BEIJING – New gTLD Security Stability & Resiliency (SSR) Update
Monday, April 08, 2013 – 15:00 to 16:30
ICANN – Beijing, People's Republic of China

Ladies and gentlemen, please welcome chief security officer, ICANN, Jeff Moss.

[Applause]

JEFF MOSS: First we need microphones. There we go. All right. We are going to get started. This is the new gTLD security, stability and resiliency update. What we are going to do here today is give you a view of how ICANN and the community At-Large is looking at risks related to the new gTLD program.

We have a number of people up here on the panel that are going to help provide their view and talk to issues. And then at the end of the session, we're going to also ask you to provide questions. And, hopefully, we'll have a thorough discussion and debate about several of the issues.

So in the end, I mean this to be very interactive. So because we're not going to take questions throughout the session, save them to the end. Write down your questions. Panelists will be here through the end, and then we will be able to address them.

Up here on stage, we have the vice president of gTLD operations, Christine Willett. We have the VP of IANA, Elise Gerich. We have the chief security officer at VeriSign, Danny McPherson. The director of DNS operations at ICANN, Joe Abley. We have director of senior -- or senior director of security, stability and resiliency on the security team, John Crain. We have the chair of the SSAC, Patrik Faltstrom. We have the chief retro phrenologist from Google, Warren Kumari. And we have the senior technology analyst at ICANN, Steve Sheng. So that's who you are looking at up here on stage.

I will just reiterate what we all know. The core mission for ICANN, security plays a very important role. It is called out in the ICANN bylaws as amended December 20th of 2012. And our mission is to coordinate at an overall level the global Internet system of unique identifiers and, in particular, to ensure the stable and secure operation of the unique identifier system. So if you look in our bylaws, this is like page 1, paragraph 1 and 2 stuff. It is right up there up front, and we take this very seriously.

As a global multistakeholder organization, we try to facilitate the security, stability and resiliency of the Internet's unique identifier system through the coordination, collaboration and cooperation.

And this is, of course, limited to functions within ICANN's technical mission.

So as I said earlier, we're going to concentrate on the new gTLD program. And if you think about it, there's sort of three roles here. The first bullet here: Managed within the ICANN organization, that's really where ICANN has an operational responsibility. The second for ICANN, the community to address, that's really where we have a coordination role.

And, finally, the global community to consider, this is really where we are trying to facilitate the dialogue and come to consensus. So it all ties back into our operate coordinate, and facilitate model.

Something I'd like to point out that's fairly obvious is that there is no such thing as 100% security. You can go broke trying to do it as well as insane. So nothing in the Internet is 100% risk-free. So what do you do about that? You try to manage risks. And as security professionals, everything is about a risk-reward tradeoff or a cost tradeoff.

So part of the situation is risks we know about, we develop mitigation strategies around. We have plans to address them. We assign costs, likelihood, impact, and we work to drive down the impact and the likelihood.

But we need to put some processes in place for unknown risks. We know things are going to occur in the new gTLD program that we cannot perceive. So we have to be flexible, and we have to be resilient and we need to come up with an adaptable process to help us address this, either ICANN as a sole operator or in coordination with the community.

So that's why I want to have participation of the community because when problems occur, it is going to be up to all of us to address the problem. ICANN most likely is not going to be positioned to magically solve all the problems.

The DNS we are talking about is a complex ecosystem, but it is not an unchartered one. We have decades of experience operating the DNS system. And if you look back, the community has dealt with several expansions. We've survived the addition of new resource record types such as AAAA records, TSIG, DNSSEC, the change in the protocol for EDNS0, introduction of IDNs, and we've done this for years. And the root system did not collapse when we made these introductions.

We've also dealt with the transition to -- in routing to BGP Anycast distribution of the roots to make the root system more resilient, and we have weathered that as well as a community.

So my message is whatever problems we have or do not have in the future, we will weather them together as a community and we will address them and we'll move on.

This next slide you'll see is an update of the slide Fadi gave on the new gTLD monthly update call where I really want to draw attention to the bottom here, sort of this brickwork of the SSR. And the point here is to show about at every level in ICANN, we recognize that security, stability and resiliency of the Internet is our number one core mission. And everything else is built on this foundation. So we are not going to do anything that puts the security, stability and resiliency of the DNS or unique identifier systems at risk. We're not going to do that. So with that said, Christine Willett just gave a session before this on operational readiness. And if you weren't there, I want to pass the microphone, so to speak, to Christine to give us a brief update and to speak to this slide.

CHRISTINE WILLETT: Thank you, Jeff. So this slide depicts the building blocks, the components, of the gTLD program building towards operational readiness. It starts with the applications that were published and posted last June. We completed our prioritization draw. We are in the midst of initial evaluation and we're publishing results for initial evaluation on a weekly basis.

As we look ahead, the additional blocks towards gTLD operational readiness, we move to the blue boxes, the contracting phase with a -- as we finalize and approve a registry agreement. Once applicants get through the contracting process, they move to pre-delegation testing. Once they've completed pre-delegation testing, the new gTLD program can say they are complete. They've completed their program responsibilities. We issue a notification to the applicant and we notify IANA that they are completed. And we issue them a token, an authentication credential, that they can then take -- the applicant can take to IANA to seek delegation.

We will be deploying the trademark clearinghouse. The verification system for the trademark clearinghouse is already live. And we will be building and implementing the sunrise and claims functionality, making those operational starting July 1 and with claims starting in August.

We have the URS, the Uniform Rapid Suspension, component of the program, which will be available and operational by the end of July.

The EBERO system, the emergency back-end registry operator service, which will be operational in August. And our SLA monitoring tools that will be operational in August. Those are the components of the new gTLD program as described in the guidebook that we're building. And that is our path, our timeline, to gTLD operational readiness.

JEFF MOSS: Thank you. So once operational readiness is complete and the applicant has processed through all of these steps, it now becomes time for delegation, a request to IANA to delegate the new gTLD. And to speak about this is the VP of IANA, Elise Gerich, who is going to talk about the delegation request process.

Elise?

ELISE GERICH: So I do want to thank Christine and the new gTLD team for doing all the heavy listing. From an IANA perspective, delegating a gTLD is much like standard operating procedure. There were a few things that we had to do to get ready for about five times the number of gTLDs that we have currently. One of them was to make improvements to the automation system that we have today. And that's called RZM. The RZM, root zone management, system was rolled out in 2011; and we have been working on enhancements with our root zone partners, both NTIA and VeriSign, to get this ready in time for the applications when they come as requests for delegation into the root zone.

Another thing that we worked on is to move from having a narrative report which we submit to say that someone has made all the requirements and met all the criteria to a checklist approach. And we've worked very closely with Christine's team so that that checklist is fulfilled during the new gTLD process.

And, finally, we've added an additional staff.

If we could move to the next slide, please.

So basically we've -- like I talked about the automation, we have implemented a new workflow that allows the creation of new gTLDs. In the past, we've had very few new gTLDs that were created. And so, therefore, that process was never seen on a high enough priority to automate.

