system that could be interrupted or prevented from working is not going to take 120 days to figure it out. I mean, by definition, if it's going present a clear and present danger to human life, you're going to know that in day one, not day 120. So I really hope that ICANN has movement for that.

The last thing is I'm assuming ICANN will waive all of its fees for the first year until someone can actually put names into their registry. So I think ICANN will, hopefully, come out with that.

It says that, Francisco, you've created an exception for nic.tld because you said it's required for services like WHOIS. I mean, really? If you're not allowed to put any names into the TLD, then why is WHOIS a valid or valuable service? If you want to know who operates a TLD, you can go to IANA. You don't need a WHOIS on a nic.tld page. And the other thing, Francisco, you mentioned about, well, you can still do sunrise and you can still do claims. I think that ignores, really, the commercial realities of what anyone should do. You know, you're not going to sell names during a sunrise period or claims period if you can't -- if you can't allocate them within any kind of reasonable time period after that.

So I would actually encourage you not to show that slide any more, because it really does ignore the realities of all the TLD operators in this room. And, sorry. One more thing is there any reason why ICANN could not just do the delegations today, petition with IANA and -- for at least today it's the NTIA. But get all the TLDs in the root now; start the controlled interruptions; point them all to 127.0.53.53; get their support mechanisms up; get the education up so that we don't have to wait until

we sign an agreement, you know, whenever that is and then another 180 days. Thanks.

[Applause]

JEFF SCHMIDT: Thanks, Jeff. So -- and thanks for the kind words.

Regarding the 120 days, you know, I hear and I understand your concern. I want to point out to, you know, everybody here and everybody in remote participation, that our report is a draft. It's -- you know, currently in the ICANN public comment period. We kind of ran a public comment period before the public comment period. But it's now officially in the public comment period.

So it hasn't been even given to ICANN at this point for formal consideration. They certainly haven't agreed to do anything in it.

So, to the extent that you want to, you know, make comments about ideas to shorten up the 120 days, justifications, alternatives, please do. We welcome any comments and any ideas that you have to that effect. So thank you. It's a good comment.

FRANCISCO ARIAS: Jeff, regarding the comment about nic.tld, I appreciate that perhaps you don't believe this is something that having registrations, even if they cannot be activated, is not something you would be interested in doing. However, a lot of people seem to be interested in that, and so far the contract allows that to happen.

Regarding the idea of delegating all the TLDs at once, I think that's a very interesting idea. You may want to think about it and perhaps post it in the public comment in more detail. At first, I think there are -- I can see a few issues with that idea. Some of the TLDs have not -- a good amount of them have not been formally approved. So requesting delegation of them, I'm not sure how that will work.

Just first thoughts on the idea.

JEFF SCHMIDT: Thank you.

JEFF NEUMAN: Sorry. Can I follow up on that?

BRET FAUSETT: Argh.

JEFF NEUMAN: I'm sorry. I just want to follow up on that just as -- what's the harm of delegating a TLD as an experimental thing even if a contract is never awarded for it? There's no harm.

FRANCISCO ARIAS: I think there are implications that need to be understood. I don't think we are going to solve it here.

BRET FAUSETT: Thanks. Bret Fausett with Uniregistry. We are an applicant for a number of TLDs, including dot home. And I make that disclosure up front because I want to address this idea in the JAS report that home, corp and mail should be permanently reserved.

As a registry operator, dealing with name collisions is incredibly frustrating, but at least I take some small comfort in the fact that through the names mitigation process, we are educating users around the world about the proper use of the DNS. We are taking people who are using the Domain Name System in nonstandard ways and telling them about the perils of this and trying to get them to take corrective action. So when we say now that three TLDs actually might be permanently reserved, we're encouraging them to continue with these nonstandard practices.

So I would not call them permanently reserved. Clearly you've made a case that I think is compelling that the period should be longer than 120 days. Clearly these are outlier cases that need to be addressed in a special manner, but we shouldn't call them permanently reserved, because when we do, we're allowing the behavior that we're trying to correct for the other TLDs to continue.

Thanks.

[Applause]

JEFF SCHMIDT: Thank you. A quick -- a comment on that. RFC 1918 IP space exists for a reason. It exists for a reason because private networks need private IP spaces.

The need for private DNS spaces has emerged during, you know, this study. It's obviously been around for a long time, but I think the need for not allowing haphazard use of the DNS but funneling correct private usage to a couple of places is -- you know, is appropriate.

