DURBAN – Security, Stability & Resiliency Update (SSR)
Wednesday, July 17, 2013 – 09:00 to 10:30
ICANN – Durban, South Africa

BRAD WHITE: Folks, we're going to get going here in one minute. We need just one minute, so just take a seat and we'll get going here.

We want to welcome you to the Security, Stability & Resiliency Update. We think this is going to be a pretty interesting session. I did not take part in this session as a moderator in Beijing. I did in Toronto. I'm not a security aficionado, so I found it very interesting.

What was unique about the session in Toronto is it was one of the few ICANN sessions where people started piling into the room and staying as opposed to a normal ICANN session where people start leaving the room. If that's any sort of indication, this should be a very good session.

We're going to basically break this into two parts, this session today. First, we're going to be talking about the CAB Forum. There will be a presentation.

The second part is a little more open, a little more broad-ranging fielding of your questions. We're going to welcome your questions throughout.

In the back over here, we're going to have an Xplane drawing/diagram, which will be posted in a sort of visual summary. We're trying to learn a new language here at ICANN, and we're trying to do this in as many sessions as possible.

So with that, let me introduce our panelists. We're going to go from left to right. Dan Timpson – Dan, wave your hand. Jump up and down. Scream. He's Director of Quality Assurance at Digicert. He's also filling in for a colleague who's a member of the CA Browser Forum. Next to him, Lyman Chapin, Interisle Consulting Group. He's also a former ICANN board member, and a current member of SSAC. Jeff Moss at ICANN, who is our Chief Security Officer. Francisco Arias, next to Jeff, our ICANN Registry Technical Liason. And Patrick Falstrom, who's the Chair of the Security & Stability Advisory Committee.

With that, we also have an announcement today, a little bit of news that we're making here. And I'm going to turn that over to Dan, who will announce what that is.

DAN TIMPSON: Thanks, Brad. So the announcement that we're making on behalf on the CA Browser Forum is a letter regarding the proposed delegation of .corp as a gTLD. This letter will be posted on the ICANN website shortly. It's not up there now, but you can watch for it in the near future. I've got a presentation that we can go ahead and queue up.

Again, glad to be part of the Security, Stability, and Resiliency Update here. I'm going to cover a few things with you, and we can go to the next slide. It's not working – there we go. All right. So for those who may not have some background on this, I'll cover briefly what is a Certificate Authority. I will cover the current situation with gTLDs and internal names, some of the actions that have been taken so far, and recommendations.

Okay. First and foremost, a Certificate Authority issues publicly-trusted certificates on behalf of users and organizations. We follow strict guidelines or baseline requirements on how these certificates are issued. At a high level, there are ceremonies for creating the roots for the Certificate Authority. Those ceremonies are usually recorded, and the keys are stored on HSMs.

I mentioned that we go under strict third-party audit requirements, and we take a lot of measures to keep those private keys protected. Certificate Authorities apply for the various browser-trusted root programs so that we can get our trusted roots into browsers like Internet Explorer, Mozilla, etc. In order to stay and remain in good favor with those trusted root programs, we need to comply with the guidelines and browser requirements for those.

So when it comes to assurance levels, there are different kinds of certificates that CAs issue. Users and organizations will go through different steps and vetting processes for these different kinds of certificates.

The least amount of vetting is a Domain Validation Certificate, which is simply some checks that we do on the WHOIS records before issuing that certificate. Stronger verification is done on organizational or OV certificates, so there's additional documentation that must be provided as the vetting process.

And then the highest level of assurance is an EV certificate, or Extended Validation certificate. These certificates stand out to the end user. In the browser, you can see the green bar that's indicated next to the scheme and URL. These certificates go through the most rigorous vetting process.

Again, just to emphasize, we need to follow the rules with the root programs with the browsers, and once our roots are in the browser, the validation checks on the user end all work.

What does it look like if your trust root is not in the browser? You have these error messages over the right that I've indicated here. You also get these error messages if a certificate has been revoked. There's different revocation mechanisms that are used. The most two common are Certification Revocation Lists (or CRLs), and online certificate status protocol, or OCSP. So if a certificate is revoked, then you'll get these error messages.

A little more information on revocation. So all browsers perform some level of certificate revocation checks, whether it be CRL or OCSP, and all CAs must provide revocation information via OCSP. Those are in the baseline requirements for CAs.

Cache time can be up to seven days for OCSP requests. Another additional security factor or recommendation when doing revocation is to turn on OCSP stapling, which bundles the OCSP status check with the certificate. Most of the server platforms these days do support OCSP stapling. It may not be configured by default, but it's there and can be used. IIS, Apache – they all support OCSP stapling. Ingenix – I'm not sure.

All right, switching gears a little bit. Some background on internal names. So there's been prevalent use of internal names. By some estimates, we've got around 11,000 certificates that have been issued for internal names. These would be things like .corp, .home, .network – many of the top internal names that you might be familiar with. It's important to note that these internal names were actually encouraged as a best practice up through 2011. So for UNC deployments or exchange deployments, this was a pretty common practice.

Why is this a problem? So we of course have many servers that are already configured this way, and there's going to be an experience that's different outside your internal network once one of those internal names is now resolving publicly. It also opens up end users for man-in-the-middle attacks and compromises.

Some of the action that has been taken so far – there has been a lot of good work with the CA Browser Forum in SSAC and industry groups. The original baseline requirements mandated, historically speaking, that all internal certs expire or revoke by 2015. Now the important thing here is that we've gotten a lot of feedback from operators and businesses that there's going to be some significant roadblocks to changing the networks that they're responsible for. There are policy issues – internal policies. There are the costs of changing from internal names to FQDNs, and of course training.

So the CAB Forum has been approached by ICANN, and in February 2013, cast a new ballot that essentially accelerates the deprecation from five years down to 120 days after the relevant gTLD contract is published. Now the 120 days is to hopefully accommodate much of the work that would be required to change those more common names and get networks reconfigured.

Also, other action that was taken is the Mozilla Foundation adopted the revised browser baseline requirements, so that as of July 31st, all CAs most comply in order to remain in the trust store.

UNIDENTIFIED MALE: Now that's important because these are CAs who aren't part of the CAB Forum.

DAN TIMPSON: Right. Good clarification. Okay, additional actions that have been taken. So there has been a council that was created, the Certificate Authority Security Council. It's essentially an advocacy group, and one of the purposes was to improve the education, marketing, and research along these topics. So there's information out there along for OCSP stapling, reconfiguring servers with FQDNs, and just engage the community at large with these issues.

Many of the CAs are reaching out to customers, as Digicert, to make customers and operators aware of what's going to happen. So we can communicate a lot with the customers, tell them what's going on, but it doesn't solve the heavy lifting that they need to do. It doesn't necessarily reduce their costs. We do alert them to the situation. We at Digicert have provided an internal name tool, so for exchange environments, a sys admin can go in and run this tool and figure out what they need to do to reconfigure their environments with FQDNs.

UNIDENTIFIED MALE: [inaudible]

DAN TIMPSON: It is, yes. We actually have a lot of requests for it, and it's on our website. In fact, this link in the deck will go to it, for your interest.

So in terms of recommendations for ICANN, what we're saying is that we're recommending to not approve the names that are most commonly used in internal certs until 2015 to buy a little bit more time. The most popular names can be referenced in the PayPal letter. As I mentioned earlier with the announcement about the .corp gTLD, this Digicert letter is also going to be posted live on ICANN's site.

