BUENOS AIRES – Security and Stability Update
Monday, November 18, 2013 – 15:15 to 16:45
ICANN – Buenos Aires, Argentina

JOHN CRAIN: Okay, ladies and gentlemen. Welcome to the security and stability update. My name is John Crain. I have a new title this year. For those of you who know me, I change my titles regularly. This week I am the chief security, stability and resiliency officer at ICANN. If you don't know how to say the acronym, we just say Cicero. Anybody who knows anything about philosophy will know who that clown was. So go look him up.

I'm just going to moderate, so I don't have to sit up there and have things thrown at me.

So I want to introduce our first speaker. And that is Jeremy Rowley from DigiCert. And he's going to talk about some of the CA issues and browsers, et cetera. Jeremy, you want to do a quick self intro?

JEREMY ROWLEY: Hello, I'm Jeremy Rowley. And I work for a company named DigiCert. We are a certificate authority, and we do digital certificates, obviously. We were asked to present today about what's going on in the CA world, specifically what's going on with the CA/B forum and on the public mailing list, what's being

discussed there and what improvements are being made. So I've got some slides here.

Okay. Yeah. So talking just a bit about what the CA/B browser forum is and what we're doing and that interplays with ICANN a little bit.

So the CA/B forum is a group of browsers, CAs, and other interested parties who are looking to improve online security by raising the bar on digital certificate issuance practices and CA operations. You have pretty much all of the major CAs and browsers represented there. You have membership of Mozilla, Microsoft, Google, Symantec, DigiCert, and Go Daddy and all the rest.

Interested parties -- any of you guys are invited to participate on the public mailing list. And you participate in working groups by signing our IPR, which is a RANZ (phonetic) policy. However, those are non-voting rights in the forum. We produced over the years several standards that are -- controlled the use and issuance of certificates. Baseline requirements control the use of sale certificates. EV guidelines dictate how certificate authorities provide EV certificates, which are the enhanced validation certificates.

We have network security guidelines that provide a minimum standard for how CAs secure the infrastructure. EV code signing guidelines and code signing based on requirements are -- the code

signing requirements are pending. They're detailing how we issue code signing cert.

Yes, you bet. Okay. Sorry about that.

So new developments in the CA world. There's been a few technology changes that have occurred since, I think, you guys all last met.

First, the 1024 bit certificates have been deprecated. You have to move to 2048 bit certificates. If you're still using 1024 bit certificates, you're going to get an email or a notice from your CA that you have to move off of that and at least go to 2048. SHA2. Microsoft recently came out with an announcement that they're getting rid of SHA1. SHA1 will no longer be trusted in their browser. So, for code signing certificates, that means you have to get rid of all your SHA1 certs by 2016. And in 2017 all SSL certificates will have to be upgraded to SHA2.

As most of you know, we're working on phasing out internal server names. That's coming up in 2016. However, we are phasing out many things early in response to the 120-day rule, which is that 120 days after a new gTLD is delegated, all certificates issues on that gTLD will be revoked. So they're no longer valid.

We're also pushing OCSP stapling. Mozilla has announced that they're looking at pushing OCSP stapling as their only source of revocation testing and also pushing using SSL everywhere.

Other things we've actually rolled out with several customers and people, CAA, which is a process of letting you select which certificate authority is authorized issue certs on your domain. And certificate transparency is being deployed by Google. And there's -- at least DigiCert has been working with them. And you can get actually your certs logged now, if you want to opt into that. So there are working log servers, and certificate authorities are participating.

So next slide. So some of the projects that are going on right now of interest at the CA/B forum is we're looking at expanding the EV certificate arena right now. EV certificates are kind of limited on who they can be issued to. Companies that meet set criteria and some of the discussions that are going on is how to expand that scope to other types of business entities. And it will include more of an international focus so companies located outside the U.S. will be more readily able to obtain an EV certificate.

Another group is performance working group. That's going to be looking at faster performance for certificates with smaller certificate sizes and the impact on certificates and some of the a best practices profile that CAs can use. We're also looking at talking about the deprecation of long-lived certificates. Some CAs

issues 10-year certs. Those are being phased out to 5-year certs. Those have been phased out to 5-year certs. We're going to be dropping to 39 month certs in the very near future. Next.

So the path ahead: Right now we're all being look to improve online security through researching new technologies and implementing better standards in enforcement. There have been some bumps in the road. For example, there's problems with legacy devices in SHA2 and 1024 bit certs. Japan, in particular, has devices that are noncompliant or that can't process SHA2 certs. So we're looking at what to do there. There's some of the old paradigms. People are resistant to change a little bit, especially where it costs money to implement new software or hardware.

And then, of course, we have the ever-present chicken and egg problem. With any new technology there has to be a first implementer. And there's first -- and whether the browser is implemented or the CAs implemented, somebody's got to be first. And several CAs are looking at how to break that cycle and be the first to implement. I think in the future you're going to see a lot of transparency, accountability, and self-selection of CAs. We're working on ways to provide registry, registrars, and EV server operators more visibility into what certificates are out there and provide enhanced information on domain operators and that information. So I think it's got a bright future. And that's the end of my slides.

JOHN CRAIN: So I want to give the opportunity to ask a few questions now, if anybody has questions. I know it's a highly technical subject. I'm looking at Chris, because I know he always likes to ask questions. Mr. Wright? No? Okay.

Any questions for Jeremy? I'm going to walk with the microphone. Mr. Moss.

JEFF MOSS: Thank you. Hey, Jeremy. So SHA2 -- SHA3, SHA4, SHA512 -- are we moving to the next lowest rung on the ladder as opposed to jumping --

JEREMY ROWLEY: When I say SHA2, I mean, SHA65, SHA512. So we're moving to the next one.

JEFF MOSS: Okay. So you're using SHA2 to reference the whole suite of cypher sizes?

JEREMY ROWLEY: Yeah.

In fact, I think the intermediates that we're doing is SHA512.

JOHN CRAIN: I think I see another hand over here.

ALEX DOUGLAS: Hey, Jeremy, how are you doing?

JEREMY ROWLEY: Hey, good. How are you doing?

ALEX DOUGLAS: Hi. I'm Alex Douglas. I represent Internet (saying name.) Since we last all spoke, there was talk about possibility of CAs revoking certificates faster than 120 days for TLDs other than corp and home. Has there been any discussion in the browser forum?

JEREMY ROWLEY: That was discussed a little bit during our last face to face. However, it was viewed that because -- and, although the 120 days was originally based on the fact that dot corp was such a big concern, it was viewed that by the time we could pass something and get it implemented in the browsers, the 120 days was -- we'd already be, you know, another six months. So almost all of them would have been delegated, so there wasn't a lot of energy to spend time on that, though I could bring it up again.

JOHN CRAIN: On the collisions? Okay. I'm going to take one online question, and then I'm going to come back to the room.

WENDY PROFIT: Question from a remote participant. Applicants that choose not to follow the APD route, when are these expected to be able and proceed towards delegation? How much of a delay are they facing until the collision issue is resolved?

JOHN CRAIN: Okay. I want to save that one until the end. Because, at the moment, we're talking about CA issues. And the next four presentations are -- or at least three of them are going to be about collisions. So I saw more questions in the room. And we will come back to that, the person who asked. Did you have something?

WERNER STAUB: Yeah, just -- Werner Staub from CORE. Question along the same line about whether there could be some faster than the 120 days. What stops the revocation of internal CAs to start now that actually the strings affected unknown?