And so, therefore, now we have a process which finished testing last week end-to-end testing with our root zone partners. And we're ready to go to production on May 1st. So if we could go to the next one.

Anyway -- and then I talked about the checklist and the report. So as the gTLD applicants go through the new gTLD process, they will be going through a number of evaluation panels which Christine has spent more than an hour today talking about and you all know better than I.

So, anyway, you will go through that process. And as you finish every panel, there will be a check in a little check box. And then when you finish it and all the checks are in all the check boxes, you will get a little credential at which point you will be able to take that credential and the check boxes will be online and you can request to be processed for delegation by the IANA department.

Can we go to the next one?

This is just a prototype. It is not exactly what the IANA Web site will look like. If you have really good eyes, you might be able to read it. Basically, what it does say is that on the bottom you put in your new gTLD string. If you're say -- the example up here says dot example. And below that, you type in the credentials that you received when you came out of the new gTLD program.

At that point, you'll have started the process for delegation. And in that process, the things that we do or the standard things we do for all TLDs, we do a technical check yet again to make sure that your name servers actually work. We do reach out to the contacts who are listed in your application to make sure they still exist. And it's standard operating procedure.

And I do want to assure you that we are ready. We worked really closely with our root zone partners, VeriSign, NTIA, and with the gTLD team and the automation system is ready to go live on May 1st. Thank you.

JEFF MOSS: All right, thank you, Elise.

So once the delegation is complete, I'd like to talk a little bit about the root server system in relation to the impact of the addition of all these new Gs that may be added. And this is to address questions around scaling, measurement, monitoring. And to talk about this, I'm going to hand it to Joe Abley, our director of DNS operations at ICANN. And I have a feeling he is going to pull in Danny from VeriSign.

So -- oh, I need to go back a slide, don't I? There we go.

JOE ABLEY: Thanks, Jeff. So what we have here is quite an old slide that's been circulated over a number of years. It came from Cisco. What it really demonstrates is the growth of traffic on the Internet is really dependent on the number of TLDs that we have. We don't necessarily expect a lot more traffic. We expect the traffic on the root servers to drive more traffic in general. We expect them to increase regardless of the size of the root zone.

What we have here is quite hard to read from the audience, what we have here is a decade or so of growth in the number of delegations from the root zone. I apologize for the size of the scale here. We see approximately 100 new TLDs added over the course of the last ten years, very modest growth. The graph looks far more dramatic and up and to the right than it actually is. The actual absolute size of the zone is still very small. So, obviously, there is concern despite the widespread experience with as many registry operators in the audience and the amount of people that run the root servers and running much larger zones than this and distributing it much more widely than the root zone. Clearly, the root zone is important. The root server system is important. So conservatism is important. We need to make sure the root servers remain stable.

JEFF MOSS: Danny, do you think that's a pretty accurate statement?

DANNY McPHERSON: I think that's definitely accurate. I think the one thing to observe in this is that over the last 14 years or so, there have been 67 or 68 delegations, I think, which irons out to 4 1/2 or 5 per year. What we are looking at now is sort of hitting the throttle and saying there are going to be 4 1/2 or 5 every 36 hours or so.

And so going from the pace that we've had traditionally to the new pace is certainly something we need to be cognizant of.

I think if you take a step back and look in 2009, the ICANN commissioned a study, a root scaling study, that looked at some of the capabilities and different implications on the root zone system to scale and what issues might arise when doing that.

And some of the things that came up were -- one of the key things certainly lots of folks in the community addressed is the ability to build some baselines across the root server system. So across all the roots, what sort of the latency, the number of queries, the consumer-concentric perspective of who's (indiscernible) in the root and how is it performing. And without those baselines, sort of stepping on the throttle is a little riskier than it may need to be.

SAC 46 expressed this, Recommendation 4 in SAC 46 talked about the need for an early warning system as well, loading up on the root scaling study. And the RSSAC in connection with the root ops have been doing some things in this area, but I think there are certainly still room for improvement here.

I think that --

JEFF MOSS: Wait. You can get your points out throughout the whole session.

DANNY McPHERSON: Sorry.

JEFF MOSS: You can get your points out throughout the whole session, not all at once.

DANNY McPHERSON: These are all related to the root server system, is that in the root scaling study and in SAC 46, there was a requirement up on ICANN that there be some visibility across the performance of the entire system. And I think a lot of the recommendations that SSAC made and the root scaling study team of experts made were move forward with new gTLDs predicated on the existence of that visibility.

We definitely need to get there before we step fully on the throttle. And somewhere in between there may or may not be reasonable. But I think it is definitely something we have long since recognized and discussed in the community.

JEFF MOSS: Thank you, Danny.

So, Joe, L-root has been collecting some statistics.

JOE ABLEY: Yes. So within RSSAC, there has been a draft set of metrics and measurements that is still under discussion. I think it is fair to say a majority of them are fairly stable at this point. The goal is to try and build a set of metrics which can be collected consistently by every root server operator and published so that long-term trends in performance of the root server system can be identified.

So the first step, as Danny mentioned, is to establish a baseline, to understand what is the baseline performance of the root server system with the current level of growth.

And then as we start to add new gTLDs, we continue to monitor that and look for any sort of trends that indicate that the system is struggling. As I said before, we don't expect it to struggle. But it is an important system. And we make sure that we operate it conservatively and responsibly.

So on the 3rd of April, we started publishing data which had been collected over the previous two months based on that initial as yet draft recommendation originally authored by Peter Cock, who is the IAB liaison to RSSAC.

As I said, it has been discussed within RSSAC for some time. Here is an example. If we back up a slide just once, the announcement that's referred there at that link has links to the live statistics, which are currently updated weekly. Anybody can go there and can track various aspects of growth in the root zone and also the various metrics recommended by RSSAC that we should collect. Again, this is just for L-root.

L-root is the first root server to publish these statistics, but we are aware of other root servers that are collecting these statistics. And we intend to publish them. We have an example here. Now, usually I think when people present graphs like this, they take great pains to point out the notable features on the graph. But really this is two months' worth of data. And what it demonstrates is the time taken to distribute a new root zone across the entire L-root system, which is around 300 Anycast nodes globally, including some nodes in some fairly unconnected areas, it is at worst case around a combined distribution time at around 4 seconds. Scale on the left is in milliseconds.

But no doubt as the months continue, we will see occasional spikes. We will see lines going up and we will see lines going down. This is just because we distribute the server over the Internet and Internet conditions change from day to day.

The important thing about this graph is not the features that you see right there from the first two months. The important thing about this graph is to track this graph as time progresses as we approach first delegation and as we continue to delegate in order to see what the distribution time for new root zones throughout the Internet continues to be.

So just to put this in context, we are looking at a four-second distribution time for a zone that's published twice a day. So I think four seconds in 12 hours is not too bad. That's where we are right now.

Here's an example from another set of statistics. We have two. One is the performance in L-root based on the RSSAC statistics, and the other set is the actual size of the root zone as measured in various ways. So this graph shows the size of the root zone in kilobytes, I believe. Bytes? Kilobytes? It is hard to see.