I just want to make the distinction, we're not recommending throw up our hands and say corp and home are beyond repair. What we're saying is there's obviously a need. Let's funnel that need, put it in a couple of places and cause people to pay more attention and do the right things.

Corp and home are also very interesting because RFC 6 -- there's an RFC that intimates that corp and home are safe to use for private DNS namespaces. It's fuzzy, about you somebody that is trying wholeheartedly to do the right thing could certainly have gotten the wrong idea.

So I just didn't want you to think that we were throwing up our hands and saying they're beyond repair. I think those are very special for a number of reasons.

JORDYN BUCHANAN: Hi, Jordyn Buchanan with Google. I have three questions.

The first relates to the 120 days. I'm just trying to understand sort of the data driving that period. And it's unclear to me from your explanation as to whether it's 120 days for detection of problems, and that would explain why we care about things like batch jobs, or if it's 120 days for mitigation of problems, which is why maybe the CA/B forum standard is useful, because it's clearly not a detection problem

there. You know what it is. And just sort of what data would underlie the 120 days as an approach or how frequent these problems are, are there quarterly problems either way, just to wrap my head around that. So that's the first question.

The second question related to what Bret was just asking on home and corp. Did you consider whether it would make sense to reserve a new namespace and use that for internal use, and then maybe wait a few years and see if people migrated off of home and corp at that point and see whether they were safe to use.

And the third one specifically about dot mail. I think for home and corp you're implying that those are used as internal networks, much like the RFC 1918 space is. Mail seems clearly to be used almost exclusively as a host name as opposed to an internal network. So I'm trying to understand, especially given that dotless mail wouldn't be allowed under the existing agreement, what's the rationale for reserving mail? It has a lot of usage, but it's all for the dotless name and it's not for internal networks.

JEFF SCHMIDT: Thanks, Jordyn. What was the middle question?

JORDYN BUCHANAN: Reserve some other namespace for internal use and reevaluate in a few years.

JEFF SCHMIDT: 120 day detection versus remediation. Good question.

The intention of the 120 days is actually both. And neither of them, in the case of a DNS namespace collision, are trivial.

So detection, you know, as I mentioned earlier, is not trivial, and we need to make enough noise so that somebody notices it in their logs. So that's kind of that phase. That's not going to go fast, in some instances. And then remediation is also not fast.

So, you know, you take a -- you know, an extensive -- sorry to pick on Lotus Notes, but an extensive LDAP infrastructure that, you know, that is using a DNS namespace incorrectly, that's certainly going to be hard for somebody to rename.

So the 120 days is designed to cover both, which is slightly different than the CA browser form 120 days, which, as you said, is all about remediation. The detection part has already happened. But we have both in the 120 days.

We thought about using, you know, like, going the RFC 1918 route and telling everybody to use dot internal or dot local which are basically already -- or dot invalid or one of the ones that are already reserved.

The issue there is purely about fighting the current. The usage of home and corp are, at present, so extensive and so widespread that fighting that doesn't system like it makes sense. That, you know, namespaces are needed, and those currently aren't delegated now. So it just kind of made sense to leverage that existing issue.

Dot mail, yes. So corp and home -- corp and home have been known issues since the first Interisle study where it was very clear that they were exceeding outliers on a number of dimensions.

We put together a multi-prong test internally to figure out whether -- you know, whether a TLD should be recommended for the reserved list in our report. And in that multi-prong test it included frequency sorts of things, like how often are collisions being seen, what is the diversity of second-level strings within that TLD, how many source IPs, how many source autonomous systems, all those sorts of things. But we also looked at evidence that the string is hard coded into systems, included in documentation, included in scripts, included in examples and all those sorts of dimensions. And that's really where mail stuck out.

You know, if you notice, mail isn't on top of any other frequency lists. It's in the top ten or so of probably all the frequency lists that everybody's seen. But when you include the other dimensions that we looked at, particularly around the example scripts and being hard coded in configurations and such, mail really -- the specialness of mail and the issues around the usage of mail start to become apparent.

We also liked mail because it's a generic term. It's not associated with a particular vendor, a particular product or something that's going to be temporal. You know, it's a long-lived generic term that is relatively consistent with, you know, other reservations. And so it just kind of seemed to make sense.

JEREMY EBBELS: Hi, Jeremy Ebbels from ARI Registry Services. Jeff, I have, hopefully, what is a simple question. For applicants who are already delegated, when would the 120 days start? And does anything need to be put in place before that period begins?