JEREMY ROWLEY: I think a lot of CAs are doing that. The problem is that the rule doesn't say -- you have to wait -- or they aren't required to revoke

within 120 days. So to be absolutely sure that all the certificates are revoked, you have to wait 120 days. Like I said, I think most CAs are probably revoking prior to that time and are looking at the revocation of those -- the delegated strings. Especially now that the 820 rule is gone, we're looking at revoking those much sooner. But you can't be sure that all CAs out there have revoked those certificates until the 120-day time period has passed.

JORDYN BUCHANAN: Hey. It's Jordyn Buchanan with Google. I guess it's a two-part question. Number one centers around digital certificates for IDNs. So two parts.

The first is I know there's a theoretical spec to make digital works digital certificates with IDNs. Could you maybe educate us a little bit about the actual practicability of that? Are there digital certs that are issued for IDNs on a regular basis? And, if so, do they work in browsers? I expect there would be a U-label, A-label confusion probably as between what the certificate is issued for as opposed to what's transmitted on the wire versus what the user is typing in, potentially.

And the second question is: Are you aware of any internal cert that's ever been issued for an IDN?

JEREMY ROWLEY: I can't speak to the IDN issue, because I'm not sure what -- I'd have to look at what we're doing. And I'd have to look at the spec. But, haven't been following it as closely as I probably should have. But I'm not aware of any that have been issued an IDN.

FRANCISCO ARIAS: Okay. Digital certificates issued by CNNIC certification authority on IDN names on the China TLD. And I believe they have the U-Label in the subject name.

JOHN CRAIN: You can come out. I'm not climbing in there again.

JORDYN BUCHANAN: Yeah, Francisco. There is a -- I don't know if a -- spec is probably an overstatement. But there is a strategy that was developed several years back that enables IDN certs. I don't know whether browsers have actually implemented it. So that's what I was mostly curious about.

And, then, secondly, I was just curious about the internal name, in particular. I know there are digital certificates issued under other TLDs. But I'm not aware personally of any IDNs that have been issued as internal names. And I'm sort of curious if there is any evidence of that out in the world.

JOHN CRAIN: Okay. I think we're going to move on to the next subject in the interest of time. Jeremy is going to be around. He loves you to corner him in the corridors and ask him complicated questions. So feel free to do so. Jeremy.

JEREMY ROWLEY: Yeah, please do. I'm here all week.

JOHN CRAIN: So the next subject we're going to go on to -- and we're going to have a series of discussions on this. And I'd like to try to hold the questions until we've done all of the presentations. Because you might get answers to your questions in the following presentation. And I want to start with Francisco Arias from ICANN staff who is going to talk about gTLD collision occurrence management plans. Francisco, if you'd like.

FRANCISCO ARIAS: Thank you, John. So next slide, please. Previous one. Oh, thanks.

So, just a bit of background, the background slide, please. A bit of background here. The name collision issue was brought to the attention of ICANN by SSAC in -- with the SAC57 discussions back from November last year to March this year. This was to a

specific issue of internal name certificate. It's part of the general issue of name collisions in new TLDs.

Following that there was a series of discussions within the ICANN community and others about what should be done. There was a proposal that was put for public comment on 5 August.

After the first report was commissioned by ICANN, the name collision report developed by Interisle. So in 4 August we put this proposal out for public comment.

And then, after the end of the public comment, we have a revised proposal that was considered by the board new TLD program committee. And it was developed as a new gTLD collision occurrence management plan. Next slide, please.

So what does this plan includes? First, describes were called the high-risk strings. Home and corp are identified as high risk. And so ICANN will defer delegation of home and -- of these strings indefinitely.

And the other element in regard to these two strings is that ICANN will look to collaborate with technical and security communities to continue to study these issues and look for a final resolution in regards to high-risk strings.

Next slide, please.

The next part of the plan is the outreach campaign, what does this contain? One of the objectives of this campaign is to make the public aware of this issue, advise the users are the things that ICANN and network operators are doing to mitigate these issues and, as such, to provide guidance to the network operators what they can do to fix these issues.

We have to remember that the root cause of the issue, of the problem here is that people are either purposely or without knowing using internal names. And those are leaking to the public DNS. And that's what has the potential to cause collisions once new gTLDs are delegated.

And, finally, ICANN is looking to collaborate with other parties and members of the community who have interest in promoting solutions to this issue.

Next slide, please.

Another element of the plan is to have a 120-day period of no activation of names for new TLDs. And so the 120 days counted from the signing of the agreement, there will be no activation of names under the TLD. And this is to mitigate the internal name cert issue as described SAC57. This wait period comes from the -- as Jeremy explained, from the baseline requirements that the SSAC has forecast. This is to allow them to revoke all the certificates related to a new created TLD.

We expect this -- the impact of this measure to be minimal in the launch of TLDs, the reason being there is a series of process that needs to happen between contracting until the TLD can be launched and activate names like predelegation testing and delegation, IANA, the sunrise, et cetera.

Next slide, please.

There is another element to the proposal. It's the name collision response. This is, in case a collision happens, what can be done. This is the last response mechanism once somebody identifies that it's been affected by a collision. So ICANN has set up this Web site that allows an affected party to report an issue related to name collision. And ICANN will act upon this. We are setting up the ICANN NOC, which will have service 7 by 24 and will look to these issues to an email validation and then would like to report to the corresponding registry that has to consider the issue and could potentially temporarily deactivate the name to allow time for the affected party to affect changes in the networks so that the issue is resolved from the root.

Then the plan contains two options for the applicant regarding how to get to the delegation. There is a primary part, which is by applying in the mediation measures required by a name collision required in assessment. This assessment is the result of applying what is called a name collision for a management framework. We

have a section to talk about this later. Then the other alternative is the alternate paths delegation.

Oh, I'm sorry. Can we advance the slides.

Advance, thank you. Yes.

Thank you, yes. And so the alternate path is -- the main feature of this is to look all SLDs seen in the DITL and all the relevant datasets that are available. Next slide, please.

Just high-level, so many of the framework there is a -- another section for which we are going to talk more in detail about this. It's to be developed in cooperation with the community. This is not -- ICANN is not developing something. It's something that we intend to do together with the community and we're starting now.

The framework will include parameters and processes to assess the probability and the severity of impact should a name collision occur. We will specify the corresponding measures regarding the type of collision that is identified. It will be focused on the second level domain names under a specific TLD. Like I said, there will be an assessment pair individual TLD. Next slide, please.

Regarding the alternate plat -- to the litigation, we believe this is a conservative approach, that allows both the progress of the new TLD program without compromising the security and stability. And as I mentioned before, this is a temporary block of all SLDs

seen in the DITL data and all the relevant datasets. The important feature is that it preserves the DNS results that is -- that exists in the domain response from the point of view of the resolver that is making the query in the public DNS. This will be December, so there is being presented now to the resolver without the TLD being delegated. Next slide, please.

Those are important things to consider about the alternate path to delegation. This is based on blocking analysis and in order to be able to apply these mechanisms, we need to be sure that we have a high certainty that the list of SLDs for the TLD do not vary too much from year to year.

So we developed an eligibility criteria, which is basically to look at the comparison year-to-year in the DITL data. Like I said, we have datasets from 2006 to 2013. What we did is removed the 2013 dataset, given that the capture of this data was done after the reveal of the strings, and we identified that there are some queries that seem to be looking at what was going on with the TLD. So we removed that. And we look at the comparison of each year, with other -- for example, 2006 with 2007, 2007 with 2008, and so on. And based on that, we identified outliers, and those outliers we -- we established the criteria for eligibility that they will not be -- that a TLD will not be an outlier in two or more of these year-to-year comparisons and one of those will have to be

2012. Why 2012? Because we are looking at something that is still a current trend. And I believe that's all my slides. Thank you.