That step you see there in July 2010 is the signing of the root zone. It is the enlargement of the root zone due to cryptographic signatures and cryptographic keys and things of that nature. We have experienced net functions before in the size of the root zone. We can imagine that orange line there continuing its trajectory up at an increased pace as new gTLDs begin to be delegated.

Yes. I will hand it over to John Crain.

JOHN CRAIN: So I just wanted to briefly show some of the statistics that we have, or eyes that we have on the root server system. There is a program out there operated by the region Internet registry in Europe, called the RIPE NCC. And it is a global program. ICANN has been one of the sponsors of this from its early days.

The graph we're showing up above, that you can see up there, it is looking at the L-root because we like to show our own measurements. But they are doing this for all the root servers. And this is looking at the query times to root servers, how fast can I query a root server. Each one of these dots is not a root server. It is a little machine on the Internet that's asking questions. And this can also be tracked over time. This is all put away in databases. This is one graphical showing of the data. There are many other things that they measure.

So there is actually a few years' of this data in databases that we have baselines from. So this is not ICANN doing this or the root server operators. This is a third-party doing this. So this kind of data is out there to -- there is an URL on the bottom. And, of course, we will share these slides. I will give you some more URLs to show you other places you can get statistics.

So next slide, please.

So we talked about the L-root statistics. What we were really showing you was a formatting of the data, which is what the RSSAC document talks about, is what is the format of the data that we present so that we all give a similar view into the data?

JOHN CRAIN: I'm not going to give people a lesson on root servers. They can do that a different day.

But these are other letters of root servers that also publish statistics. To my knowledge, they all collect statistics. But these are actually public interfaces that you can go and look at or in the DNS-OARC, which is an operational analysis resource center for the Domain Name System, where there is a lot of data. This is not in the format from the RSSAC document. That's not finished, but the data is there.

Next slide, please.

We also have agreements to collaborate with some of the root server operators, not all of them, although we do all collaborate. For example, for F-Root, we have a mutual responsibilities agreement and with ICANN and we also have letters of intention and agreements of what they will go ahead. They are out there saying, This is our responsibility. We take seriously. And we are willing to collaborate together and do these things. Other letters just have Web sites that say this. These are ones that they have something with ICANN.

Next slide, please.

So if you have not met the root server operators, people who do this responsibility, they do come to ICANN meetings. There is discussion about moving the Root Server System Advisory Committee meetings to be more in ICANN meetings so you will see more of them around. But these are operational people. They are operating the DNS servers and they collaborate regularly. They actually meet three times a year in person to discuss operational matters.

JEFF MOSS: They also exercise their response systems.

JOHN CRAIN: They do exercises. They do all the things you'd expect. That's part of the collaboration to do things like exercises together. In fact, our friends from VeriSign are our biggest partners there in helping us fund those and actually do technical talks, et cetera.

We have seen all kinds of things in the past. We have seen change before. This Internet thing seems to still be working. So this is good news. We are not expecting things to fall over the day after we put a few hundred TLDs in.

Can we go to Danny?

JEFF MOSS: Yeah, go ahead.

DANNY McPHERSON: I wanted to add to that. You didn't see A and J-Root which are two of the 13 roots that VeriSign operates. And we definitely intend second half of this year to have public Web sites set up with the statistics associated with that. We do some very extensive collection of those statistics.

Regarding the contractual aspects, we do operate a root server agreement contractually with the U.S. Department of Commerce.

From an agreement with ICANN today, a lot of the letters of commitment as opposed to regulatory compliance frameworks -- I mean, from a CSO function of VeriSign, I have 1,385 different controls across eight regulatory compliance frameworks that are continuous monitored, audited from FISMA high security to SOX and so forth. They are audited and verified and continuously monitored. We are certainly all for some contractual obligations with ICANN and certainly that can be related to root operations and what's acceptable from a performance or data publication or other perspectives. I think getting those frameworks in place would certainly be valuable to have some baselines and common methodologies for infrastructure and collection and so forth.

JEFF MOSS: John?

JOHN CRAIN: Microphone was off, sorry.

So these maps here, these are actually root servers. It looks a lot like the last map. They are all based on some stuff that is out there that we're all familiar with, a map provider.

This shows you the extent of the root server system. So sometimes you hear worries about whether or not it has capacity. It has a lot of capacity. I can't tell you exactly how much, but I can tell you that L-root has 300 nodes out there. And (indiscernible) we operate. Each of them has plenty of capacity to do oday's load and quite a lot by 18 individual machines. And we have 300 and other operators have similar infrastructure. So there is a lot of infrastructure that's continually upgraded. This is today's status.

If we look at six months' time, we will probably have a lot more. Those who remember when there was a DDOS attack against root servers many years ago, we only had 13 physical locations at that point due to the technology.

This is evolving. The root server system will continue to evolve as the DNS and the Internet evolves and we'll continue to keep up and hopefully stay ahead of the needs and requirements.

JEFF MOSS: Thanks, John.

So now we are going to move into another section where we are going to talk about the ICANN's approach for dealing with those issues that were not known. These are things that come up, maybe catch us off guard, maybe there's a new issue with a new wrinkle. And we have to deal with it. We have to come up with a mitigation plan.

So for this next case, we are going to talk about the report many of you are familiar with, SAC 57. And this is going to really illustrate the way in which ICANN plans to -- continues to deal with these situations as they arise.

So I would like to hand it off to Patrik Faltstrom, the chair of the SSAC.

Patrik?

PATRIK FALTSTROM: Thank you very much. Before I hand over to Warren Kumari that will explain the details of SAC 57, let me explain a little bit how we came up with this report.

First of all, SSAC is operating based on actions that can be triggered either by external questions we get from the board or any other part of ICANN or maybe from the community. But it can also be the case that we have a self-initiated action where a person in SSAC wakes up in the middle of the night and starts to think about something.

This is one of those self-initiated issues. So, Warren, I don't know what you did when you came up with this. But he brought it up to SSAC, and we thought this is serious enough to actually talk about it.

Another thing that I also can say about this specific report is that, yes, we all know that these are the kind of things that should be detected when you do an investigation and you do the risk analysis that we just heard the people talk about earlier, but we also all know that are working with these kind of security issues, that regardless of how thorough or detailed you make these kind of reports, regardless of how much you assimilate, there are always different kind of things that you will detect. So it's very, very well important to be able to act.

So readiness doesn't -- isn't to be a hundred percent -- to be a hundred percent sure that you're safe. It is to be able to take care of things as they happen.

This is one example where we did detect and come up with this report.

And the third thing that we did a little bit differently than normal was that instead of making the report public immediately, we handed over the report to the ICANN security team that -- because we thought this was serious, so ICANN needed to have a disclosure policy to follow. And we'll come back into that how that was handled.

So there are various things that we were able to -- that we had to do in a new way because of this report, but just because we managed to do this, I think -- I claim that the community together actually are ready for moving forward.

So with that, let's hand over to Warren Kumari, and -- that will describe what's in SAC57.

WARREN KUMARI: Great. So I've got a lot of material to get through and not much time, so I'm going to be going fairly quickly.

Next slide.