FRANCISCO ARIAS: I think I take this one.

So this is something that we need to define one of the details. This came up in the past Webinar, and we mentioned that this could be as simple as given that we are currently monitoring all the TLDs, (indiscernible) TLDs for the SLD block list, we could perhaps use that as soon as we see that there's the SLD block list has been -- is now policy name in the DNS, this special record, then we could start the clock. But in so many, this (indiscernible) has not been defined yet.

JEREMY EBBELS: Okay. Thank you.

DANNY McPHERSON: Danny McPherson. I had a couple of questions, a couple of comments. Jeff, I -- I do share Jeff Neuman's comments on you guys engaging with the community during this. I think you have done a fantastic job on that, so thanks for that.

One of the questions or one of the comments or observations I wanted to share is a lot of the issues or the concerns people had relevant to name collisions -- for example, in SAC 57 with internal name certificates -- isn't necessarily the collision itself. It's a combinatorial effect of the collision and the fact that someone can get a certificate for that namespace or that the public suffix list doesn't include that so super cookies and privacy issues may emerge as result and so forth.

So I think making sure that we continue to analyze those systems and deal with the residual effects -- or, I'm sorry, the residual risks in SAC 57 that really haven't been followed up on quite as much directly, as well as the things like the public suffix list, which I know SSAC is working on and ICANN is already doing a lot of work on. They've done some work with an ICANN fellow and others related to that.

So I think those external systems, and making sure that when you sort of close the time frame from when a new gTLD is delegated to when all the external systems are caught up is going to be an important thing to do as we go forward with new gTLDs.

One of the other things I wanted to ask, actually the lady who presented at the start of this said you had reached out to a hundred or so organizations. And I did receive from you, actually, Francisco the press kit related to name collisions and I think the materials there are getting much better. I was wondering if you've seen any measurable impact as a result of that.

You know, for example, when we did the CBA and then we reached out to the primary source of that, we saw about a 60% aggregate drop in the number of queries to our root servers as a result. And that's something that's immediate and measurable and means that, you know, their risk profile goes from this to this and allows applicants to move forward that much more quickly.

I know half of the answer to this, which is the measurement apparatus in the root server system doesn't have the capability to provide that today, but I don't think we're making any progress towards getting that, and I think that's a disservice to the community. And I think that some

work needs to be done there. And, you know, I think that, you know, in this role with IANA and the root operations and other things, the measurement apparatus, the data collection -- for example, VeriSign has been collecting this since we provided our initial reports. And so I think that's something that should be expected of the root server system so that these people that want to move forward more quickly can do that.

One other point. SAC 63, KSK rollover in the root, has a lot of recommendations in it like community outreach, like measuring things at the root system, like seeing what queries at what frequency and so forth. I think three or four of the five recommendations in SAC 63 are recommendations that also address some of the issues with name collisions. So evaluating that in that manner and making sure that we capitalize on the programs you guys are putting in place for community outreach will be really valuable as well.

My last question is when are we going to see the rest of the report? And what are the timelines from that point forward?

JEFF SCHMIDT: Thanks, and thanks for the kind words and thanks for all your help on this issue. You have been working on collisions longer than we have so thank you.

Kind of going through those, I'll take two and Francisco will take two.

You're right, the collision in and of itself isn't the problem. There's something that happens after that that makes a problem.

We found that there's two specific things that are the root causes of the problems afterwards.

One has to do with not being able to find a resource, so something causes, you know, a DNS lookup that was succeeding to fail. And that can be any number much things. DNS lookup says, as we all know, we're not as straightforward as many people think they are, particularly with the interaction of search list and stuff. So that's kind of one of the core issues.

The other one, we talked about this in London a little bit, is, you know, the improper reliance on DNS as an authentication mechanism. So, you know, just because, you know, I asked for the IP address for, you know, Jeff dot JS advisors.com and somebody gave me an answer, hey, talk to one dot two dot four, that doesn't mean I should just believe them.

So getting to, then, your other question about when is the rest of the report coming out, you know, we are going to talk about those two issues in great detail as well as release the rest of the data that we are not able to release at this point as a part of the complete report.

It mentions in our report the rationale why we split it up into two phases. It has to do with we did identify an issue that did not materialize in new space that existed. It materialized in existing space but it's serious enough that we're working with a vendor on a responsible disclosure process and want to give them time before we release all the data. You know, our data, once the vulnerability is known, will unfortunately help somebody do bad things and so we don't want to do anything there.