JOHN CRAIN: Next slide, please. So I see people looking, and I imagine there's a few questions there. There was a lot of slides. So I'm going to ask you to kindly take note of them and come back at the end of this series of slides on collisions. Okay? So we're not running away. None of these people up front get to escape. You get to ask your questions to them. We're just going to do it in a little bit.

So the next issue is the Name Collision Occurrence Management Framework, and we have Jeff Schmidt from JAS Global Advisors, LLC, who is going to talk to this matter. Jeff, would you like to take the microphone.

JEFF SCHMIDT: Guess I need that. Thanks, John. Good afternoon, everybody. My name is Jeff Schmidt. Very briefly, I run a boutique consulting firm for the last 11 years, JAS Global Advisors. We focus on information security and risk management. Our clients tend to be banks, critical infrastructure, government. We have been involved in a couple of ICANN projects. You may have seen us or seen our name around including tertiary reviewer for initial evaluation and extended evaluation. Next slide, please.

So it's -- it's important to understand where -- where the project that we've been engaged to work on fits in the overall scheme. So there's three big chunks. The -- the SLD block list approach, which is a temporary approach that Francisco talked about, the creation of the framework which is the mechanism for bringing closure to this issue on a permanent basis. And then the application of that framework to all applied-for strings.

So those are three steps. What we're talking about now and the project that we've been engaged to do is number 2, which is come up with a framework that can be applied then to bring closure to all of the strings. Next slide, please.

The -- back up one, please, I think. Yeah. So from our perspective, you know, the high-level objectives, we understand that we need to bring closure to this issue and we know that it needs to be done in a very timely basis. The timelines are aggressive. We -- we will have our report out for public comment in January. And as Francisco mentioned earlier, the March ICANN meeting will be the, you know, official roll-out and completion of the framework.

We need to have a framework that is repeatable, that is applicable to all of the strings, and we recognize that the inputs into our framework are going to come in in various states. So we need to take into consideration that some of the strings may have been delegated already, some may not. We'll have to take into

consideration the block list. Some -- some registries will have chosen to implement the block list and will be operating in that state. Presumably some may not choose to do that and so will be operating in a different state.

So we need to recognize these, you know, various input situations and take them into account in our framework. And again, bring closure and bring a deterministic path to the issue. Next slide, please.

You know, our background is security and risk management. And this is a risk management exercise. You know, our top priority is to get a better understanding of the consequences. This issue has been discussed, you know, fairly energetically over the past couple of months. But the entire focus has been on frequency. There is no concrete situations where there's a -- you know, damage that can demonstrably -- or that can be demonstrated that we're aware of. We are actively looking for experiences that can give us some idea of how to -- how to start thinking about the consequences. A lot of work has been done on frequency. We think the next level is to understand what the potential harms could be. We're also approaching this understanding that not all collisions are created equal. There are a lot of collisions, there are a lot of collisions in every name space that we deal with. And not all collisions necessarily lead to harm and not all collisions

necessarily lead to the same degree of harm. So we need to better understand that. Next slide, please.

So the next couple of slides, just kind of giving you a sense of the parallel paths that we're working on internally to get our arms around this issue. Again, a lot of work has been done understanding the frequencies. We were trying to structure that into a taxonomy to get a -- get a more structured way of discussing and categorizing the -- both the strings and the types of queries that we're seeing. The datasets -- for those of you that have looked at the data, you see that the datasets are absolutely dominated by garbage, by random -- apparently random, algorithmic, those sorts of strings. 30 to 40% of the dataset, depending on how you calculate it, is -- can be explained by these strings. We've been doing some -- some work to try to understand those. The strings that Chrome generates, the so-called Chrome 10s, have been getting a lot of attention, but there's a couple of other patterns which we'll talk about later. We're trying to get an understanding of what is generating those. We know that there's a couple of families of malware and click fraud tools that seem to be responsible for a significant proportion of these strings. So again, understanding where they're coming from helps us understand very large chunks of the dataset.

Looking at the effects of collisions in previous delegations, so there have obviously been a lot of delegations, Gs and CCs. Some of them have data, DITL data in particular, that existed before and after the delegations. So strings like XXX, Asia, there's a couple of CCs where we can actually look and understand what -- what happened before and after. We also want to understand what happens inside delegated name spaces with respect to collisions. Collisions are not limited to TLDs. They also occur, obviously, in delegated spaces.

One of the things that we found helpful in other projects when examining kind of long risk chains, you know, in order for this really bad thing to happen, these seven other really bad things have to happen first. Those sorts of risk chains, you tend to be able to get a sense of how probable they actually are and be able to compare them to other complex risk chains with Monte Carlo-type analyses. So we are going down a path to help compare the actual risks associated with name space collisions, with other risks that IT departments face every day. Everything from hardware failure to com link line failure. Risk of patching is something also very interesting from an IT operational perspective. Every time an IT department deploys patches, you know, there's clearly a risk that something catastrophic could happen. They manage that risk and understand that risk. And so we're trying to put the -- put these in some kind of context. Next slide, please.

The -- The -- Yeah, that's the right slide.

One of the things that we really want to do is get outside of the ICANN sphere into the I.T. operational folks that are dealing with these issues, and will be dealing with these issues on a daily basis. So we're trying to get as -- you know, as broad of a call as we can to help ask some questions, help understand what the implications may be, what the actual operational impacts would be, how they would respond, how they would know, how serious they think it is, et cetera.

So we will be developing a survey and we'll ask for everybody's help to -- to complete the survey as well as to forward it to as broad of a community as you can.

We do want to also ask for specific case studies that anybody with operational experience with respect to namespace collisions, we'd ask that they reach out and give us some additional data.

And then finally, at the end of our rather long task list here is the -- developing the options to actual deal with the collision issue. So we expect that, you know, there will be a menu of -- you know, of possible mitigation and remediation options, and, you know, our objective is to give those meat and certainty and deterministic applicability as well as match what strings in what buckets in what situations need to apply which mitigation measures in order to bring closure to this issue.

Next slide, please.

So a couple of things we're specifically asking for. We're asking for case studies. I mentioned this earlier. Please, if you have experienced a namespace collision or know anybody that has experienced a namespace collision of any type, variety or sort, we would like to know about it.

The -- To the greatest extent possible, we're going to, you know, have a -- have an open report, but if people do need to protect names, operational details, et cetera, we're so interested in case studies that we'll do almost anything we can to get that data.

Next slide, please.

The -- I mentioned that we will be sending out a survey. We'd appreciate your help.

One more slide, please.

We'd appreciate your help in sending that along. Wow! Halloween color scheme.

So another request for -- for participation is information on these algorithmic patterns. The top five patterns are listed there, aside from the Chrome 10s. If you know anybody or have any experience with software that may be generating these sorts of queries, the queries listed on the screen here plus the Chrome 10s explain more than 30% of the data set, more than 30% of

collisions. So it is really important to understand where these are coming from.

Next slide, please.

I mentioned the survey. Next slide, please.

Data. So we're interested in data. One of the nice things about the DITL data sets is that it is available for anybody to use. We like research that can be reviewed by others, repeated, embraced and extended, et cetera. But we recognize that there is -- there are limitations to the DITL data, and so we are asking folks that have data, I may have already asked you if you have data and have an interest in this, please reach out to us.

Next slide, please.