So when we make an SSO or TLS secure connection to a Web server, basically something that starts with HTPS, your browser gets a public key and it uses that for the encryption. And it gets this public key in a certificate which is signed by a certification authority, and the certification authority's signature on this basically binds the public key to an identity. The identity is something like www.example.com. And when your browser comes to use it, it checks to make sure that the signature is correct and it's a signature from a CA it knows about. It checks to make sure that the certificate is still valid, it hasn't expired, and it also checks to make sure that the name it's connecting to is the same as the identity in the certificate.

So when the CA hands off the certificate, when they sign it, they first need to validate that they're giving it to the correct person. So the way that the CAs do validation, especially -- or I guess only for domain validated certificates -- is they send e-mail to an address at the domain that's being applied for. So it's (audio problem) example.com or the e-mail address that's listed in WHOIS, and this e-mail contains a token, and the person who receives the e-mail replies back to the CA and that proves that they control or (audio problem) own the domain (audio problem).

There is also another class of certificates called internal server name certificates and these are designed for things that are internal only. Hence, the name. They're often used by things like Microsoft Exchange, Active Directory, mail servers and a bunch of other things (audio problem). And the identity of these certs is of the form www.corp or www.accounting or mail.test. And the bit that makes the internal server name certificate different to a regular certificate is the fact that the identity doesn't get in the TLD. This means that you can't use the certificate on the Internet and it also means that the CA doesn't have a place to send a validation e-mail.

So what happens when the ending label on one of these internal server name certificates suddenly becomes a real TLD? So once this gets delegated, what happens?

The short answer is: Bad things happen.

Now, to demonstrate this, I applied for a certificate for www.site, and because I realized this was going to be validated by a person, I made up an interesting name, Dulles Steel and Forge, which (audio problem). Next slide.

And then I submitted my request to my CA and they popped up a little box saying "Warning: The common name www.site is not going to work on the internet. Do you realize that?"

I click "yes." Next slide.

And then three or four hours later, they e-mailed me my certificate. And you can see there the (audio problem) name is me. The subject is www.site. There are also two additional names, or actually there's a subject (audio problem) that has www.site and also just dot site. So big deal. I've got an internal server name. What can I actually do with this?

So to demonstrate this, I set up a fake instance of the root in the lab, I delegated dot site to myself, and then I set up a Web server (audio problem) and I browse to it in Safari and, sure enough, I get the lock icon and Safari says the certificate is valid. I mean, that's fair. It is actually valid. It works.

I did the same thing in Chrome and Internet Explorer and Firefox and Opera and a bunch of other certificates.

So what are the other implications of this?

Well, an attacker can go along and take those two applied-for TLDs and they can go and get certificates for well-known names in those TLDs. Then he just holds on to the certificate and waits for the TLD to get delegated.

Once this happens, he hangs out in his local Starbucks or coffee shop or a hotel where (audio problem) names or (audio problem) whole bunch of other (audio problem) type attacks and when somebody browses to a site that he's got the cert for, he presents the cert, the user gets the lock icon, and then he runs away with all of your money or your banking credentials or your cookies or whatever else he can get his hands on.

So (audio problem) an advisory and we had some recommendations that the security team should reach out to the CA/B forum, the CA and browser forum, a sort of industry group that represents the CAs, (audio problem) vulnerability disclosure policy on how to handle this sort of information, a communication plan to all affected parties, and a contingency plan (audio problem) before we had a way to mitigate it.

And I'll now hand it back to the security team who will explain how they executed this.

JEFF MOSS: Thanks, Warren. I want to pass it over to Steve Sheng.

STEVE SHENG: Thank you, Jeff. When DNSSEC raised this issue to ICANN in early January, we took this issue very seriously. So shortly after the teleconference briefing, we formed a cross-departmental mitigation team.

So the mitigation team met regularly to plan the mitigation steps in anticipation of the SSAC report.

So between January and February, that's what we did. We held several teleconferences, including with the CA/B chairperson, with the major CAs, alerted them of this issue, and they invited us to talk at the meeting in February. That's where we formally presented this issue to them.

They also took this issue seriously. I want to note that this is not a new issue. This issue about internal name certificates was pointed out to CAs as early as 2010 by the Electronic Frontier Foundation.

I think the CAs in this case, they act very swiftly. They advance Ballot 96 which I will talk about in the next slide.

So the key date is on February 20th, they have a voting procedure and the Ballot 96 passed. That significantly reduced the vulnerability window for this experience.

After that, SSAC finalized its advisory and as recommended by the SSAC, on March 15th, we notified this issue to all the new gTLD applicants.

Next slide, please.

Previous one. All right. So Ballot 96 really recommends the CAs to stop issuing internal certificates immediately, and within 30 days after ICANN has approved a new gTLD operation -- that means ICANN signs a contract with the operator -- the CA must cease issuing these type of certificates, and within 120 days after the publication of the contract, the CAs must revoke each of these certificates that ends in a new gTLD suffix.

So following up with the obligations, we set up a notification service to CAs to notify them what are the strings being applied for, and whenever a contract is signed between ICANN and a TLD operator, we'll post a notification to the CAs, so that -- to aid them to discover these time lines.

Next slide.

So there are some remaining risks on this issue, and in the next couple of slides we plan to talk about what these risks are and how we are planning to mitigate them.

And we want to invite the community to give us input as well.

So the first risk is we expect most certificate authorities to abide by Ballot 96, but not all the CAs are on the CA/B forum. So it is possible that some CAs will not abide by Ballot 96 until it formally becomes a (audio problem) requirement. So for example, for WebTrust, that's for North America, and ETSI, that's for European standard, the major browsers use those standards to -- all the CAs to include them in their root list.

So we -- the strategy for this is we communicate this risk and we're actively working with parties that can effect the changes.

So one of those parties we're actively working with is the browsers, to -- to get them into proactive action in requiring CAs to abide by Ballot 96.

So we think this risk is mitigated in the process.

Next slide.

Second remaining risk is for a variety of reasons, I think mostly for performance reasons, there are versions of some browsers do not -- do not check real-time for revocation.

So you could have a CA revoke the certificate but if a browser does not check the revocation in real time, there's still a vulnerability window where the certificate is -- still shows to be valid.

Our strategy is to communicate these risks. We already communicate this with the browsers, and we are in discussion with them how to best address these issues.

There is a variety of things being proposed, and we're in active discussion with them.

Next one.

And thirdly, it is still possible there exists a vulnerability window between signing a contract of ICANN and a TLD operator and the TLD operator activate second-level domains.

So this graph shows kind of a time line. If you start -- if you count, 120 days is roughly 17 weeks, so we start contract signing as week zero.

After that, we move to pre-delegation testing, and then IANA delegation, there's a sunrise notification period for 30 days, and then followed by the sunrise.

There are -- so I think all of this shows that there could still exist a vulnerability window on that, and on this issue, we really want to -- next slide -- to seek the community input on how, as a coordinator, this risk should be best mitigated.

So as Jeff mentioned earlier -- and I think Patrik, too -- that it's sometimes impossible to anticipate every risk ahead of time. All you need to do is to be ready to act and to have a process, so I'm going to hand over to Jeff on this.