An alternate recommendation would be to approve the application of some of these domains, but to delay the delegation until 2015. We feel like that the remaining 90%, the less common names, we can move forward with those, and there's not as much of a security impact relating to those less commonly used names.

BRAD WHITE: Let me ask you a question. Just so people in the audience understand, that 2015 date is based on maximum expiration time of three years based on when the cert's issued?–

DAN TIMPSON: That's right. You're exactly right.

BRAD: So the idea is, even if there's no OCSP, the certs will expire and all browsers treat expired certs pretty much the same.

DAN TIMPSON: Right. Those errors that we had in the deck earlier, users would see that at that time when they expire, it would just stop working, stop validating through the trust chain. Okay. So with that, I believe I'm going to turn it over to you, Lyman.

LYMAN CHAPIN: Okay.

UNIDENTIFIED MALE: Wait, wait, I've got a question.

DAN TIMPSON: Sure. Yeah, go ahead.

UNIDENITIED MALE: In totality, those are the issues. There's no other secret issues, you think, lurking behind the scenes?

DAN TIMPSON: I think those are the main issues. There could be others that just aren't as present, but those are the top issues as we've looked at this.

UNIDENTIFIED MALE: And does a browsing more of a move toward mobile browsing? Does that impact any of this – the way that the mobile browsers validate or behave, or is this just a clear transition from desktop to mobile?

DAN TIMPSON: I think that the mobile doesn't add anything that would not be covered in the desktop experience in terms of revocation checking and validating certificates up the chain, so I view it largely the same.

UNIDENTIFIED MALE: To clarify, I think that what you've really asked was this will resolve the [name collision] issues related to the certificates, right? Because we will go into other name space collision issues as well.

DAN TIMPSON: True. This is not just a certificate problem, I guess to highlight that. If there's internal routing, host files – things like that – internal names can still cause havoc and problems in the network.

BRAD WHITE: Dan, we also have an online question. Will the CAB Forum publish per TLD, internal name statistics for aggregating all CAs?

DAN TIMPSON: Unknown at this point. I would need to check with the contacts of the CAB Forum and ask them that question.

BRAD WHITE: Sir?

UNIDENTIFIED MALE: There's an applicant fee to use the microphone?

UNIDENTIFIED MALE: We just delayed the deployment.

UNIDENTIFIED MALE: Dan, thanks for the presentation. I think that's great. I'm really impressed actually with the Mozilla trust store issue and saying that if they implement the Ballot 96 recommendation, that's the only way they're going to be added. I think that's great. My question is sort of –

UNIDENITIFIED MALE: We can hardly hear what you're saying, unfortunately.

UNIDENTIFIED MALE: Can you hear me better now? Is that better? Okay, so I was saying the trust store thing with Mozilla is a great progress. I'm really pleased with that. That's the first I've heard of that. I was going to ask about the 11K estimate that you said, Dan. I think that you really – unless you look at the aggregate corpus of internal certs across everyone that makes it in the trust store, you don't know. So do you have that information or is that just a guess?

DAN TIMPSON: That's the information we have at Digicert. However, looking at the SSAC report – I looked at those numbers last night – and that was more in the realm of around 30,000. I defer to Patrick on that, but it's certainly in the tens of thousands and could be much, much more. But just what we've been able to gather far, that's what we've seen.

UNIDENTIFIED MALE: So the SSAC report says 37,000, but those are misconfigurations? Those are Internet facing internal certs? So all those are misconfigured and observable as a result of that. So I think it's important to not underestimate this. Netcraft or the entire corpus of someone going into the Mozilla store and maybe ask Mozilla to say you must publish the yearly statistics on in internal names or something, so you can actually have an accurate number. I think you're probably off by a couple orders of magnitude there.

DAN TIMPSON: Okay. Good feedback. Thank you.

MARILYN CADE: Good morning.

DAN TIMPSON: Good morning.

MARILYN CADE: My name is Marilyn Cade, and I want to express my appreciation for the opportunity to ask a couple of questions. It's a very general question in one way. Are you able to distinguish the size of the corporation that's involved? Because obviously reaching a number of large corporations is very different from reaching a much, much larger number of mid-size or small corporations to try to work with them.

The barrier to cost that we experienced when we did go through the extensive process of dealing with the small companies and suppliers during dealing with the Y2K I think taught us a lot of lessons about how distributed and long-term the education and awareness process may be.

DAN TIMPSON: Yes, certainly, I agree with your question. I think, with regards to "are we evaluating the size," we would need to analyze the data and figure out what the average corporation or company looks like that's going to be grappling with this issue. But I completely agree with your assessment on the difficulty and costs with trying to make these changes, particularly in the SMB space. One of the things we're making recommendations to ICANN to be considerate of is the tax and possible hardship on those segments.

UNIDENTIFIED FEMALE: Hi, it's [inaudible] JPNIC. I'm really glad to see the actions taken after SAC057. I've been keeping a close eye on what's going on. I have one question. According to SAC057, I think it mentioned it over 60-something or 100-something internal certificates that clashes with the applied gTLDs.

I've understood the CA Browser Forum has taken action to revoke after 120 days, which sort of makes sense if we are able to reach those people who have those certificates that receive the clashed names with applied gTLDs. I'm wondering if somebody has taken any actions, or have any plan to reach to those people who receive those soft internal certificates with the same name as the applied gTLDs.

DAN TIMPSON: Sure.

JEFF MOSS: I'd think customer management.

DAN TIMPSON: Yes, that's certainly, like Jeff was just alluding to, it is customer management. So I can speak on behalf of Digicert in that regard in that we are looking at our database of customers. So those who have requested certificates with those gTLD internal names, we will be contacting them, letting them know what's going to happen in terms of revocation or possible changes to their network.

UNIDENTIFED FEMALE: Thanks. That's great to know. Thanks.

BRAD WHITE: Let me just note that after this gentleman asks his question, we're going to move onto Lyman's presentation. We will have plenty of time for additional questions in the second half. Sir?

PAUL STAHURA: I got a lot of questions, but I'll just ask one. That's Paul with Donuts. Is the solution you provide, which is delayed until years from now to delegate .corp – is that the only solution you guys came up with, or did you come up with other ones?

DAN TIMPSON: Well, so the recommendation to delegate .corp is really our recommendation to say, "Reserve that as private," because we feel like that there is enough use of that and a strong enough case to suggest that is an internal name like .local or .localhost or one of the others. Other than that, yes, the delay does buy businesses more time to reconfigure, retool, and redeploy in some case.

PAUL STAHURA: There's also a cost, which is I believe to competition by delaying that important TLD for years. There's a cost to not bringing competition sooner by delaying that. But another solution, possibly, is we delegate it, and then if somebody presents a certificate which we could check to the registry for a certain subdomain, then we could put that subdomain in the zone for .corp. So therefore we could go ride with corp sooner and get the benefits of this competition and the benefits .corp would bring to the Internet sooner. Why can't we do that solution?

DAN TIMPSON: Well, I'm not prepared to debate all of the technical reasons; however, I think that your line of reasoning is good in that we need to look at other possible solutions that could help with that. I think just the massive scale and nature of this problem – the time is certainly going to help either way.

PAUL STAHURA: Thank you.

DAN TIMPSON: Yup.

BRAD WHITE: Lyman? Sir, can we just hear this presentation, then we'll open it up and we can take more questions?

JORDAN BUCHANAN: I have a question about the certificate problem [inaudible].