In an effort to, you know, keep the project open so that everybody has a sense of where we are, where our thinking is and be able to provide feedback and commentary along the way, we are using the DNS-OARC collisions list as the focal point for continued list-based discussions of this project. So the DNS-OARC collisions list started with a few researchers that were looking at this issue. It quickly became kind of the spot and we're going to continue using that list throughout the process.

If you're interested, please subscribe. Please review the archives. We will be on it, we'll be on it actively, and we'll like to direct list-based conversation there.

Next slide, please.

We also maintain a couple of pages on the DNS-OARC Web site that are public. It's mostly the technical nitty-gritty, but please know that this location is out there. We're actively maintaining it. Again, it tends to focus on the more technical things. Our source code is out there.

We've made our data sets available to other DNS-OARC members. The data sets start very large, and tend to get whittled down fairly quickly, and so not having to repeat that effort has been helpful.

Next slide, please.

In an effort to keep the project open, we encourage feedback and participation. Find us in any way that you can. Email us. Participate in the collisions list, participate in the other DNS-OARC forums. Stop us in the hallways. We're excited to hear ideas and commentary.

And one more slide.

The -- Actually, there's two more slides.

Our draft report is expected in January. We will look at the -- we will look at the public comments carefully, and the March 14 -- or the March 2014 Singapore meeting is when we expect to have closure.

One more slide.

We are also blogging on domainincite again in an effort to keep the project open and try to expand the community of folks that are aware of and tracking. We expect about five blogs over the course of the project on single issues. And we will be participating in the -- you know, in the commentary there.

I think that's it.

Thank you.

JOHN CRAIN: Okay. Thank you, Jeff. Once again, I know it's a lot of information, but we've got two more slide sets, both of which -- I think one of them is one slide and one of them is -- I think Patrick is like me. He doesn't really like PowerPoints, so he probably won't have any.

So with that, I'm going to move over to Dave Piscitello, who's got a fancy new title just like I do. He's the vice president of security and ICT coordination at ICANN, and works in the SSR group.

Dave, would you like to take your subject?

DAVE PISCITELLO: Thank you, John.

I'm going to approach this subject from a slightly different perspective, and the goal of the work that I'm going to describe,

which is still ongoing, is to provide awareness to; in particular, to network and operators and I.T. administrators, and to also assist those operators in understanding how to deal with new TLDs in those circumstances where the namespaces that they use are the sources or causes of leaks, which we call collisions on our end. From the perspective of a network that is connected to the Internet that is using a private namespace, these are leaks. These are not intentional in most cases.

If they're intentional, they're intentional by virtue of complacency more than by intentional configuration.

So together with a subject matter expert, I have been working with a gentleman by the name of Paul Hoffman, who is very well-known among the Internet firewall and enterprise network community and security community as having significant expertise in private namespaces, in firewalls, in managing security from an enterprise perspective.

And what we are attempting to do is to begin by describing for network operators what this landscape is because while this is a very, very interesting and often-discussed subject here in the ICANN community, you would probably be very surprised at how little this may be of importance until raised directly to the I.T. operators. They have smaller budgets than any of you, most likely. They are always struggling for time and prioritization of their work. And DNS issues aren't always at the top of their list.

So one of the things that this report that Paul is working on, and I have been enthusiastically commenting on, is to describe the problems of these namespaces and describe the situations that organizations may encounter when they use internal namespaces, and when those namespaces leak into the global DNS.

Now, there are three kinds of scenarios we will describe. One is where names are leaking into the global DNS because -- if they are branching off the global DNS. Another is where organizations are using their own private TLD as the root of their namespace. And a third is when organizations are relying on or are still using search lists.

So for each of those scenarios, what we will be doing is identifying exactly why and under what circumstances the leakage is occurring. And then for those, we'll identify or recommend measures whereby operators can mitigate some of the problems for private namespaces that are being used.

Now, in many of these cases, the recommendation is to consider using fully qualified domain names instead of short, unqualified names. This is one of those scenarios where I often compare it to people living in a hurricane belt or living in a 100-year flood plain. Many people have been configuring in networks rather than complacently and they have been using these names for many, many years. If you go back and look at DITL data even when the SSAC was looking at the DITL data in -- I think it was 2008 or 2009,

many of these names were already being leaked and others were being leaked as well.

So the goal for us is to explain you have been on borrowed time. You have been using these names. The probability that you can -- you can have a collision is going to increase as we introduce new TLDs, but the opportunity for you to go and correct this once and forever by using fully qualified domain names exists, and here is how we do it.

And the report, honestly it's my fault it is not published this week. I will take full blame because I got very excited about getting it right and having all the details, and every time Paul gave me a version, I said, gee, shouldn't we also tell people this, and, gee, shouldn't we also tell people this.

So I'm very optimistic that it will be comprehensive. I will apologize to you because I'm the one, you know, who's responsible for it not being here, but I'll hold myself accountable. And I hope that when we do get the paper that you will find it very, very important reading.

Much more importantly, this paper is not actually for most of the people in this room. This paper is for people who actually have to contend with the networks that are causing the leaks and it gives them very, very good advice in how they can fix the problem once and permanently.

Thank you.

JOHN CRAIN: Okay. With that, only one more presentation, and then you get to ask your questions.

I'm going to be rude and pass out the translation devices to the speakers while Patrick is doing his presentation so that you can ask questions in your own language, if you want to, because we do have translation in here.

Patrick, would you like to go ahead with your topic?

PATRICK JONES: Sure. Thanks, John.

So I'm going to talk about sort of the umbrella of activity that falls within security, stability, and resiliency for TLDs in general. It's not so much related to name collision but I think it sets a pointer to some of the conversations that we may be having at the Singapore meeting and the London meeting and the meetings beyond. So at some point, name collisions may stop being the topic of the day and we'll begin to talk about some of the other areas of activity and concern within TLDs.

So I'm going to talk about risk and incident escalation processes. This is not a new area for ICANN. This is the umbrella of activity that some of this is lessons learned out of RegisterFly from 2007.

So for those of you who aren't familiar with that name, that was a registrar that had approximately 2 million names under management. The others in the room can correct me about how many names exactly that registrar held. But when that registrar collapsed for business reasons, there was a need to transition those registrants to new registrars and to handle that in an orderly way.

ICANN and the community learned quite a lot about the need for a registrar and registry data escrow, the need to have an incident management process, the need for a coordinated vulnerability disclosure, ways to contact ICANN registrars, registries and others in the space. And that led to enhancements that went into the Applicant Guidebook, that went into the registry and registrar agreements. Quite a lot of work has been discussed in the community since 2007, all -- you know, primarily from that event, but also it was just timely that those security and stability issues were discussed within ICANN and by the community.

So the work that we're doing to look at this unified approach for TLD risk and incident management is to have a process. Right now it's an internal process. I think at some point, this will be a process that will be more readily available for the community to read and follow, but this includes all the steps from the evaluation of TLDs, the delegation process, our coordinated vulnerability disclosure information that ICANN has published, Service Level

Agreement monitoring process, and then everything around the emergency back-end registry operator program.

There's another step that has not received as much attention in the community as perhaps it should and that's the need for an undelegation process for TLDs. The ccNSO has the document that is currently out for public comment out of the delegation -- it's the Framework of Interpretation Working Group, I believe, on what types of categories or scenarios might be invoked for revocation of a TLD.

Within the EBERO process and within other discussions at ICANN, this has been looking at potential scenarios, what scenarios might qualify for undelegation of a TLD.

I think there's very few examples where that might be invoked, but that could include where a TLD poses significant cybersecurity, stability, resiliency harm to the DNS. Something that's learned post delegation.