JEFF MOSS: Thank you, Steve.

So I wish this was a little larger, but this is essentially a flow diagram of our coordinated vulnerability disclosure process that ICANN has adopted and we published it last month and we applied it to the SAC57 issue. That was sort of our run-through, to make sure it worked and fine-tune it.

And so going forward, this is the way in which we will deal with vulnerability disclosures.

And there's a couple ways in which we may deal with disclosures.

So think about it this way: You in the community may find a problem with, say, a root server, root server software, name server, and you disclose that information to ICANN.

We would use this process to determine how we then disclose that information to the affected parties.

In another situation, ICANN, we might be the affected party. You might find a vulnerability in one of our Web services or our Web applications, so you come to us and say, "ICANN, I want to tell you about a problem I found. How is this going to work? Are you going to publish my name? Is it going to be transparent?" Maybe I don't want my name published. So we would follow this process in disclosing our own vulnerabilities.

So it's a generic way of dealing with notification of affected parties and coordinating between those who may not wish to be put in direct contact with -- with the affected parties.

Oh, I see.

And then this is a fairly new graphic. We published it in our SSR to the community, but in case you didn't get a chance to see it, this is essentially a visualization of how -- ICANN's overall approach to risk and communication with the community, and this is our final slide in the presentation and now all of you who have been taking questions, writing them down, this is your chance to ask us anything.

I know we have a couple representatives in the audience from the -- yeah, Danny, from --

DANNY McPHERSON: I just wanted to go back actually to the internal name certs before we take open questions. I had a comment on Slide 45, actually.

JEFF MOSS: Okay. Let's go back. Which one? That one?

DANNY McPHERSON: Actually 44, maybe, then.

JEFF MOSS: That one?

DANNY McPHERSON: There we go.

JEFF MOSS: Okay. Good.

DANNY McPHERSON: So one of the things that I wanted to point out here as well -- and I think Warren made this point in the ccNSO or ALAC meetings -- is that, you know, I don't think there's agreement -- and Patrik can correct me if I'm wrong, as chair of SSAC, but I don't think there's agreement on if this is acceptable or not in general, and certainly for SSAC but for the broader community.

Additionally, I think that the window of vulnerability is unknown, actually. It's at least until 2016 for some of these things because a lot of applications, as stated previously, don't actually support revocation at all, or if there's actually a man-in-the-middle attack, then an attacker would certainly stop those revocation functions from working, and I think the gentlemen from DigiCert, Jeremy, made that point in the ccNSO. He may be in the room to clarify if I'm wrong there.

One other aspect that I wanted to point out here is that -- and I know you're certainly aware of this -- is that there are four things you can do with risk.

You can -- you can avoid the risk, you can control or mitigate the risk, you can accept the risk, or you can transfer the risk.

And anything we do here that doesn't fully address this issue is ultimately unilaterally transferring that risk to the consumer, right? It's going to be the people that consume things within this namespace that are going to be impacted. It's going to be, you know, Warren's example of the person at Starbucks that goes on line to update some financial or some health records and gets man-in-the-middle attacked because of this new namespace and some things that we're introducing in the gTLD.

So I don't think there's any -- to Fadi's earlier discussion, I don't think there's any magic ingredient here that's going to solve this problem immediately, other than a lot of coordination and a lot of work in the community.

And I think we -- you know, we sort of moved forward leaps and bounds in the work that the ICANN security staff did in three months to get to where -- you know, the things Jeff pointed out. I think that was -- that was a tremendous effort, but it still doesn't -- and, you know, there's such a huge amount of residual risk and unknown -- unknown there that, you know, we're ultimately, if we forward with that, transferring that risk to the consumer and that's certainly a concern.

The last point that I'll make is that in the RSST study that Patrik and a number of other folks did, and in SAC45 and SAC46 and SAC57 and so forth, there's a lot of discussion about interdisciplinary studies. And what those are talking about is that the DNS enables users to access something on the Internet.

Ideally, it does that in a safe, stable, predictable, secure manner. And -- and so the, you know, users don't normally go, you know, on the Internet, access content in the DNS. They use it to get somewhere else. And so when those dependent systems tie themselves back to the global DNS for safety and security, it's an obligation for us that we don't, you know, unilaterally make changes to that system without coordination that would impact the security or usability of the system.

And so that's sort of the crux of the point that I wanted to make related to this, Jeff, so thank you.

JEFF MOSS: Thank you, Danny.

I also would like to point out we do have some representatives in the audience of other root server operators, the CA/B browser forum and CAs, so I'm hoping to get a lively discussion going with other experts that are part of this community.

So let's go on to the questions slide and go to the microphone on the left. If you could just mention your name and where you're from and your question, please.

JEFF NEUMAN: Yes. My name is Jeff Neuman. I'm from NeuStar. I just have a question for -- I asked Mr. Faltstrom yesterday a question. I think it was yesterday. I'm getting my days confused. Maybe it was two days ago. There was an SSAC presentation to the GNSO Council, at which point Mr. Faltstrom said that the SSAC was not advising -- as an advisory group to the board, they were not advising the ICANN board to delay or slow down the new gTLD program.

It seems like from the comments that Mr. McPherson made, you still believe that there's significant risks, and I think we've heard that, so I guess my question is directed at Mr. McPherson.

Is there -- do you have any concrete proposals on the table for mitigating this? I mean, so you brought the risks to your attention. So what do you think is the next step? How quickly can you get it done? And if you were running the project to mitigate this risk, what would you do?

DANNY McPHERSON: Fair question. I would study the problem space like has been asked for by an ICANN chartered study team of experts since 2009 for interdisciplinary dependencies.

I think it's interesting that in the last two or three months when people are seriously looking and saying, "Hey, we're going to be wheels-up soon in this new gTLD program, what are the implications of that," people are starting to say, "Well, what's it mean if this TLD is delegated and I use that internally and sign my network?" Or, you know, "What level of visibility do we have to the root server system as an aggregate, and do we have an early warning capability that would allow us to identify threats? Was there an assumption that we would have that when we were wheels-up?"

And so I don't have a magic ingredient. I do know that there are a lot of smart people in the community and a lot of excellent work has been done on coordinating with all those dependent systems that rely on the DNS and that, you know, that some work will need to be done there.

As far as time lines or anything beyond that, certainly that's a community effort to determine those things, and if there -- if there are delays.

But I think that for my organization, if this were my decision, I certainly would consider the implications of that and, you know, personally I wouldn't use a new gTLD if I thought that -- you know, for my own personal health or financial transactions or other things, I wouldn't use that if I thought these issues existed inside the infrastructure I was using.

I would -- I would use something more stable that we -- we know we have predictable, safe, secure performance with. So that's all I've got for you.

JEFF NEUMAN: Patrik?

PATRIK FALTSTROM: Yeah. So I would like to make a clarification. The questions you asked was -- were -- that I responded to you was twofold.

One, whether SSAC was working on any follow-up on this work, and the answer was at that point in time, no.

And the second question was whether we were doing any action based on the letter from VeriSign, and the answer to that was no.