BRAD: I know. What we'd like to do is just get to this presentation, then we'll open it up and you can ask on any subject.

JODAN BUCHANAN: So the problem is you've scheduled this at the same time as the NTAG meeting.

BRAD AND OTHERS: Go ahead.

JORDAN BUCHANAN: Sorry for showing up late. So I just saw with this remotely, so it's possibly I've missed conversation. If I have, just tell me to go away. In the recommendations I saw for the certificate problem, it said we should avoid delegating strings that have commonly-issued certificates until 2015, which it seems like a thing we should definitely think about. But then the specific references I saw to which strings might be affected, where the PayPal letter and .corp – and .corp seems like that's probable that there are in fact certificates issued for that. But the

PayPal letter was mostly about estimating traffic at the root, as opposed to certificates that are issued. So I guess I'm confused to how that is relevant to which TLDs might have certificates, like internal corp certificates issued.

DAN TIMPSON: I think it's relevant in the sense that it's still an internal names issue. That's really why it's relevant.

JORDAN BUCHANAN: Maybe because it's like an indicator of how much internal certs or internal names are out there looking at the leaks that hit the root?

DAN TIMPSON: Yeah, that's true. The other thing is that it's interesting because it comes from business like PayPal. It's not coming from the tech side, from the CA side. It's just additional information to consider in the general context of things.

JORDAN BUCHANAN: I guess I was going to ask you, wouldn't like the SSL Observatory or data like that be more relevant than the NXDomain traffic?

UNIDENTIFIED MALE: I think we should see the next presentation, then [inaudible].

UNIDENTIFIED MALE: There will be more information in the study results.

BRAD WHITE: Also, folks, Marilyn just reminded me of this. If you would be so kind, please, to identify yourself when you go up the mic.

JORDAN BUCHANAN: Sorry. Retrospectively, that's Jordan Buchanan with Google.

BRAD WHITE: Great. Thank you. That's as much for the record and for remote participants as anything else. Lyman, go ahead sir.

LYMAN CHAPIN: Thank you, Brad. This presentation provides an update on a study that my company, Interisle Consulting Group, was commissioned to perform on the incidence and potential consequences of name collision and the DNS. It includes some of the topics that we've just been talking about, so we'll end up in roughly the same discussion space, even though the titles don't like necessarily completely similar. I hit two buttons at once.

I want to start out by describing what we mean when we use the term "name collision." It's not necessarily familiar to everybody, and I apologize for those of you who understand the way the DNS works at a fairly detailed technical level, because the next few slides will bore you, and you will come up with at least 50 reasons why they are wrong with respect to details. But if you'll bear with me.

In the world in which we live right now, where we have not delegated any new gTLDs for a long time, you can imagine someone using a computer – either an individual or application – uses a local name to access a local resource. Typically, in local environments, you would use something like printer.myname. You wouldn't necessarily use a fully-qualified DNS name. Your local network resolves printer.myname and knows where to find your printer.

That printer.myname string looks like a DNS name. It has the dot and so forth, but it's not. It belongs to a local namespace that is only meaningful within the context of your local network. That namespace is not a global namespace the way the public DNS is. And if you ask the DNS about printer.myname, it will tell you that name does not exist because myname is not a registered TLD, and so it comes back with what we call an NXDomain or a Notauth response.

If you look at what situation will be after we delegate a bunch of new strings as TLDs – for example, should ICANN myname as a new gTLD, and I deliberately obviously chose one that is not on the list of current proposed TLDs, and then someone comes along the name "printer" at the second level in the new TLD. He goes through his registrar and says, "I'd like to purchase this name." So now, printer.myname is a global DNS name. Now if you ask the DNS about printer.myname, instead of going to something on your local network, you'll get a pointer to, in this case, [inaudible].

That is, essentially, at the highest level – the 30,000 or maybe even 50,000 ft. level – is what we mean when we talk about name collision. You can think about it also as trying to interpret a name in one semantic domain – the context of one namespace – when it properly belongs somewhere else. So it's really a namespace collision, as opposed to a name collision issue.

So we were asked to conduct a study to find out, first of all, how likely is it that we're going to see this in the real world. Up until now, this has been a theoretical possibility. Everyone can imagine how it might happen. Is it really going to happen?

Second, what effect might that have if we do see name collision after we delegate new gTLDs? What effects might that have on security and stability? And then, of course, given that we might find such effects, what options do we have to deal with them, either before or after the fact, where the fact in this case is the delegation of the new string?

So we looked at the best data sets that are available – the historical data sets that are available – for queries to the root servers. This is questions that are being presented to the root of the DNS, basically saying, "Can you tell me information about the following name?" We had two good, large samples. There's an exercise that was started by an organization in 2003 that has been carried on every year since then called, "A Day in the Life of the Internet." It's an exercise in which an organization called DNS-OARC captures or receives captures of packets that have been sent to each of the root servers over a continuous 48-hour period. Actually, it's a minimum 48-hour period. Almost all of the root servers – the 13 individual root servers – participate in this exercise, and it's a very good, uniform place to go looking for things that might be happening in the query stream to the root.

So we took those two data sets, which together compromise about somewhere between eight and ten terabytes of information, and we looked for proposed TLD names in those data sets. We also then, after we did that, did some work to investigate some of the potential consequences. The two areas we focused on were what happens when the resolution of a name is ambiguous when it's not possible without additional contextual information to determine how to resolve a name. And the other issue is the one that has to do with the internal name certificates – X.509 public key certificates that Dan has just been talking about.

The last part of our study was to investigate some of the options that might be available not just to ICANN, but to the community at large, to mitigate some of the effects of these collisions in those cases in which the consequences are severe.

So if we look at the query stream to the root – this is for 2013 – we'll see that there's actually some good news here. About 55% of the queries that the root servers receive are questions about actually existing TLDs. This is what you would expect. In a perfect world, of course, that would be 100% because nobody should be asking the root about things that don't exist.

In this slide, what we call a proposed TLD is one that has been proposed as a TLD in the current new gTLD program round. So one of the 1930 original names that we in the pool. And 3% of the query streams – keep in mind, this is before any of these have been delegated – 3% of the query stream consists of questions about those strings that are on the proposed list. 19% are what we call potential TLDs. Those are strings that don't exist as TLDs, they haven't been proposed, but they're syntactically valid strings. In other words, they could be a TLD in the future if someone proposed them. Their syntax is correct. And then 23%

of it is just garbage. It's invalid strings that could never be a TLD because they don't obey the syntax rules for top-level domain labels.

This is a list of the most queried TLD strings. This covers existing proposed and potential TLDs, so everything expect the 23% that are invalid. It's important to point out, first of all, these numbers are in thousands, so in the 2013 data sets, what we see for .com is roughly eight and a half billion queries for .com. Not surprisingly, .net and .org are on that list as well. If you leave out .local, which is a special case that got lumped in with queries for the root itself, the next item on the list is the string home with just over a billion queries. You go down the list. If I continued this list out to the top 100 most queried TLDs, there would be 13 of the proposed TLD strings on the list.

Of those the proposed TLDs – the ones that have actually been proposed in this round – this chart shows the rank and counts for the most queried proposed TLDs. There's some interesting things to note about this chart, not just the fact there's a fairly accurate power law fit to the distribution. .home, .corp, .ice, .global – as you go down, it roughly follows a power law. Although for .home and .corp we have a good pretty good sense of what's calling those strings to appear in queries when they're not actually currently delegated as gTLDs, it was kind of a shock to see. ice number three on the list.