Another example may be a TLD with government support that loses that support, and the government requests that the TLD be removed from the root.

Another example may be an operator with legal rights in a TLD label requests removal of that label from the root.

So this could be an example where a brand owner decides that they don't want to be a registry operator anymore, and so this

may be an example where there's an orderly wind-down of a TLD, and quite a bit of awareness for the community that this is happening.

The final example may be that a TLD goes into the EBERO process and there's really no interest in the community in having it be transitioned to someone else. And so then it needs to go through a wind-down and removal.

We'll be looking forward to having further conversation with the community about these topics. I think this could be a topic of discussion in Singapore.

And with that, I turn it back to John for questions.

JOHN CRAIN: Okay. Thank you very much.

So we're going to open up the floor for questions.

I do ask that if you've got 20 questions, please try to only ask two so that others get a chance.

Feel free to speak in your own language, especially if it's one of those written up there, because we do have translation. Everybody has headsets, I hope.

It looks like we're going to set up a queue. That's very posh. Okay. That means you have to get out of your seats to ask a question.

But before you get out of your seats, I believe we have a question from earlier online. And I'd like to give the online people a chance to go first.

REMOTE INTERVENTION: Thank you, on behalf of remote participants question, from Reg from Minds+Machines. Question: If NIC appears on the list indicating that there is a name collision issue there, why is NIC allowed in the gTLD and why is this risk more acceptable than the risk of any other name collision?

JOHN CRAIN: Okay. I believe Francisco is going to take this one.

FRANCISCO ARIAS: Yes, thank you.

Regarding NIC -- and shall I add there is not really NIC. It is WHOIS.NIC.TLD because that' the only name that's required to be active. That is a fair amount of TLDs that have that in their list. And what we did there is a balance between risk and usability.

In this case, with WHOIS.NIC TLD is important because the WHOIS service is being used for, for example, Jeremy to my left in the CAs, they need the WHOIS to be up and running, and they need an easy way to find so they can help resolve the internal name certificates so they can know, for example, who has been allocated a new name under a new TLD so they can know if it's a holder of an internal name certificate, it's the rightful owner of the name under the domain name in the public DNS.

JOHN CRAIN: Okay. We have one more online question that I want to let go first.

FRANCISCO ARIAS: Sorry. I forget about something else.

We also have, of course, the name collision response mechanism which can be used in case there is an actual issue with WHOIS.NIC.TLD, so that that can be reported, and then of course we will act upon that.

REMOTE INTERVENTION: Thank you. The next question from remote participant Michael Flemming. Question: Is there an overall timeline for the name collision mitigation process? Perhaps one that can be posted on the ICANN Web site?

JOHN CRAIN: Francisco is popular.

FRANCISCO ARIAS: Thank you.

So the timeline is regarding the alternate path. We are already in the process.

We finalize publishing the alternate -- sorry, the list, the (indiscernible) list we block for all the TLDs that are find illegible. That's already there.

What we are doing now is starting the name collision course management framework, and the timing for that is to have a draft for public comment by January, and then the final framework by March.

Then the next step will be to apply that framework for -- to each TLD.

We don't yet have a timeline on that. We think it will be very quick once we have the framework.

Thank you.

JOHN CRAIN: Okay. With that, Mikey, if you would like to state your name and start the questions.

MIKEY O'CONNOR: This is Mikey O'Connor speaking into the mic until the mic lights up. It's not lighting up yet, but maybe it will get -- Can people hear me okay? What? Once upon a time, 33 purple birds were sitting on a curb, chirping and burping -- I can hear myself. Oh, good.

JOHN CRAIN: Okay. This one does work.

MIKEY O'CONNOR: Okay. Take two, this is Mikey O'Connor, and actually, mostly what I want to do is tell you a bunch of stuff that I like a lot. This is much better. And I really appreciate all the work that all of you are doing. So this is not the accusatory, put-you-in-a-box question.

This is more the question of we, the ISP constituency, have a lot of members, as ISPs, and, in turn, a lot of customers who could really benefit from a lot of this information really fast. You know, as soon as you develop your stuff, great. To the extent that you can develop your models and your study and make them shareable,

you know, to the world, great. The more the merrier. The more Web sites there are.

And so the one thing that I'm sort of coming to the question on is -- and it's really kind of an extension of the online question, is is there going to be sort of a name collisions portal that, you know, everything you want to know about names collisions, all the way from how do I figure out how I've got them to how do I fix them, something like that would be fantastic.

So that's -- that's the one thing.

And then sort of the supertactical request is the ISPs are meeting tomorrow at 2:00, and if there was sort of a channel into this process -- You're on my agenda. Cool.

So then the last one is if there's a steering group that's sort of drawing the communities participation part of this together, I'm sure that somebody from the ISPs would love to be a part of that. You know, both your project, but then all these other ones, too.

And with that, I'm done.

JOHN CRAIN: So I'll answer your first question, and that's if that's what you're telling us is needed, a portal where we gather this, I'm sure we can find a way to do that. And you can work with us on it to tell us if it actually meets the needs of your constituents.

MIKEY O'CONNOR: Well, I was trying to build one and I don't want to do that. So anything that's being done is better than mine.

JOHN CRAIN: That sounds amazingly doable. I like doable things because then we can say look how good we are. We did something.

Okay. The second question, I don't know if anybody from there wants it or -- I think you said were you involved already.

MIKEY O'CONNOR: And then the third one is sort of a steering committee, just a desire to help. You know, because what I see is lots of things going on. Just make sure that it's kind of coordinated. That would be cool. Thanks.

Great job.

JOHN CRAIN: I did say two questions, Mikey.

ALEX STAMOS: Hey, Alex Stamos, Artemis Internet NTAG, I guess.

So my comment, of a comment and a question. I want to thank Jeff for taking everybody's input, and, also, I think what's

especially awesome here that I've never seen anywhere else -- sorry; I'll slow down -- I haven't seen anywhere else in ICANN is that you guys are publishing all your source code on github, which I think is a huge deal because in source code veritas; right? And I would suggest that whomever is doing the final calculation, I don't know if that's JS or if it's staff, like first we would love Francisco to see the source code that was used for these lists. It seems to be more complicated. We thought it was going to be just a unique on 2006-2012. It seems it's more complicated than that so it would be interesting to see what code was used so -- not so much that we can verify but so that we can understand what happened, and so that should definitely happen for the next one.

And then so my question is, the timeline's a little fuzzy as well as responsibilities to us on the outside. So it sounds like January, JS is going to come out with a report. Jeff referred to providing a menu of options. So I assume that means you're going to have a bunch of different options to be chosen from.

Is there a comment period? And then who makes the choice? Is it staff? Is it the Board? Or are we going to end up -- Because it seems to mow we're ending up in this little Ping-Pong game where the Board said they cared a lot, they shot the ball over the net to Interisle. Interisle said, yeah, there's a problem but we're not going to make any judgments. They hit it back to the Board. The Board said We don't want to make any judgments. They hit it

to JS. Are you going to hit the ball back and say we're not going to make any judgments and then we go through the cycle again like through, like, London or -- Yeah.

JEFF SCHMIDT: So let me take the last part of that. I think it is incumbent on us in our report to bring closure and deterministic outcome to this issue. Kind of philosophically, our report needs to do that. So I don't believe I have a ball to ping at this point. The ball has stopped pinging.

So our report -- you know, the way I -- the way I see it is we'll provide, you know, several mitigation packages, for lack of a better term; right? Probably some number less than five. I don't know the number. Relatively small number, and then some kind of deterministic way to map strings and states into those.