JEFF NEUMAN: Okay. Just one quick follow-up, if I can.

JEFF MOSS: Okay.

JEFF NEUMAN: Is this with respect to all TLDs or are you just worried about a subset like dot site, dot corp, dot home?

I mean, I know the problem could be expanded to all of them if someone in theory were to do it, but are you really concerned with that or is this just a subset?

DANNY McPHERSON: So I think some may be more problematic. I believe absent the entire corpus of certificates issued by all the CAs, which you're never going to get, the -- well, that would be my guess -- then there are probably varying levels based on the usage of these -- you know, of these new gTLDs.

You know, there are also other dependencies. Some of the things that PayPal -- Bill Smith, I think -- pointed out and so forth. And so I -- I think the short answer is there are varying levels.

Warren actually did some of the analysis so he may want to comment on that.

WARREN KUMARI: So yeah. While doing some of the analysis of the SSL -- sorry, EFF SSL observatory data, which is a corpus of publicly seen certificates, we saw some for dot home and dot corp and the things that you would expect in the queried cc's from the root. But there's also a whole bunch for dos ads, which we couldn't figure out what those were. Eventually, it was figured out it was Active Directory services. But those were things that you wouldn't actually know unless you could actually see the certificates there.

So it's fairly much impossible unless you can get a representative sampling from all the CAs to know what's been issued.

JEFF NEUMAN: And internationalized characters, are they the same kind of threat or is it mostly, you're thinking, ASCII?

WARREN KUMARI: Not a clue, but I think the person behind you might be able to answer that.

JEFF MOSS: I'm going to go over to the microphone here first.

CHRIS WRIGHT: My name is Chris from ARI Registry Services.

Similar question to -- to Jeff's, but perhaps a little bit different.

So the session was great for bringing us up to speed on the SSR issues that exist around the instruction of new gTLDs. It's good to see all the information brought together in one place and made simplified enough for us to understand.

However, what I didn't get from this session that I wanted was an understanding from ICANN as to what the action plan is from here, what activities ICANN's going to undertake or has left to undertake to address these issues, what are the time frames for those activities, what are the metrics that we'll be using for each issue to indicate that we're comfortable and we have it under control and it's safe to proceed, what are the specific goals that need to be achieved, and ultimately, what is the overall impact all these issues have on the new gTLD program?

So to put that another way and perhaps a little less political than Jeff did, what's ICANN's position on each of these issues? Is ICANN of the opinion that each of these security issues has been adequately mitigated at this point and that the residual risk is at a satisfactory level to be accepted or transferred?

JEFF MOSS: So I'll start with that. I think it's probably safe to say, and correct me if I speak out of bounds, but the technical issues that were raised are fairly well understood by the technical community and none of them were big shockers. I think what's happening, though, is if you look back over the last month you're seeing continued incremental improvement on all of the issues. So, for example, if there was a concern that ICANN didn't have a system in place to notify the CAs when a delegation -- when a contract was signed. That system has now gone live several days ago. There was a concern that no EBERO providers had been selected. Well, they're selected. And so, we have a list where we're tracking every one of these issues and we have the remediation or the response plan to them and we're picking them off one at a time.

Now, the areas where ICANN is operational where we can control our destiny, then that's under my remit, and it's up to me to mitigate those issues. Where we're collaborating with the community, say for example on the SAC 57 situation, then we're doing a whole lot of collaboration. And Steve Sheng spoke about all of the work he's doing working with the major operating system manufacturers and browsers, but we're not going to be able to tell you about that until we come to some sort of conclusion with them. So it may look like, for example, that there's a pause, the browsers aren't doing anything, we reach a solution and then we announce, much like with the CA/B browser forum. So this is not the sole extent of what we're doing. I don't know if anybody wants to add to that. No? Okay. I'd like to go to Jeremy.

JEREMY ROWLEY: Jeremy Rowley from DigiCert and representing CA. But I wanted to commend you guys from adopting a vulnerability reporting mechanism and if you could enhance that and make it more publicly apparent then that would be better. I'm sorry, am I not speaking well enough? Because we discovered this issue not long before you guys came to us and we weren't sure where to take it because it really didn't count as an objection or anything like that. So if you have a better place to report that, I think it would -- it would make it easier to identify these kind of issues.

To answer the previous gentleman's question, it's actually a problem with every gTLD since an attacker can get any domain they want, provided that they can meet the requirement -- meet the browser requirements which simply say that you have to show control over that domain and they can because they have the box with the server on it and it's not a gTLD. So yeah. So it's a potential for every single domain. However, we've seen the top four domains that are problems are dot corp, dot ads, dot mail and dot bank. So there's that problem.

And then I just wanted to comment briefly on the revocation problem, and I think that that one is primarily solved if people will start implementing OCSP stapling because you can't block the revocation response on that.

JEFF MOSS: Are there browsers that use OCSP stapling?

JEREMY ROWLEY: Yes, all the major browsers support OCSP stapling, I believe. It's a matter of turning it on on the client's -- the server side which IS has it I believe enabled by default and then you have to turn it on Apache and Genex. So it's a matter of not being turned on, which a lot of the industry groups right now are pushing towards recommending that as a solution.

JEFF MOSS: So maybe as a mitigation you would suggest that then we work with server manufacturers to get them to turn on OCSP stapling by default.

JEREMY ROWLEY: That would be great. Yeah. We're pushing on that, and if other people pushed for that it too, it would probably be adopted even faster. And that's all I have.

JEFF MOSS: Sir.

BILL SMITH: Bill Smith from PayPal. So yes, PayPal did send in a letter to ICANN about the -- basically the top 13 of these private top-level domains that represent on the order of 10% of the traffic to the root for DNS resolution requests. It's a significant number. And our suggestion was to not delegate these. You know that that may cause some concern here. We are also concerned about other names, but we are particularly concerned about those. And then the -- and that the -- the impact of transferring this risk from -- well, effectively nowhere today, these things don't resolve. They cannot resolve. The problem can't happen. And then suddenly we're going to turn a switch and 10% of the requests to the root for resolution potentially are resolved to the wrong place. Okay? And that this is a -- the practice of these private TLDs has been going on for decades. So that's -- we believe it's a serious issue. We actually commend ICANN and the CA/B forum for responding quickly. We are concerned that not enough, though, is being done. And I guess my question is really around, how do we ensure, between all of the players in these, you know, both the DNS system, the certificate authority system, and all the other systems that are out there, that this will work. And the last thing is, it's great that we're talking about browser vendors doing things but they are not the only people or applications that we use to connect to things on the Internet. So we still have all of those things to worry about. So this -- this is a -- it is a significant issue. It's actually one that's been known for quite a while. And, you know, my question really is, how do we ensure that we don't repeat these things.

JEFF MOSS: So I've got a question for you, and I'm just -- first, let me make one point. There has not been a final decision made as we responded in our letter to you that ICANN is still continuing to investigate the situation. So just free form thinking here, instead of it being a binary decision, delegate or not delegate, do you see any in between? Maybe not delegate for two years, give people time to remediate?