To give a sense of how to think about this, just occurrence – the number of times you see a string in a query – doesn't tell you a lot about whether not that's a good thing, a bad thing, a neutral thing. It just tells you how often a string appears. The additional information you need to make any kind of risk assessment is how serious a consequence might ensue if you saw that string and it collided with a delegated new gTLD string.

As an example, an event that occurs very frequently but has no negative side effects is one thing, and an event that occurs very infrequently but has a really serious side effect, like a meteor strike or something like that, it's always a product of those two factors that leads you to an assessment of risk. So just the fact that a string occurs a lot, it looks scary – not necessarily so.

So if you go down this list, you ask yourself, "Okay, .home we pretty much have a handle on why that's occurring. There's lots of routers and DSL modems and so forth that are configured to use .home in a local environment." .corp we've already talked endlessly about the way in which active domain name configurations are frequently set up. Active directory name configurations are frequently set up using .corp as the top-level domain.

.ice turns out to be the Electric Utility Co-op in Costa Rica, which for some reason, is blasting .ice requests out to the root in the third position on this string. You can imagine that the occurrence and consequence product for that would be very different than it might be for some of the other things on this list.

This list obviously goes on for a long time. In 2013, there were only, I'd say, 14 of the proposed TLDs never appeared. So this is a distribution with a very large head, a very long tail. There are lots of things out in the negligible occurrence tail, but there are only 14 of those strings that are currently proposed that never occurred. If we look at larger data sets – we have had some informal discussions with individual root server operators – if we looked a larger data set, it's almost certain that every single string that's been applied for would be found somewhere in the query stream to the root.

JEFF MOSS: [inaudible] HSBC.

LYMAN CHAPIN: Oh, yeah, Jeff just pointed out that, as you go down this list, you also see some entries on here, like at number 14 for 2013 rank, HSBC, which is highly likely to be restricted to that particular company – that bank. However, if you look at rank number ten, you see Cisco and you would think, by the same logic, "Well that must be the Cisco corporation," but in fact, it's most likely an artifact of the way in which people tend to name their routers – router1.cisco.corporationame.com, or something like that.

So the kind of triage that you could imagine doing by looking at the strings and making sort of informed guesses about what they might mean in terms of origin, you have to be a little bit careful. You can do that in some cases. If you saw IBM on this list, which is down below 15 but is on the list, you might assume it's the IBM corporation, but you would want to follow that up with some investigation to be sure.

So if you look at the potential consequences – and again, before I go through what looks like a very scary list of bad things that might happen, I want to point out once again that we need to look at both the potential consequences and the likelihood that they might occur before we make any judgment about what the delegation risk might be.

But the most obvious potential consequence right off the bat is that it's likely to change the way in which local namespaces resolve. So if there's a collision, you might find yourself accustomed to a particular behavior on your local network that might change if the label that you were using as the top-level domain in your local network were suddenly to start resolving in the public DNS after it was delegated as a new gTLD label.

Search list processing is likely to change as well. Search list is a list that's maintained by your operating system or by a piece of application software which adds suffixes in order to the string that you might enter at a user interface to try to create a fully-qualified domain name that will resolve. So it might try the name you entered first, and then it might try it with a .example.com extension, then it might try it with corporationname.com. So it's going to try lots of different suffixes until it comes up with a fully-qualified domain name that actually resolves. Obviously, because today a lot of these strings are not delegated and tomorrow they will be, the way in which those search lists cause resolution behavior to happen in your local environment could easily change.

One of the consequences of that is the possibility that various kinds of application streams and packets could get misdirected. When we look at the databases of packet streams to the root, we see not only requests for what you might call standard requests for A or address records, we also see requests for .mx, which is mail exchange records, and .srv, service records. Mail exchange records are typically used by mail processing systems, obviously. SRV records are most often used by SIP, the protocol that's used for Voice-Over IP. The fact that we see queries for those suggests that there are systems that are configured to look for that information from a local root that are escaping onto the public web.

So when you start to resolve those globally, it's possible that mail voice-over IP calls and so forth could get misdirected. Again, I'll just add that this is not something that is going to be a wholesale. We're not going to turn on a new TLD and suddenly see mail across the entire Internet suddenly going to the wrong place. These are things that could happen, and it's important to point out that in order to determine whether or not they would actually happen in practice would require more investigation than we conducted in this study.

The public key certificate issue is the one that we just covered when Dan gave his presentation. It's also covered in the SSAC Report 57. The final one on this list – these are the ones that we've spent some time on in the study; there are probably other consequences that we don't know about – but the way in which web browser cookie data are stored, cookie data are always coupled to the fully-qualified domain name, so it would be possible under certain circumstances that cookies stored on your machine to be accepted in a different environment where the name resolved differently in such a way as to expose your cookie data, which would enable someone else to essentially take over you, and find out your identity information to masquerade as you in other situations. Again, this is not something that's going to happen to all of us as soon as somebody turns up a new TLD.

The options that we came up with for potentially resolving some of these issues – these may apply only to a very small number of strings, or they might apply to the entire list of strings as a whole, depending on how we decide to deal with some of these issues.

The obvious one is to permanently reserve a string that you've decided is both likely to occur – name collision is likely to occur – and the consequences are very serious. So if both of those variables have very high values. The product, which is the measure of risk that we're using, would be very high as well. There have been some suggestions over the past few years, particularly within the IATF, to in fact permanently reserve some of those strings to prevent the name collision issues that we've been talking about from happening.

That's a pretty radical step to take. So again, these are options, not necessarily things that anyone is going to do. In particular, we don't, as a result of our study, explicitly recommend that ICANN or anyone else follow one of these options. These are presented genuinely as choices that can be debated within the community.

Another obvious option is to study the impact, either of an individual string or of name collision in general, more than we were able to do in this study – get better data sets, do more investigation into what the consequences of name collision might be in specific cases by going out into the world and asking people who maybe have experience with some of these things. There's a lot of avenues that you can imagine exploring a lot more thoroughly than we did in this study. Obviously, that would delay delegation and you would have to have some kind of a termination condition, because whenever you suggest that further study be done before you can take an action, you have to say, "Well how much study is going to be enough, and when will I know that I've studied the problem enough to be confident that I can now make a decision?"

A third option, which we call "wait until everyone's left the room" or "wait until everyone's gone" is very similar to what the CAB Forum is suggesting with respect to the strings that are most commonly used in internal name certificates, which is delay delegation of the string until the colliding use has stopped. In that case, it would be until all of the internal name certificates that have been issued with that string in the subject name or subject alternative name field have either expired or been revoked. That same approach could be adopted for other strings if the uses are such that they could readily be changed.

So, for instance, if the Electric Utility in Costa Rica could be convinced to close whatever hole is causing the .ice queries to escape out into the public Internet as a query to the public DNS root, that would presumably solve the problem with name collision for that string.

So there are many ways in which you can – wait until everyone's gone doesn't necessarily mean wait for ten years until you're certain that every possible use of the string has died out. In cases where you can identify, very specifically, the reasons why a string is occurring, you may have options that don't require that lengthy delay.