So, you know, I've got dot Jeff. Dot Jeff needs mitigation package 3, you know. That's the way I see it.

So it is -- it is deterministic. Everybody will know exactly what to do.

ALEX STAMOS: Okay; great. And the idea is that the Board would then accept or not accept that as an entire package after a comment period?

Does anybody know what happens after Jeff's report gets spun out?

JEFF SCHMIDT: Now I have to turn it --

ALEX STAMOS: Is that the idea, Akram? The record will note Akram says the Board will accept the whole thing. Okay; great.

JOHN CRAIN: Does Akram want a microphone?

[Laughter]

AKRAM ATALLAH: You said it right. So after the report comes out, there will be a public comment period. After the public comment period, sunrise comments, we put the report and comments in front of the Board. The Board might accept it as it is or they might actually consider changing some of the, you know, comments that we -- you know, that were accepted or not accepted. The Board does what the Board wants to do, and then eventually we will be able to provide you with the final way forward.

That's the whole idea.

ALEX STAMOS: Thanks. And thanks, Jeff, for putting yourself in a position that there's going to be a seven-piece blog -- seven-post blog about why you hate puppies. So good luck.

WERNER STAUB: Werner Staub from CORE.

I just looked through the list of labels that have been published for dot Paris, and 16-, 17,000 or so. And it's striking the number of brand names that are on that list, including, in many cases, misspelled brand names.

So it gives the clear indication that many of them are, really many of them are just people simply typing, of course, the brand name together with the word.

And now we don't have any more information, each one of those lines that we have on the list. It just says, whatever, the brand name. That's it.

We already had the data that actually went with this, like in the year such and such it was, you know, with the frequency, like it appeared five times or ten times or once.

Now, we would be able to leverage quite a bit of help from other people if you were less parsimonious in terms of publishing data that has no risk whatsoever, you know, in being published.

So, for instance, would it be possible to update that list and add columns yearly year to that list and say, look, it appeared once, you know, in 2007 and so on? That would be relatively easy to do. And you would give people quite a bit of, you know, of a picture of what we're actually dealing with. And especially help other people help, rather than just say let's wait until the JAS advice comes up with a real snazzy solution.

JOHN CRAIN: I don't think anybody really has a direct answer to that now.

Anybody up there?

Next question, please. Thank you, Werner.

KIRAN MALANCHARUVIL: Hello, my name is Karin Malancharuvil from MarkMonitor. I am -- As a representative of brands, I am obviously concerned about the number of exact matches, trademarks that are on the collision lists including some very large trademarks. Google, Microsoft, Facebook, Bing. And so my question is how long will the TLDs -- or the registries be required to keep these names on a block list?

And then what will happen to them once they're released? Will there be a sunrise period? Will there be claims? And will the RPM requirements be robustly applied at that point? And how will that happen?

Thank you.

FRANCISCO ARIAS: I can take the technical part, but I will defer to someone else on the sunrise question.

So on the technical side, the requirement is temporarily block those names until the name collision -- the assessment is delivered to the registry. Once the assessment is delivered to the registry, the registry will be required to implement the mitigation measures per type of collision.

And when those mitigation measures have been implemented, then the name can be released. That's the plan.

Now, the other part of usual question I believe is whether there will be a sunrise on claims for those names. For that, I am not qualified to answer. I'm looking to my colleagues to see if someone can help here.

If not, what we can do is come back with an answer to you later.

KIRAN MALANCHARUVIL: Well, certainly we anticipate that in order for -- to fulfill the requirements that you're speaking about; that it would be, then, long past the sunrise period and long past the 90-day claims notice period.

So just register now our concerns, if there's no answer yet. Register our concerns that we would like those names to then be subject to all of the rights protection mechanisms that they would have been subject to had they not been placed on a name collision block list.

FRANCISCO ARIAS: Understood. Thank you.

JOHN CRAIN: And that's good because we do obviously scribe everything said here so we're getting down these concerns.

Thank you.

JORDYN BUCHANAN: Hi. Jordyn Buchanan with Google again. Just one quick note on that last topic. The block list may be a good block list to block delegation. You might still allow registries to perform registration, and then people could still protect their brands during sunrise if they wanted to.

Two questions, since I'm limited to two. First one is for Francisco. I want to understand a little bit better the 25 names that were not eligible for the alternative path delegation. If I understood you correctly, and I hope I didn't, you said if any -- if there was any

year from 2006 to 2011, and then comparing that to 2012 there was, like, a big increase, that wouldn't be eligible for delegation.

FRANCISCO ARIAS: Yes.

JORDYN BUCHANAN: Or the alternative path.

FRANCISCO ARIAS: Or know.

[Laughter]

Let me try to explain it. The comparison is year-to-year, back to back, only consecutive years. So 2006 is compared to -- the other way around. 2007 is compared with 2006, and we look at the increment. And so if a string is found to be an outlier in any of those, there will be one match. And the other is they have to also be an outlier in the 2012 comparison with the 2011 in order to be considered an outlier. That would be the minimum criteria to be considered ineligible.

JOHN CRAIN: So another way to look at that they have to appear in at least two years, and at least one of those must be 2012, the most recent data set.

JORDYN BUCHANAN: Right. But I mean just hypothetically, like, dot snapchat had been applied for. You might see that trending up really fast recently because it didn't exist a couple of years ago. That would still be caught in your analysis right now?

JOHN CRAIN: No, because it wouldn't be two times. If it's only recent --

JORDYN BUCHANAN: Snapchat is like two years old; right? So --

JOHN CRAIN: If it only appears in 2012, it's not in the list. It must be at least two years.

JORDYN BUCHANAN: I will move on to my second question, which may be more a comment, but maybe just to get a better sense from you guys. There's talk about case studies and informational materials. There's a lot of talk about reaching out beyond the ICANN community. But I get the sense that a lot of us within the ICANN community are maybe not sort of holding up our end of the bargain in terms of helping understand these issues. There are, in fact, a lot of big companies that are concerned about name

collision that are present and knowledgeable about the topic that I haven't seen loft data coming from there. There's ISPs that come here. I haven't seen a lot of data from ISPs about the scope of the problem.

One thing that we're doing internally at Google is we're in the process of doing a test delegation of every one of the new TLDs to see what breaks, if anything; right? And we're hoping to publish a paper about the results of that relatively soon.

But that seems like all the companies here that are in this room and paying attention could all do that and try to figure out whether the theoretical awful things that might help actually happen, and give you guys some -- you know, both a procedure that could be used, and maybe we could work with you to help document that, as well as to much better understand the potential risks. Instead of just talking about what might happen, the bad things that could happen, be much better to try to get a better understanding of what would happen when these TLDs are launched.

DAVE PISCITELLO: So the methodology that you're applying by creating a new TLD for all the -- you know, all the TLDs is a variation of what we're actually talking about in the name -- name collision identification paper that Paul Hoffman is writing. We believe we have a little bit simpler way to do that, which has much more to do with logging

and monitoring the existing private name server and the firewall, or any perimeter devices that are forwarding and are able to log DNS traffic.

So we're in the same -- we're in the same wavelength as you are, and I think that's encouraging to hear that you're doing this. And possibly it would be useful for us to try to find a time off-line with Paul to say, well, you've done this and where are you, and then we can compare the rest of the methodology.

I know one of the things we have been talking about is eating our own dog food and implementing this ourselves so we can actually see how -- you know, how much work it is and where the work becomes somewhat troublesome so we can add that to our advice. Because I believe this kind of advice is not static. I believe this kind of advice is adaptive.