BILL SMITH: So I'll speak personally at this point because I know PayPal may have a different decision or opinion on this. I think the answer is yes. I think we don't have to be binary. But if we're going to do something, it needs to be more towards the don't delegate and take longer, at least on these 13, and we need to be cautious about the others to see what we can do. And the real issue, as Jeremy pointed out, is we have no idea, because there are no records that we have access to, of what certs have actually been issued as private TLDs for all of these top-level domains or any other string that might be out there. So this is a problem that's going to be with us for a while, even if we take action right away. So I -- I think there are choices between don't ever delegate and delegate instantly, but it would be closer to the don't delegate.

JEFF MOSS: Conservative side.

BILL SMITH: Be more on the conservative side, especially for things like dot corp. It makes sense. It's just common sense. Don't issue that one right now. You're going to get 10% of all requests to the root are going to be for those and they won't be valid. So if they resolve, they're going to resolve to the wrong place.

JEFF MOSS: Right. I -- first, does anybody else have a follow-up comment to -- nope? Patrik.

PATRIK FALTSTROM: Yeah, I just want to thank you very much for this comment and also that you -- that you participate in the dialogue because just like you referenced in your letter that (indiscernible) SAC 45 that we released on 15th November 2010 when SSAC talked about these issues and yes, we do have a dialogue with Jeff and others on whether it is the case that we and SSAC should redo SAC 45 or other kind of things to look into these gray area issues to investigate that. So I think it's really, really good in the community that all of you go the microphones, and if you don't have time, or are just too late or something, just try to reach out to us because this is why we're here and trying to solve this problem together. Thank you.

RUBENS KUHL: Rubens Kuhl, dot br. Question for Steve Sheng and possibly Warren. Have you considered browser modifications such as looking at this issuing date for the certificate. So if the issuing date for the certificate is before the actual date we know the TLD has been delegated, this was an internal certificate so we can now disregard the certificate because this domain has now been delegated. So is -- it is an option.

STEVE SHENG: Thank you for the question. When we interact with the browsers, this is indeed one of the options that's being proposed to us actually by our browser vendor. So we're working to see the feasibility issue here. Because, you know, there's -- we need to be able to push that to them, you know, when things are -- when things are delegated. And so we're studying the possibility of that. So we'll -- we'll have more update on that for you.

WARREN KUMARI: One thing to keep in mind is there is still a huge number of people who are running IE6 on Windows XP.

RUBENS KUHL: They have security issues.

WARREN KUMARI: They have many other issues. But you'd only manage to get people from a specific date forward. There are a lot of legacy things. Also, not everything is the web, yet. If you haven't looked through the RFC series, there are around 850 references to certificates and around 870 something to TLS. So there are all these other protocols like AAA and e-mail and Jabber and all these other things that rely on certs as well. And so you have the same issue with that.

RUBENS KUHL: Okay. Thank you.

MIKEY O'CONNOR: Hi, Jeff. My name is Mikey O'Connor. Most people in here know me as sort of this earnest plugger in the working group world. But I'll reveal another facet of my interest in ICANN. I own the domain name corp.com, and when I turn any kind of routing on for that domain name, I get a lot of traffic. I get so much traffic that in about 20 minutes I saturate my borrowed link at my friendly ISP. And I would be more than willing to point that traffic at anybody who wants to study what's coming in. I can tell you that it's not web requests. It's all kinds of active directory stuff, exchange server stuff, weird ports, unnamed ports that people are using for strange things. It's a horrendous amount of traffic. And that's not the one that's causing the trouble at the root. That's the dot com version of it. So I can only imagine what happens when you start routing dot corp to addresses out on the Internet. So I would just like to join, you know, Bill from PayPal, but I would also quite cheerfully offer a data stream for anybody who wants to take a look at what's in there.

JEFF MOSS: So you don't think that's something auto completing dot com on the dot corp?

MIKEY O'CONNOR: No, it turns out that in a lot of the Microsoft documentation for how to set up your little server, the default is corp.com. So that's just an example of the kind of problem that you're going to run into when you go to dot corp, which is by convention used that way. But, you know, for years Microsoft FrontPage, for example, resolved to company.com as the default, which I also happen to have at the time. So I've got a lot of experience with sort of the secondary effects of this. And one good example of the kind of problem that this could cause is that when I first lit up the mail server for corp.com, it took about ten minutes for me to get a misaddressed e-mail to joecorporatefinance@sun.corp.com that contained SEC filings prior to release, which caused me to run screaming over to the DNS server and turn all that stuff off because clearly that could cause all kinds of trouble for all kinds of people, including me. Well, I'm a good guy. But if a bad guy is doing that, there could be all kinds of things going on that, you know, right now we don't know a whole lot about.

JEFF MOSS: Thank you for offering it for researchers and not to some organized crime group because they probably would pay more than we will.

MIKEY O'CONNOR: Yeah. But anyway, I'd be happy, with all the appropriate safeguards, to offer that. But I want to sort of point out, this is a real world problem. This isn't some hypothetical that, you know, is perhaps not that big a deal. It's -- it's a lot of traffic.

JEFF MOSS: Thank you. Does anybody have any comments on that? Anybody? No? Okay.

ANDREW SULLIVAN: My name is Andrew Sullivan. In the previous session to this we had a presentation about pre-delegation testing that suggested we needed some sort of interaction between the people who are being tested and the testers because apparently this is not a completely mechanical process that either passes or fails. And here we just heard the observation that, you know, people are suddenly realizing that we're only a couple of months from wheels up and maybe it's time to start worrying about this. And I wonder if what we're discovering here is this is all quite a bit riskier than we thought it was. We're going to turn this on and people are only thinking about the real risks, you know, a few months before we're going to start to go. So I'm wondering if the panel has -- I mean, stepping up in kind of one level from this particular CA problem, which is bad enough, I'm wondering if the panel has something to say about sort of general risks in this area and whether we ought to be reconsidering our stance on what the risks are of delegating this many new TLDs without considering other effects. Thanks.

CHRISTINE WILLETT: I can speak in part to the pre-delegation testing comment. The desire to have more interaction between the testers and the applicants as part of pre-delegation testing was based on the need to make sure that the applicants are capable of providing the documentation necessary for beginning the test. It wasn't about a lack of automation or the automated testing not working. It was a communication issue, perhaps a language issue. So we're going through an effort to expand and become much more clear and specific about the documentation being required. I have not gotten any feedback that the tools or that the automated testing was problematic in any way. So I just wanted to address the pre-delegation testing portion of your comment.

JEFF MOSS: I'd like to point out that this is not the first time we've expanded the G space, this will be the third time, and all of these problems we're talking about have existed since the first time we added a new G to the system. I think what's happening here is the size of the routable space is increasing much larger than it has in the past. And now we're delegating names potentially like dot bank which might be more attractive to some miscreant than, say, dot info. But the problems we're facing now are the same problems we've always had. We're just a lot more aware about them now. Danny.