And then there's a fourth option which has been discussed, and in particular it's been proposed in the letter from VeriSign to ICANN execs, which is to do what VeriSign calls an ephemeral delegation, and we called it a trial run, and that's to delegate the name in such a way that it isn't being operated by the eventual registry operator – the applicant. It's being operated either by a third-party – some trusted party –

carefully deploys it, established the kind of monitoring that would be necessary to determine if something bad happens, and if something bad happens, withdraws it very quickly. And obviously, you would advertise those names, for instance, with very short time to live, or TTL fields, so that if some negative effect were observed, you could withdraw it very quickly and the consequences would hopefully be limited. There's some obvious pros and cons to that, and I won't go into a lot of the details because the devil is definitely in the details with that one. But that is certainly an option.

UNIDENTIFIED MALE: So when you just say no, your first option you could also apply that not just to the whole TLD string, but to a subdomain. So if it's just powermeter.ice that's causing the query, you could potentially register just the powermeter, reserve that, and those queries [wouldn't] be impacted by collision.

LYMAN CHAPIN: That's right. Yeah. And of course, in that case you'd have to be sure that whoever was operating with .ice as a TLD would never delegate any further names. Right. I'm happy to take questions.

BRAD WHITE: Before we take your questions, sir, Kevin Murphy from Domain Incite in London has had a question here posing for some time, Lyman. "Did the study look at the sources of the requests for the proposed TLDs, i.e., is there a way to figure out how many of the users or networks might be affected if these TLDs are delegated?"

LYMAN CHAPIN: We did have an opportunity to collect data on the source. We collected those data using IP address prefix, so we do have a sense of where the queries are coming from. What we did in the study report is we looked not so much specifically at where a request has come from, but for each of the TLDs that we looked at, how many different sources do we see queries from, because the interesting thing is really how widely distributed is the query stream for a particular string.

If it's coming from one very particular place, as it is in the case of the .ice string that I've been talking about, then you can be pretty confident that wherever this string is escaping from is limited to at least, from a topographical standpoint/IP address prefix, it's limited to a relatively small area. That suggests that it might be easier to figure out a way to stop that conflicting use than it would be for a string that appeared from many different IP address prefixes. So yes, those data are in the report. I don't have them in the presentation today, but they are on the report.

BRAD WHITE: Great, thanks, Lyman. Just by way of a time check, we've got about 35 minutes left in this session. We'll start taking questions now. Again, if you would just give us your name and who you're representing, if anyone, that would be greatly appreciated. Sir?

ALEXANDER MAYRHOFER: Thank you. My name is Alexander Mayrhofer from nic.at. I have two questions. The first one is a simple one. The numbers that you gave –

were those queries per day or per months – the billion queries that you had. That's my first question.

LYMAN CHAPIN: No, those were counts of the total number that we saw in the day in the life data set. So it's aggregate number. It's not a per day number.

ALEXANDER MAYRHOFER: It's about two days, or something like that.

LYMAN CHAPIN: Right. So it would have to be compared to other numbers. It's not really all that meaningful as an absolute number.

ALEXANDER MAYRHOFER: Sure. The second question is you mentioned cookie leak before. I was wondering, could you give us an insight about work regarding the public suffix list? I haven't been following that one, but I understand that this is a major problems with cookie leaks as well. Are you aware of any work in that direction?

LYMAN CHAPIN: I'm sorry, I didn't hear the question well enough to answer.

ALEXANDER MAYRHOFER: There was work regarding the public suffix list that defines what is a valid registry in all the browsers, and I understand that the Mozilla people have been working on redefining what the public suffix list does in the light of thousands of names, getting into the root. The problem is essentially [inaudible] cookies, as far as I understand.

LYMAN CHAPIN: Right. There is a separate issue that we didn't cover in the study but is definitely of concern, which is how are we going to ensure that these new domains – these new suffixes – are going to be recognized as valid? But that was not something that we looked at in this study, but is definitely something that ICANN and SSAC and a bunch of other folks are worried bout.

ALEXANDER MAYRHOFER: Okay.

FRANCISCO ARIAS: This is Francisco. So regarding the public suffix list, Mozilla agreed to use [inaudible] mailing list announcement with anyone interested to know TLDs that have signed a contract with ICANN and will still agree to use that to populate the public suffix list.

PAUL STAHURA: Paul Stahura with Donuts again. I have a couple questions. I'd like to do them in order, so could you go back a couple slides to the pie chart? To follow up on the prior gentleman's question, you said you got to compare those numbers to other numbers. I agree with that, and I'm curious about whether – in this pie chart, I think you said you looked at how many NXD records, for example, produced the 3%, and how many lookups in .com produced the 55%. But I think you're comparing two different things, because one is NXD, and one isn't.

The first thing I'm curious about with this pie chart is, did you take into consideration the TTL for the NXD being different for NXD domains coming back from the root, because for example, those networks home boxes might query, and query, and query with a short TTL, which would increase that 3% compared to the 55%. So I'm wondering if you factored the TTL into this, because I'm pretty sure the TTL is different.

LYMAN CHAPIN: I would agree with you that it almost certainly is. We did not look at TTLs. This is total number of queries. But I also want to point out that what we're looking at is only the query stream to the root. This doesn't include the query stream to intermediate resolvers – third party resolver operators. So a large number of the queries that actually make it to the root are actually coming from those recursive resolvers. So there's a whole section the report that is nothing but, "here's why the data might not be accurate."

PAUL STAHURA: Looking forward to your report. I think this pie chart is misleading, in my opinion because if we look at the TTL and the other factors like you just said, because they're coming from recursive resolvers – for example, if those home networks that just do one query but across a large amount of recursive resolvers, that's going to be different from Yahool – people go into Yahoo! or Google or whatever at one ISP. That one ISP would just make one query to the root, where – okay.

The other thing I think you should do, since you have the data, is look at what percent of that 55% is NXD queries for .com names, because .com names get NXD queries currently.

LYMAN CHAPIN: Right.

PAUL STAHURA: They could get a trillion a day, too. So why aren't we talking about this name collision for not-taken .com names?

LYMAN CHAPIN: From the data set that we have, you could certainly do that, yeah.

PAUL STAHURA: So I would compare that information – NXD to .com – to compare apples to apples, to NXD for subdomains and let's say .corp. So that's my first question about this graph. I got two more. Could can you go to the—

BRAD WHITE: Hey Paul, could you possibly make it one more, because we have a huge queue behind you.

PAUL STAHURA: I got a lot to say about this issue, but yeah, I'll make it one more. But I'd like to maybe to talk to you guys after, or individually. I'd like a seat at that table, actually.

BRAD WHITE: We can definitely carry this on.

PAUL STAHURA: Okay, one more then. Go to I think one or two more slides.

LYMAN CHAPIN: I'm sorry, which slide are you looking for?

PAUL STAHURA: I think one more – nope, the next one. Go forward. That one. So as I said, you can't really rank them like this because you're comparing two different things. So this ranking…

LYMAN CHAPIN: We're not comparing two different things. If you want to look at the study methodology, we are definitely not looking at all possible ways at looking at the data, so we're not claiming that this is the last word on the subject, but we definitely feel that it's a valid comparison for the purpose of at least looking at the way these things rank out.

PAUL STAHURA: Okay, I understand your opinion.

UNIDENTIFIED MALE: I think that Roy pointed out, too, that looking at the root queries, you'll never be able to find out the NX queries for .com because only .com can do that.

PAUL STAHURA: That's not true, because you can get the .com's own file and look at the queries that you get for .com, compare it to the zone file, and find out which ones are in there and which ones are not. So you could predict, or you could just ask VeriSign to give you the data. But you can do it without asking VeriSign. So I disagree with you there.