As people find how to do this, often people find ways to do it better. Somebody is willing to post a script and says, by the way, if you use this script, you cut a week off of your time off, as an example.

I think what you're doing is the right idea, and we have to have more experimentation on the operation side inside enterprises and even small, medium businesses.

JORDYN BUCHANAN: Thanks. That's really helpful. I'll just briefly say, I'd love to hear if anyone else has thoughts about how we get these people, everyone here, engaged in that process as opposed to just relying on ICANN to solve this, because ICANN, frankly -- like, you can speculate about what other people's infrastructure looks like, but clearly it would be better to have people with the infrastructure doing this as opposed to you guys doing it.

>> Yeah, and Jordyn, to the extent we can, we're tapping our client base and trying to gather data from them in tabletop mode as well as some technical instrumentation as well.

DAVE PISCITELLO: So just one more point. Actually, if Google were willing to write a paper that describes their process and explains the results, that's exactly the kind of community paper that we would -- we could put on this portal that Mikey recommended. And I think that having a name like Google associated with a paper that identifies a remediation goes a long, long way for people to believe that there's a credible way to do this.

JOHN CRAIN: Next, please.

STEVE DELBIANCO: Steve DelBianco with Net Choice.

I have some questions for Jeff. Near the end of October, I helped organizing an event in Washington, D.C. on collisions, and all we looked at was ICANN's management collision occurrence management framework and spent an entire day thinking about how to make it better, more responsive.

And I learned a lot that day. I learned from many TLD applicants who attended that the alternate path is not so good, really brute force way of blocking a bunch of SLDs that ought to be eligible for registration.

So I think that your primary path that you're working on is going to be very popular. And we're going to need to get it quickly. So I look with interest to see the workplan that you came off with after you won the project.

And I believe it's missing two things. And, if we don't get them in there, it could extend by a lot of time putting them in. In ICANN's board resolution, it wasn't just the mitigation elements in your workplan, but also the consideration of other relevant data sets, which I didn't know if I heard you mention that you would work that in, but, more importantly, the work of the SSAC and the SAC62, which I'm sure you've read by now. And Patrick is right here. But it's got some excellent sections in there on trials. And trials came from the ICANN original proposal. But I didn't see that you had studying a trial as part of your workplan. And, if you did,

Patrick and the SSAC have given you a roadmap of questions, considerations when you design trials. And I think they have four different kinds of trials and maybe more than are needed. But those need to be designed so that we can all react and comment on them. ICANN needs to pay you or pay somebody to build that into the framework. The framework isn't just the Monte Carlo simulations. But, once you've determined the risk category of dot kids, say, and you map it to something -- you say dot kids, what you really need to do is you need to do a trial using services and applications. Well, it will be ridiculous if we haven't designed appropriate trial method by that point of time. Were you thinking you'd do those elements as part of your project?

JEFF SCHWARTZ: So is your question whether trial delegations are on the table for a remediation plan?

STEVE DELBIANCO: Not just on the table, Jeff. But would you, as part of your deliverable to us, be able to design trial delegation processes and paths? In other words, how the trial is set up; who runs it; what are its criteria for how long do you run it; what do you do with the answers and information you get back? Because that becomes other relevant data set for not only that TLD but others. So it's more than just you putting the words "consider a trial." No, it's

the framework should consider the entire design trial with a lot of the SSAC questions answered.

JEFF SCHWARTZ: So, I mean, with specific request to trials, they're certainly on the table and out there and have been suggested in various forms. In fact, there's a lot of very good ideas for potential mitigation paths in the public comment period, for the initial. Other people have put data and put ideas out there. There's the ad network Internet draft. I consider all of those on the table for consideration, you know, as a potential mitigation measure that could very easily appear in our report, if we thought it was appropriate.

STEVE delBIANCO: You said there was a package approach. If dot kids was part of my TLD; and, based on my risk profile, my Monte Carlo simulations, you put me in package two, risk package two; and risk package two has within it a trial of a certain kind, a certain duration with parameters in it, that has to be baked in if you're really going to give me a packaged solution.

JEFF SCHWARTZ: Yes, that's right. If a trial delegation appeared in package two, it would be incumbent on us to define how that would work; if there is data collection; what the data collection looks like; what is good; what is bad. Yes.

STEVE delBIANCO: When will we, the community, then get a chance to see if trials are part of the project work you'll be delivering? January?

JEFF SCHWARTZ: Trials are a part.

STEVE DEL BIANCO: They are a part. Excellent. That's what I'm asking.

PATRICK JONES: I'm going to add that we should be very careful what you're talking about trials in the terms of something that Jordyn was describing as sort of a honeypot network versus a real live delegation in the root zone. We should be really careful about that.

STEVE delBIANCO: So, Patrick, I'm trying to be careful and use the words that the NGPC used. So the word "trial" showed up in the NGPC resolution and document. And the SSAC used the same word "trial" and described in several pages the kinds of trials they had in mind. I'm using it in that spirit of the word, not any other.

Those are both ICANN definitions. That's what we should all use.

JOHN CRAIN: Okay. Thank you for that. In the interest of time, I'd like to get the next gentleman a chance. I'll either cut the line off right either before or right after Mr. Wright. I haven't decided yet.

JONATHAN ZUCK: Jonathan Zuck from the Association of Competitive Technology. I work primarily with small businesses, who I think are the ones who I think might be the most impacted and also the ones that will sort of be the last to find out. So I'm very interested in things like the outreach plan, the mitigation thing that Dave is working on. I guess I'd recommend not using terms like "you're living on borrowed time" in that kind of communication to those folks, potentially, but I'm very interested in that. And I'm also interested in the timing of that. Because so much of what we're trying to do is study, you know, with Monte Carlo simulations and others find this multiplier for probability of consequences versus likelihood, right? So what is the degree to which that's going to inform the outreach plan? Or is the outreach plan going to start right away? And I'm also concerned about very small details like someone coming to you and saying they have substantial harm, how you plan to define that? Because millions of dollars are being invested by a small number of players in these things. And it's going to be very easy for any problem to seem infinitesimal by comparison. So I'm wondering if that's going to be some sort of

absolute thing thing, or is it going to involve dialysis machines not working, as Mikey suggested, as potentially happened with corp.com? What is that going to look like? And what is it going to take for a small business to be able to make a case that a pause needs to happen for mediation to occur? So that's sort of a bunch of little questions. But I'm very curious about the outreach and also that -- you know, how you're going to measure a severe harm or substantial harm to undelegate.

FRANCISCO ARIAS: Thank you. So I will be interested on getting your contact after this session in that respect.

The questions on the time limit, as they mentioned, this document that we believe will be helpful for people that are implementing private names and are linking to the DNS, we hope to have that available next week. So that will be the first thing in terms of the time.

And the outreach that you talk about, we will leave this up. And it will be the primary source of information that we can provide to people that are doing that outreach. And we are -- we are working on preparing a plan to have this outreach ongoing. We have to have it soon. I don't have a date yet. And I forgot what was the other question.

JONATHAN ZUCK: The nature of that outreach, too. Because, I mean, obviously, the philosophy of build it and they will come has not historically been successful for ICANN. So I'm curious about what outreach might mean and the timing of it. But my third thing was proving sufficient harm for at least primary undelegation and what that might look like.

FRANCISCO ARIAS: Yes. So defining what is a significant risk is not an easy thing. Defining what a risk is is an exercise that requires human intervention in order to be able to define what is the risk that's important or severe.