Right now our estimated time frame, again, it's driven by the vendor in this particular situation, but our estimated time frame is, you know, in the June time frame.

FRANCISCO ARIAS: We're running out of time. I would like to close the lines and ask the rest of the people in the line to be quick on the questions.

ANDREW MERRIAM: Andrew Merriam, Top Level Design.

I want first wanted to thank Akram and Francisco for acting so quickly on NTAG's request for an Webinar. That was great interaction. Really appreciate it, so thank you. And during that Webinar, I believe Francisco answered that having new TLDs blocking all domains instead of the APD block list would be better. So could you clarify objective criteria on how it would be better and why a solution that is acceptable to previously delegated TLDs would not be acceptable to address the same problem for those TLDs that are still working towards delegation?

FRANCISCO ARIAS: We believe the control interruption option is better, is an enhancement, and, therefore, it's why we think we should -- if this proposal is approved, to retire the other option. This is already considered in the contract --

ANDREW MERRIAM: Further details on how and why it would be better and for the communication would be ideal.

JEFF SCHMIDT: Yeah, so the intent of the block list approach, the block list just returns NXDOMAIN. It doesn't change the behavior. It doesn't do a negative acknowledgment. The controlled interruption actually initiates a negative acknowledgment. So the SLD block list preserves the existing behavior.

DMITRY BELYAVSKY: I'm Dmitry Belyavsky. I have two questions. First one, is there any procedure to premove the particular string from the collision list?

The second one is should we use the control interruption for the strings which are not allowed by the TLD policy?

Thank you.

JEFF SCHMIDT: Thanks. There's no discussion in our report about removing specific strings from the SLD block list.

DMITRY BELYAVSKY: What about nonallowed strings?

FRANCISCO ARIAS: Nonallowed strings? Like?

DMITRY BELYAVSKY: Well, for example, IDN strings in domains which policy forbids IDN. And vice versa.

FRANCISCO ARIAS: Yeah, the extension to the control interruption idea is to mitigate the name collision issue. I don't think that will apply to the use case that you mentioned. Eleeza.

REMOTE PARTICIPATION: I have several questions from remote. I will just list them all off quickly.

This one is from Rubens Kuhl. To Jeff Schmidt, how many algorithmic sources of SLDs, notably rogue applications besides Google Chrome 10, were identified during the research?

Would they explain what other sources tried to classify as a growing entropy namespace, and the growing number of SLDs to block that led to the 25 SLDs ineligible to APD?

And then I have two more if I can ask them, Francisco.

JEFF SCHMIDT: Yes, hi, Rubens. Yes, and many. We identified probably on the order of 10 or 12 algorithmic sources of queries.

FRANCISCO ARIAS: Go ahead, Eleeza.

REMOTE PARTICIPATION: The next one is from Limei Liu with CONAC. The SLD block list under some Chinese TLDs result from the registries operation of TLD test bed.

In that case, any single SLD on the block list will be resolved to only one corresponding address and is under the registry's sole management. This is fundamentally different from the ICANN's definition of name collision, plus the controlled interruption is unworkable for such TLDs as which will stop the normal resolution of the block SLDs and have negative impacts on user experience. Can you indicate how shall this situation be properly dealt with.

FRANCISCO ARIAS: So I think we responded to that question in a letter, but briefly, as it has been shown by research done by VeriSign, for example, in the CBA string, it is difficult to exactly define that all the quarters to a string come from a certain source.

Thank you.

Mikey.

MIKEY O'CONNOR: Thanks. My name is Mikey O'Connor, and I, too, participated in this study, and wanted to -- and continued after this study was done. I want to characterize one thing. There's a lot we don't know in the London work, in the seminar that was in London several weeks ago. When I

wrote my notes at the end, it was more a whole lot of really smart researchers saying, gee, there's a lot of stuff we don't know yet.

So let me give you one more thing, you community folks, that we don't know.

I've been running -- I contributed the DNS to corp.com to the JS study and to the Interisle folks because it's a delegated domain that's already got this kind of traffic, and maybe we could learn more about things from doing that.

One of the things that Jeff and I did was we tried out the 127.0.53.53 loop back address thingy to see how it would work. And I tell you true, given that I'm something of a scaredy cat about name collisions, that first time we put that in the DNS, that was a bad day for me.

Mikey cranky that day.

We did if for an hour that day, and we got no email. So the next day we did it for six hours. That was another bad day. And we got no email.