But back to this. Since you have the data, I assume you looked at, let's say, the .belkin, or anyone of these names, did you look at the second level? How many of those queries of the trillion queries for .home, could they all be for one second-level domain in that home, correct?

LYMAN CHAPIN: We did. We actually went down to the ninth level. We've got counts for – but mostly what we have, and someone else has suggested this would be a useful additional look at the data, what we have is the occurrence of the string in each of those positions. So the occurrence of, for instance, .belkin for the TLD, SLD, third level, fourth level, and so forth. An interesting additional study, which you just alluded to, would be to look at when .belkin or any other string occurs in the TLD position, what is occurring in the SLD position.

PAUL STAHURA: Right, so, for example –

BRAD: Paul, can we give – we got a huge queue behind you. These guys – you can engage these guys afterwards.

PAUL STAHURA: Last statement. That .corp one, they could all be Microsoft or a corp. for all I know. Thanks.

TOM HASH: Hi. Tom Hash from Yahoo! A couple questions. The first is about the data set. So the day in the life data was around 48 hours. It's more of a comment than anything else – I think it does a good job covering time of day, but not really day of week. So there could be a whole host of issues that proliferate that we aren't seeing, whether it's on the weekend – possibly they could kick off on the weekend. Just a suggestion you might want to do, instead of ten terabytes, 100 for a whole week or something. Thoughts?

LYMAN CHAPIN: It would definitely be better if we had a data set that spanned a much, much longer period of time. In particular, one that spanned things like the first day of the week, first day of the month, first day of a fiscal quarter. I completely agree. We were limited to the day in the life data primarily because, first of all, those are the most complete data that are readily available, and also it enabled us to do a little bit better apples to apples comparison because of the uniformity in the way in which that exercise is conducted.

An alternative would be to do your own data-gathering exercise with the cooperation of the root server operators and so forth. But it's daunting because these data sets are huge. These files are enormous, and running processing – we're using very large multicore processors, and it's still taking day to go through some of these data sets.

But your suggestion is absolutely correct. A better coverage of time would give a better read on what's actually happening.

TOM HASH: Second one was you mentioned the big head/long tail. I think a lot of the things in the middle would be good to make all of this data available so people could practically go out and individuals could look at their own company and say, "Oh wow, I've got this, that, and the other that are a long tail, but my impact could be huge." So just the frequency of queries isn't necessarily the impact it could have on an individual company.

LYMAN CHAPIN: Yeah, and all of those counts will be in the report. You'll have access to all those data.

TOM HASH: Right. Thank you.

BRAD WHITE: Before we take the next question, we have about 25 minutes left – the queue is quite long, so if you folks could just ask your question and if you got a follow up, we can confine it to two bites of the apple just to accommodate everybody, that would be greatly appreciated. And Lyman, we had a couple of queries online. When will the study be available? When will it be released?

FRANCISCO ARIAS: This is Francisco from ICANN staff. You may remember there was announcement when the study was started and we were aiming to have this study available before Durban. However trying to have a more solid study, we realized more time was needed. For example, to have data from the CAs, we have a letter of intent of corporation with the CAB Forum that was signed at the end of June. We have incorporated with the CAs for all this time. So we need more time to get the data, and that’s the reason why this study has not been published. We’re aiming to have this study within two weeks after.

BRAD: Sir?

UNIDENTIFIED MALE: Thanks. So I just have two questions, but I think one is really short. In your pie chart, you had potential future and invalid, and you’re purely distinguishing based on syntactical correctness? For example, RC6761, I think there’s various reserved TLDs that in theory also could never be delegated, but you’re calling those potential TLDs as opposed to invalid TLDs? Is that correct?

LYMAN CHAPIN: That’s correct. It’s only syntactically invalid, which includes labels that are not valid as DNS labels, but also labels that are not valid as TLD labels according to the AGB.

UNIDENTIFIED MALE: Okay, thank you. So my follow up question is related back to certificates, because this can actually help me understand my previous question, which is that, as I understand the recommendation that I saw from the previous discussion about certs, that discussion said we should maybe think about delaying the delegation of TLDs where there were potentially a large number of internal certificates issued.

And then we talked about the PayPal letter, and then you guys referred me to this discussion, but none of that is actually about the incidence of certificate issuance. It does make sense, I guess, that if no one on the public Internet ever resolves a name – or it's very rare that they resolve a name – that it's very unlikely that you'd be able to do a man-in-the-middle-attack with one of these existing issued certs, so maybe that's an additional factor.

But I guess it seems we need a data set, and it sounds like maybe Francisco is working on it on actual certificate issuance or certificate probabilities, as opposed to NXDomain traffic, which tells us how often people are navigating to this, but not whether there's a certificate or not that's been issued.

LYMAN CHAPIN: Yeah, that was in fact a part of the study and those data will be in that report.

UNIDENTIFIED MALE: Oh, awesome. Great. Thank you.

BRAD WHITE: I might add, before this lady asks her question, by all means keep going on the study if that's where your interest lies, but if you have other questions regarding issues of security or resilience, feel free to ask them. We'll go where you want to go is what I'm trying to say. Ma'am?

AMY MUSHAHWAR: Thank you. I'm Amy Mushahwar with the Association of National Advertisers. We're an association that represents the Fortune 100s, so we represent billions of dollars in advertising revenue and trillions of actual revenue globally.

The problem that I have here with the name collision issue is that it is a publicity issue as much as it's a technical issue. What very much concerns me is when you say .net is one of the top 13. Have you done any further study to determine whether or not wireless area medical networks, wireless medical devices, or other extremely-sensitive, very high risk applications to the DNS may be impacted?

Also, you said SIP communications are one of the series of traffic that you're seeing within your study. Have you done any study to determine whether or not any of the SIP may be communicating with PSAPs, Public Safety Answering Points, like 911? There are some dangerous consequences that we have here that's not just impacting revenue, but could impact lives.

LYMAN CHAPIN: Yeah, that's a very good point. The study gives us a lot of information about how frequently we see these things. It doesn't tell us a lot about what the consequences of name collision might be for any individual string. I would be just as concerned as you are to make sure that before I made a judgment about the product of those two things, which is what we used to assess risk, that I do some serious investigation on what the consequences for strings like the ones you mentioned would be.

AMY MUSHAHWAR: Yeah, and these types of occurrences may only occur episodically, and you only have a six-week longitudinal study. So it's very concerning to our large companies, and we encourage you to use us to help you investigate this dilemma.

We also liased with the Chief Information Officer's Executive Network, which is 1500 CISOs globally, so we offer up that network as well to help with the publicity issue. Thank you.

LYMAN CHAPIN: Thank you very much.

BRAD WHITE: Let me just add, we've only got about 15 minutes left, and clearly we're not going to get to all the questions, but I wanted to tell folks is if you don't get a chance to ask your question, the public forum has a large open mic session tomorrow at the end where any subject is fair game. So if you don't have an opportunity to ask your question here today, show up tomorrow, and that will give you another short? Danny?

DANNY MCPHERSON: Danny McPherson, VeriSign. Again, thanks for this. I wanted a make couple of statements and a note on the ephemeral delegation point, as you mentioned, Lyman, that was a recommendation from VeriSign. But that was a recommendation only subsequent to recommendations that have already been made in 2009 by the RSST Expert Team, which you were actually the chair of. In SAC 45 and SAC 46, there's a list of outstanding recommendations. One is about an early warning system and capability. One's about the capability to instrument the root system to get this data over a window larger than "A Day in the Life of the Internet" because things like TTL expiration does certainly have a huge impact on the queries that come to the root from the data we looked at.