So the way the process is defined is the report is sent to ICANN. And ICANN makes a first check to see that everything matches, for example, the domain name Actelis (phonetic) that is an new TLD and that the request seems correct, then we contact the person that is reporting that all of this went in to make it in a timely manner. I'm talking about minutes, not hours, not even days. And then relay the report to the registries who have to act upon that. And we will follow with them to see where the report is and what they have done.

JOHN CRAIN: And I just want to say on the external communications, we are working on contracting a well-known communications firm and

hope to have those contracts in place in the next weeks. That's in our communications department, not in our SSR department. But they just made me aware that they are working on this. So you will see more soon.

JONATHAN ZUCK: Great. I look forward to hearing that. If there's a business out there that does a million dollars a year in revenue and their business is at risk and, yet, the business on the other side that would be affected negatively by undelegation is a $10 million business, I'd be concerned that that might be the criteria. That's all.

DAVE PISCITELLO: So one of the things that I begin to get upset about when we talk about risk is that we are sitting on the outside of every one of these organizations. And it's a fair amount of hubris on our part to measure risk for an organization.

Now, one of the things I call attention to almost everyone -- yeah, everyone here is that these organizations have, in fact, been doing this for many, many years. And there has been a certain amount of risk in any kind of leakage. And we're focusing the kind of leakage on a name. But, historically, I can tell you, having worked in this -- worked in small and medium business space for years, I've gone into organizations that used to take a Windows

2000 server, connect it to a broadband network with no firewall, with nothing else on. And they were advertising their Windows workspaces in every direction that they happened to have an interface. So, you know, that is an example of where people were in small business 15 years ago. You spin forward today. And people are in small businesses, and they're putting up Word Press sites with default configuration. They're putting up SQL databases behind web portals that allow command level access. This is one of a very, very -- let me finish. And I, by no means, want to diminish that. It's just that, at some point, at some point you have to take responsibility of your own network. And, at some point, you know, especially, I hope, after we produced the document that I've produced, we'll have someone in small businesses or someone like you to help us go to small businesses and say, you know, if you're running active directory or if you're using a search list, you have this business. And, if you do it just this way, this risk goes away forever. And so that's the goal that we have here. I don't want to run around looking under rocks for risks and looking under rocks for --

>> Dave, sorry.

DAVE PISCITELLO: Sorry.

JONATHAN ZUCK: That's a fair point. But one of those things, using it as a root, a substantive root, was actually a recommended practice, not just a coffee make on top of the server kind of problem. Let's not be too -- risk isn't what I was talking about. I was talking about substantial harm that could be used later after the fact to justify an undelegation. So I wasn't talking about risk assessment up front. I was talking about what kind of harms would be needed to justify an undelegation, if only temporary, downstream. That was the question I was asking, which I think is a different kind of question.

JOHN CRAIN: So I think that's a good distinction. But we've only got five minutes, but I want to give at least this gentleman and maybe Mr. Wright a chance to ask a question, too. He can ask me in the bar, if he wants. Go ahead, sir.

RUBEN KUHLS: Ruben Kuhls, dot BR. I have a comment and question. My comment is that trial delegations can be done both on TLD level and on SLD level. So we could figure out some ideas of trial delegation of sample strings and so forth in some of the packages.

My question is about the data sets that are being brought to the discussion. Wouldn't bring more data sets bring also risk of -- that

the plan ends up blocking strings that are now allowed so that would turn into a liability that say, oh, you have this domain that you solved? But it's -- it's now forcing that to be revoked? So is that safe to do? Bringing new data sets into the picture?

JEFF SCHWARTZ: So let me make sure I understand. The SLD block lists now I believe are static. I don't believe there's any plan to change those. Is that -- I -- yeah. So was that your question?

RUBEN KUHLS: Imagine that I have a 1,000 names block list. And that, after the final assessment is prepared, there are only 10 strings on those. But one of them wasn't one of the original 1,000.

JEFF SCHWARTZ: I see. Right. So asking it -- refactoring, could a different block list appear in the final report? I, actually, don't know the answer to that. I would like to say that I doubt it, but I don't know the answer because we don't have a final report yet.

But I understand the concern. And that would have to be addressed in the report -- right? -- if somehow a new block list materialized, what do we do if they, you know, weren't on the old block list? So I understand your concern.

JOHN CRAIN: But, if I understand directly, Jeff, your mandate is to look for the mitigation strategies and not necessarily produce a new list, right?

JEFF SCHWARTZ: Yeah. I mean, that's right. I would -- you know, kind of speaking philosophically for a second because I don't know what the final report is going to look like, I would not like it if the final report contained SLD block lists. I can't rule that out, because we don't know. But I would not like that. That doesn't seem like a -- you know, like the right long-term approach.

CHRIS WRIGHT: Chris, ARI Registry Services. I have two comments I'd like to finish quick, so John can finish on time.

JOHN CRAIN: We can go over time a little bit for you.

CHRIS WRIGHT: The first one is a comment for everybody, but especially ICANN. I would implore you to ensure that this doesn't go the way of the trademark clearinghouse where there was an issue that needed to be resolved or something that needed to be implemented and ICANN went away in a magical black box and came back with a process. And then we -- the registries and the registrars and so forth that had to implement that process had to fight to get that

process changed and something that was a bit more palatable for all of us. I'm hoping, as we go through this process right now that at the end comes out with something that all the registries will have to implement one way or the other, that there is enough consultation with the affected stakeholders -- and not just the registries, because, of course, there will be other stakeholders that will be affected by this -- registrants, ISPs, et cetera, et cetera. But make sure there's enough consultation throughout this process so that we're not just disappearing into the black box and coming with something at the end and say, bang, here's a magical solution you're now all stuck with. I know we said there would be a comment period, but I'm not sure a comment period is enough. But you need to balance that with the overall time frame. Obviously, we all want a solution to this as quickly as possible. And, obviously, understanding that extended consultation would extend that out. But just implore you to strike that right balance there. So that's just a comment for consideration.

The second question is one that you probably aren't going to like me for asking. But I want to know if this is your name collisions going to be extended to IDN ccTLDs? So we keep talking about this in the con -- pardon me -- in relation to new gTLDs. However, there's no technical difference between any new entry going into the zone file into the root zone, whether it be a new gTLD or an IDN ccTLDs and that IDN ccTLD fast track. So shouldn't these

name collision rules such as blocking and so forth, apply to those IDN ccTLDs coming through the fast track? And the example of this, of course, is the Iranian IDN ccTLD that was delegated just in September, recently, September/October, when ICANN did know about the name collision and so forth and somehow found fit to go through with that name delegation without considering any of these issues and risks whatsoever. So, yeah, I'd like to know what we're going to do.

JOHN CRAIN: Is there anybody here who wants to answer that? It sounds like a policy question.

JEFF SCHWARTZ: I want to echo here your first comment. You know, we take the -- the black box or the anti-black box issue very seriously.

I hope you've seen so far that we're actually trying very hard to avoid that pitfall. And, if you have ideas or if at any time you're really scratching your head at something, please reach out. I really do want everybody to feel we're doing this in the open and that there aren't surprises at the end.

CHRIS WRIGHT: Great.

JOHN CRAIN: So we'll try to be as open and transparent and inclusive as always, Chris. I promise you that.

On the overt question you ask, there's nobody on this panel that has the answer to that. But it's noted, and we'll make sure it gets through to the people who can probably answer that. Thank you.

And with that, ladies and gentlemen, thank you very much. We're only three minutes over time. So, hopefully, there's coffee somewhere. Thank you, everybody. Enjoy the rest of your day.

[Applause]

[END OF TRANSCRIPT]