And remember, we had all this cool stuff in the root about, you know, if there's imminent harm to human life, here are 7,000 things you can do to get ahold of us so we can stop this right away. Silence.

So the next day we did 12 hours, and, you know, I was getting a little calmer about that. And since Jeff got done with the research part of his study, I've left corp.com delegated with a wildcard to 127.0.53.53. So if you want to test it right here in the room, go to anything dot corp.com and see what happens, because that's what's there today.

Now, the interesting thing is that corp.com has something between a bazillion and a jillion sources of traffic. I can't remember the number now. And in the intervening six weeks or so, I have gotten zero emails. And so this contributes to the "I don't know." I think this is a wicked clever idea and I'm really hopeful that it works but the hope for that idea was that we would alert people to a problem, and out of the somewhere between a zillion and a bajillion I was expecting at least one contact from somebody who looked in there because by now Google and other search engines are, indeed, producing results on that address that are pretty good. You know, the ICANN page is pretty close to the top of the list.

I was expecting some commentary. So I'm not sure -- This is just one more thing I don't know. I'm not sure that this alert mechanism is actually working right.

JEFF SCHMIDT: Thanks, Mikey. And by the way, thanks for all your help in this process. You've been very generous with the use of corp.com and have made all of our research products better. So thank you.

So one of the things that we did, we wanted to understand how the 127.0.53.53 IP address affects different systems. We have a sense of what kinds of systems, what kinds of software are generating the majority of these queries. Certainly not everyone, and there are bajillions of them, to use the technical term. But there are a couple of common denominators. LDAP based directories are a huge, huge source.

So one of the things that we did, we actually bought a couple of domains in com where we found collisions and put 127.0.53.53 in it actually before we did our experiments. And these were folks where we already had a relationship with -- a bidirectional relationship with the colliding system.

And we asked them, in some cases before and in some cases after, what they saw, what the behavior change was.

One of the main organizations that is querying corp dot com we actually have a bidirectional relationship with. And we asked them how 127.0.53.53 affected their systems. We actually even linked them up with the vendor that wrote the software causing them to query the -- do the bad queries. So while certainly we don't know everything, you're absolutely right on that, there are some things that we do know that that just unfortunately we can't write all the details in the report yet. The systems fail in different ways, some more noisy than others. One of the things that will really help is when some of the vendors start to build in instrumentation into their software to better respond to getting this special IP address.

So right now it's an opaque failure, but if vendors started noticing and raising alarms when they see this special IP address, it will become a lot more clear.

ELEEZA: Okay. This comings from the chat. Sorry.

I would like to raise one case and hope to get some clear guidance for such case. Suppose a registry operator has operated a unique local

resolvable zone for years and you detected that a lot of SLD queries that might cause collision in your analysis. However, the operator now applies for a new gTLD. They clearly express their intention to it get their original registrant to get the original SLDs under the new TLD. That means some collisions detected by you are not going to happen. If a new gTLD applicant agrees to have such contract terms, would it be possible for this TLD to have alternative mitigation methods?

JEFF SCHMIDT: As Francisco said earlier, we said in our report and as the VeriSign CBA study also very clearly indicates, it's really difficult, even when you really think you're sure where a query is coming from, it's really hard to be sure.

So there's no recommendations in our study along the lines of if we think we know or if we're responsible for a query, have some different approach.

FRANCISCO ARIAS: Last question.

JIM BASKIN: Yes, thanks. My name is Jim Baskin. An earlier commenter today mentioned that the full report has not been released yet, and you responded to that but I want to just reinforce that without the full report so that we can actually see better how you came to some of your conclusions, it's difficult for us to have a comment period that ends at the end of this month.

We really need to be able to comment on the full report once that comes out.

And the longer it takes before we get to that full report, the less value there is in the comments.

And I understand the reasoning for not producing the full report. There's a vulnerability that has to be looked at, that has to be -- some solution has to be worked out before you publish all the data. But nonetheless, we need the data if we're going to really understand the report and we must be able to comment once we have that data.

Thank you.

JEFF SCHMIDT: Thank you. Yeah, absolutely. We look forward to releasing the full report as soon as we can.

One thing that I'll mention, though, on that front is that the more delegations -- or the farther down the road that we get, the fewer instances where we're able to use wildcarding for the controlled interruption. And the wildcarding is a superior approach. And so it's valuable to have this out there and being considered in advance and as quickly as possible.

Thank you.

FRANCISCO ARIAS: Thank you very much. With this we close the session.

[END OF TRANSCRIPTION]