So I think there's a number of recommendations there, and sort of out of that and that set of recommendations, and public forum comments in 2010, the requirement to notify potentially-impacted parties was something that people said is important. In other words, if you're going to break somebody's network, we have an obligation, and it's certainly in the public interest to notify those impacted parties.

So VeriSign in particular, being in one of the roles where we help effectuate new delegations feels like there's a lot of risk associated with this. So what we were wondering is what sort of actions ICANN might be taking to address the liability issues for changes or delegations that you might be requesting be effectuated, absent notifying potentially-impacted parties. We have grave concern over that.

FRANCISCO ARIAS: Thank you. This is Francisco. Regarding SAC 45, I would just like to point out that that's included in the Applicant Guidebook as a note to the applicants that they should be aware of those issues, and I think they may [inaudible].

DANNY MCPHERSON: And consumers, and liability? Anything?

FRANCISCO ARIAS: I think that there is a part that has to be in the applicants [inaudible] here and they should be part of the solution, and that is what is in the Applicant Guidebook.

LYMAN CHAPIN: So I'll just say that the volume of VeriSign reiterating previous recommendations because of liability is certainly something that lots of folks in this room that are paying attention here, and I don't think we can restate that enough. I think that it's public interest, and we have obligations there as this community before we impacted other consumers and users.

JONATHAN ZUCK: Hi. Jonathan Zuck from the Association for Competitive Technology. My background is coming from the business where I would have received the calls from all these corporate network administrators when things started to wrong with access to their HR system, 401Ks, etc., when these local queries start getting resolved externally. So I feel this in a very real way.

One of the questions I had about the data in this particular slide, actually, you said that there's an appendix or something that explains the reason the data might not be accurate, but one of them seems to be the various types of cashing that happens to these requests along the way out to the root. I was wondering if there are some tests that could be done that could then be extrapolated from to at least come up with some estimates. That seems to me that the amount the caching takes place is such that this issue could be very much larger than is reflected in these studies of caching taken into consideration.

Finally, I guess I would like to just recommend the release of these studies, even in draft form. I don't think the community needs the prose to be massaged, etc., unless there's really data inaccuracy issues. Let us start to engage with you sooner rather than later. I think this a really prominent problem. I'd say it's true for this study and also for the dotless domain study that we haven't even heard about. Let's get this stuff out there. Say it's in draft form, but get it shared with the community. There's no reason that the staff should be a filter for this information given the severity of this. This is like, you say, risk is any assessment of percentage of likelihood and high consequences, and it's like flood insurance. The consequences on this are potentially high, and even the infrequent occurrence could be very severe. Let's get this data out there. Let's let the community help to digest it, etc., and stop massaging the prose. Thanks.

FRANCISCO ARIAS: Thank you for the comment. I just want to clarify we're not massaging the prose or anything like that. We're just trying to ensure that we have a sort of report that was put out for comment. Particularly on the namespace collision, the kind of things we're looking at to have more details is, for example, the options we're trying to have more detail. For example, what will entail to have a [inaudible]. What will need to be done in order to have something like that I think is a very complicated subject and we're looking to have more detail before we can polish.

Regarding the dotless domain study, I should mention that it's also another study that was intended to be released before Durban. However, again, trying to have a more solid report, we expect to have that published, again, shortly after Durban within [inaudible].

JONATHAN ZUCK: I guess I'm suggesting release them in draft form, but don't open them for public comment so they can at least begin to be looked at and consumed, etc. That's all I'm suggesting. Thank you.

JEFF BRUEGGEMAN: Hi. Jeff Brueggeman with AT&T. Thank you very much for the very clear articulation of these issues in your study. It raises two questions for me. One is, with respect to the issue that you raised here, and given all of the comments that you've heard, it seems to me that at a minimum, this is saying we need further study, and then we're going to need a clear implementation plan. So does SSAC support having some kind of an urgent working group to figure out how to implement, whether it's some combination of change in the new gTLD program, or figuring out how to do the outreach and the notices that are required? What's the handoff going to be as the next step? What can we do to help support the fact that there urgently seems to be the need for much more activity on these issues?

My second question is about the interdisciplinary recommendations that you raised in SAC 59. Is there a clear next step there? Having served on the SSR Review team, it was something that we certainly had supported of doing a comprehensive SSR review of the new gTLD program and all of the implications, and I think the SSAC letter was very consistent with that. But again, I have a question of is there a clear next step to follow through, and if not, does SSAC support that, and can we do something to help get this going, because I think the questions on those are, "what happens from when the studies are issued to follow through and implantation?" Thank you.

FRANCISCO ARIAS: Sure. So in terms to what is the follow up to the study [inaudible] public comment so we can assess what should be done for when the comments from the community. I believe there was a question about what SSAC should be involved with. Do you want to address that, Patrick?

PATRICK FALSTROM: So SSAC's involvement, as many of you have seen, is that many of these studies are based on various recommendations and reports that [inaudible] has issued. Specifically regarding the namespace collision, we in SSAC are requested by the board to corporate with – I don't remember exactly the actual wording of the resolution off the top of my head – but basically we are asked to follow specifically the work that is now ongoing. They deliver the report to the conclusions that are drawn from it, the feedback from the community, and we are asked by the board to report back of what we see of the process, whether our recommendations actually are taken into account in a way that is satisfactory to us. We are in the midst of that, and to some degree, many of these reports are new to us as well, so we don't know why or what the actual outcome would be, but that's one of the reasons I'm here and why many SSAC members are in the room – to listen to your comments to be able to write a much better conclusion later on.

But we in SSAC don't feel like we're part of the flow from the report to the actual implementation. We're sitting on the side looking at what is happening to come with recommendations on how to potentially do adjustments on what direction this is going.

STEVE DELBIANCO: Steve DelBianco with the Business Constituency. What an incredible difference in this session than the SSR session in Beijing, because it's apparently moved from denial and defensive to data and discovery, and now we're having discussion. So I'm really gratified for that, and I can't wait to see the studies.

My question for Jeff: what would advise the CEO to say, because he's constantly confronting the public all around the planet, and up until now, has been saying there are no significant security or stability issues with the launch? I'm sure he looks to you for guidance on, "What can I be saying right now?" He may not be able to wait until all the studies are done and the public comment. But if he's giving a speech tomorrow, what would you tell him to say about SSR?

JEFF MOSS: Oh, just that it's one of the number one priorities of ICANN. It's our core mission statement. So we are not going to do – I'm not going to recommend that we do anything that has any substantial SSR impact. That's not a worthy risk. So I'd just keep reiterating that, if find any sort of showstoppers – if we find anything that suggests impact that will impact global DNS – we don't do it. It's not worth the risk. So I'd tell him to just reassure people that we're doing studies for a reason, and we're not going to take any risky steps.

People sometimes get hung up on the deadline, but how will you know before a deadline? Well, deadlines can move. If there's something that we find that we think's a showstopper, the deadlines will have to move.

BRAD WHITE: It's more finely grained than that. I'd encourage you to think of about other vocabulary. "Showstopper" is like a binary go and no-go. It may be the show goes on, but some actors are not allowed to be on the stage until we figure out how to fix the problem.

JEFF MOSS: Correct. Well, if you look at the work Lyman did, when you're talking about the potential – I alluded to it earlier – you could potentially block, say, a second-level string and allow a TLD to still be delegated. So there's many options, and I think it gets complicated when you try to apply a single measure across every possible impact.

MIKEY O'CONNOR: I'm Mikey O'Connor. I'm a member of the ISP Constituency, but I'm speaking in my own capacity. I've got a couple of things. I'm mostly into looking forward into the what's-next. So first, an offer. I have a source of data that isn't quite the same, but it's corp.com, which is a second-level, and I can tell you that when I tog along any kind of record in the DNS, I get a boatload of traffic that I have no idea what it is or what to do with it, but I can saturate my link to my server in about 15 minutes as it propagates out over the net, and I would love to give that data to somebody who knew what to do with it, like Lyman, and see if we can go another layer down. That might be a source of data that's a little more current, or that could be spread across a week. So anyway, just a little offer – a toy to play with.

I think the main question that you've heard from the last few folks, the answers were getting, at least to me, are not terribly satisfying on what's next. I've observed elsewhere that there seems to be a broken pipe between SSAC recommendations and action. I think yesterday I made a funny commentary that I'm going to save for the public comment period, but we really need a better plan than one side of the house saying we've got SSAC involved, and SSAC on the other side saying, "We're standing on the side."

So I think we just heard and saw that same broken pipe. So when Jeff Brueggeman is saying, "You know, we saw this is the SSRT, and when SSAC says, "You know, we say this in 2009 and 2010," the answers we're getting right now are not satisfying." So work on that, okay? It's just a comment – no need to respond.

ROY ADAMS: My name is Roy Adams. No SSAC hat this morning. I have one question and two statements. I'll start with the question so you can think about the answer while I do the statement.

The question is, you have a list of thousands of top-level domains being queried for. In the slide, you showed the top 15. Could you let me know, or if that information is available, what is the first undelegated two-character ASCII code? For instance, for a country code, that's not that delegated. The reason I'm asking that is I see an enormous amount of concern about new gTLDs branch strings being deployed, etc., etc., whereas ICANN, or IANA often in a manner, has been doing this for a couple of years on this two-character strings, and I've never seen such a study being deployed on this two-character strings. So that's my question.

The statement is: when you publish the reports, I understand you're going to give a reference to the source of the data, which what I understand is DNS-OARC. I would like to let people know that this data in DNS-OARC is actually available for others as well to analyze, so in order to reproduce the results that you have, or you can do a more fine-grained study of what you've done, and in that data set, there's not just the root server data, there's also a data for my employer, which happens to be [Nominas], so there's UK data. I understand there's also .net and .com data in there, and various other top-level domains. This is in order to Paul's question about doing this stuff on the [inaudible] that he had before.

The second statement is about the TTLs. You currently have the graph on the screen, and I think it's very accurate. It basically exactly says what you've done. This is not a statement about TTLs. I know how TTLs can influence things. For instance, if it's an invalid top-level domain, as an [impartial], the TTL has no influence. It cannot be cached. If it's an undelegated string, you get a no-such-domain, TTL is one-day. The cache is going to have to adhere to that. If it's delegated, it's two days. But most every application – not every cache, not every system has to adhere to the TTL. So if we do the TTL game on that graph, then it's going to be very expensive because that's an enormous amount of data to analyze. Thank you.

LYMAN CHAPIN: Just to answer the question asked at the beginning, we do see in the data stream some of the two-character codes. The first one that we see is way out in the tail, so it didn't show up on any of the summary graphs that I've put up before, but you will see it in the – the report is 198 pages long because it does contain tables of these data that go all the way out to the end of the tail. But it is there.

BRAD WHITE: Let me just add that we should be shutting down now, but I've talked to the technical guys over there to ask them how much time they need for a room reconfig for the next session. Bottom line is we can run a little bit long – five, six minutes or so. I don't know if that will give you enough time, but we'll stretch it out as much as we can. Again, if we have to cut it off, public forum is a venue open to you tomorrow. Sir?

PAUL MUCHENE: My name is Paul Muchene. I'm an ICANN fellow. What didn't come out clearly in your presentation, or perhaps I wasn't full at attention, was did you actually look at potential collisions with respect to using IDNs to nationalize domain names, and if so, what are your findings? Thank you.

LYMAN CHAPIN: Yeah, we do see almost all of the IDN-proposed strings in the data stream. They're not in the top 15 or 17 that made it onto the slide, but they are in the data stream and they're part of the study. So you'll see them in the report. The first IDN sting occurs pretty far down in the list, but almost all of them are there.

KEITH DRAZEK: Good morning. Keith Drazek, Verisign. I just want to say that I think that, in light of the results of the study, once we have an opportunity to review the actual study itself and to sort of take it in and evaluate it, if there are risks and we move forward with the delegation of new gTLDs, I think it's critical that ICANN – the organization and the community – really together develop an outreach plan, an education plan, to inform anyone who could be potentially impacted in a negative by the introduction of new gTLDs, whether it's name collisions or anything.

When you're assessing liability, responsibility for any potential impacts from name collisions or any potential negative impact from the introduction, it's really critical that ICANN, the organization, take a leadership role in educating the end user community – folks who don't participate necessarily in this room or in this venue and don't know anything about what we do, but could actually be impacted in a negative way. I think it's important that people understand that if things go wrong, if things go badly, that there is an information resource available to them, or that we notify them ahead of time to the extent possible.

So I thank all of you for the work that you're doing on these SSR issues. I thank ICANN staff for having this session, having the session that we did in Beijing. This is critically important, but looking ahead, when we're ready to delegate these new gTLDs, we should all work together to make sure that if there are going to be negative impacts that are known and have been potentially known since 2009/2010 based on the SSAC studies, that we work together to make sure that the people are aware of those issues so that we're not transferring risk to people that are unwitting or unknowing. Thank you.

LYMAN CHAPIN: I'd add to that that, not just putting people on notice, we need a mechanism to receive complaints and possible concerns in case something is occurring. To that end, ICANN published our Coordinated Vulnerability Disclosure Process a couple months ago. That process is being developed or augmented to include a 24-hour capability to escalate questions or concerns not only from root operators, but from the community as well. So if somebody is seeing something unusual, there will be a 24/7-hour capability for us to ingest it, look at it, analyze it so we can determine next steps. So that would be the other side of notice is actually do something when we're on notice when something is potentially occurring.

AMY MUSHAHWAR: I'm Amy Mushahwar with Associate of National Advertisers, again. I just encourage everyone at the table to think of this just not as an SSR issue, but a health, safety, and welfare issue. We had – the ANA had – a webinar with our CISO group where it was 1500 CISOs from very large companies, and CISOs are not aware of ICANN. They're not aware of how ICANN impacts them. These will be the groups that deal with the operational security that helps you resolve these issues. I mean I know we're very good with networks and networks understandings, but ops and servers is where you need to be to help resolve this issue. Consider that you might need months of delay for publicity of this issue.

So I would very carefully dovetail with what Keith had said – the communications portion of this will have to be very carefully monitored, and it will need to be more than ICANN's usual sticking something on their website. You're going to have to touch bodies, and you're going to have to do a lot of publicity. Thank you.

BRAD WHITE: So we're out of time. Gentlemen, thank you very much. I think this was a very good session. Again, if you have questions, we concentrated a lot on the study. If there are other broad-based questions, ask them at the public forum tomorrow. Gentlemen, again, thank you.

[END OF AUDIO]