DANNY McPHERSON: Yeah, so I wanted to respond to both Andrew and Christine. We -- you know from an operational perspective we've got a great deal of infrastructure and institutional knowledge associated with operating registries at my company and we found a lot of the documentation for the pre-delegation testing problematic, to say the least. I think actually a lot of that, for those of you that are interested, is reflected in our February 8th letter and a March 18th letter and exchange between Akram and the registry stakeholders group that cataloged a large number of issues there that are being dealt with but there are still a lot of residual aspects of that as well, just for pilot pre-delegation testing, for example. I think simply to delegate a zone inside of our company, we have 110 tasks that are tied down to the minute, right? It's like here's exactly what happens for each aspect of this. And that sort of rigor is important when we put sequence deployment plans and project planning together for our own internal operations. And so we certainly expect the same kind of rigor in the pre-delegation testing aspects of this. And I think it's certainly getting better and we're delighted to see the improvements and look forward to more of those.

To Andrew's, you know, sort of meta point, I definitely agree with that. I think it's simply that, you know, until you make an expressed effort to take a step back and say here are the dependencies between -- well, here are the implications on the changes we're making in this system to systems that actually, you know, users want to access content or eyeballs on on the Internet or that sort of thing. Until you expressly take that step back and start to look at those, you don't see a lot of these things. And I think that that's exactly an artifact of what we're seeing. And what seemed like a late hour but it's because people are finally saying hey, here's the implication of my operation of this thing because quite frankly not everyone follows ICANN or the DNS or the IETF or whatever, they participate in their daily operations. For us and for most -- many of the people in this room, network and DNS is what we do. That's our business. We focus on that. We care, we pay attention. But for people that, you know, don't do that, it's completely different landscapes.

JEFF MOSS: Joe.

JOE ABLEY: So I thought I would just comment on one aspect of the scaling as far as the root system is concerned. I think it's important to note that we've had much bigger structural changes in the past. What we have going on here is a growth in the root zone to a size that is large compared to now in relative terms but in absolute terms is very small. We're not talking about adding new resource records, we're not talking about changing the priming period, we're not talking about changing the protocol and returning signatures or anything else like that. We're talking about business as usual in the root zone with a root zone that's slightly bigger.

If we imagine the current root zone is here and it might end up here, we already have in the room a breadth of experience with dot info being up here and dot org being up there somewhere. All using the same protocols. All using the same software. So while again, I think it's important, I think Danny certainly would agree, that we take all the precautions necessary to establish this baseline, do this measurement, and track the performance changes as we grow the root zone. I think as far as the changes go in the root service system, these are actually very modest. We've heard other things here that are not to do specifically with the root zone but, you know, one of the things we're working on is on that very important system. And I think, you know, we have -- Danny can certainly comment on this, but I think we expect the reaction of the system to be negligible to this proposed growth.

JEFF MOSS: Patrik and then Warren.

PATRIK FALTSTROM: Yeah. Just because -- with a different (indiscernible) I'm also running I-root and let me just add what Joe was saying that we also not only running the root zone but also several ccTLDs which are -- all of them are like bazillion times larger than the root zone so this is absolutely not any problem whatsoever.

JEFF MOSS: Warren.

WARREN KUMARI: So actually responding to something that Jeff said, yes, we have launched new gTLDs before. But when we've done that we've often run into some sort of issues. For example, Ram has a thing on dot info and how people didn't actually recognize it when it launched. And it seems the sort of concerns we're seeing now are actually interactions between the DNS itself and other systems. So sort of the interdisciplinary studies that SSAC has been asking for.

JEFF MOSS: These are the issues around universal acceptance, in ICANN language?

WARREN KUMARI: Well, that was part of the dot info thing that Ram was talk about. But it's more also an interaction between the DNS and the applications that expect to consume that. So yes, that was actually universal acceptance but it's sort of also a cross application issue. And so I think universal acceptance now might be close to solved but it's more what other things that rely on the DNS are going to become unhappy.

JEFF MOSS: So in those situations, I think you can expect to see ICANN operating more and more as a facilitator or coordinator where we don't control how python resolves or how a Microsoft application looks up a name. But we can certainly point out to them, we've determined that there's an issue here and it would be great if you would address it and we have experts that are willing to spend their time to help you. And so I think as the things that ICANN operates and has direct influence over are dealt with, you'll see us moving out more and more dealing as a coordinator and a collaborator. I'm sorry. Sir.

PAUL STAHURA: Yeah, my name is Paul Stahura. I'm with Donuts. And I read with interest the SSAC report that referenced a survey that was done in I think 2009 or 2010. And then I --

JEFF MOSS: The EFF SSL observatory report?

PAUL STAHURA: Yes. And I thought that was kind of old. You know, because 2013, that's like four years ago.

JEFF MOSS: Yeah. So they just haven't updated.

PAUL STAHURA: So I did one.

JEFF MOSS: Oh, good.

PAUL STAHURA: So I'm here to report the results. So we looked at -- we looked at the com and net zone, we took every name in there, and then we looked at every server that those names point to, and we observed 25 million certificates. And of those certificates, we looked through each one and we found out what the top level label was. And we found that there was 51 new -- proposed new gTLDs in this round that appeared in those 25 million certificates. 51 TLDs appeared in those 25 million certificates. And we also concluded -- we found out the biggest one was dot corp, the same as that survey reported. And then we looked at the sub domains that were mentioned in dot corp, and we found 102 unique sub domains out of all those giant space of dot corp names that could possibly happen, 102 sub domains. And by the way, we found home was the second one with 42 unique sub domains. Offline was number 3. Inc. was number 4 and so on. So I have the list of dot corps here too. So the dot corp we found, for example, park dot corp, digi love dot corp. You know, there's 102 of them. And so why don't we just block those for a period of time, maybe until freakin' 2016 as the VeriSign guy said. We just block those sub domains, whoever gets dot corp, we don't have to block the whole dot corp name space. We could stop issuing these certificates now and then block some until sometime in the future?

JEFF MOSS: Thank you.

DANNY McPHERSON: Can I respond?

JEFF MOSS: Yes.

DANNY McPHERSON: So Pete?

>> Paul.

DANNY McPHERSON: The interesting artifact here is that in SAC 57 what was highlighted with the 37,000 from 2010, I think it's excellent you actually did an update. I would love to see the results of that published somewhere.

>> Sounds good, I'll send it.

DANNY McPHERSON: But the thing I wanted to point out is that those are actually internal name certificates and they're not actually supposed to be used on the Internet, right? And so the ones you find on the Internet are the ones that people are leaking out to the Internet. So that's absolutely a lower bound and you've got to expect that people that configure their systems correctly have orders and orders of magnitude more of those. So without the CA corpus, it's impossible to identify the extent of what certificates have been issued for what streams.

>> I agree with that but it is a pretty big sample size.

DANNY McPHERSON: For Internet facing internal certificates, totally get it. Absolutely.

JEFF MOSS: Jeremy.

JEREMY ROWLEY: Yeah. I want to say that Bill Smith's problem that he identified is related but it is different from the CA problem in that most of these domain -- a lot of these networks are set up so that they are internal, so they're not going to be resolvable even after you make the domains resolvable. I mean, these are internal networks. So you're going to have a lot of conflicts on the Internet with all of these internal servers that are set up and things like that. People are going to go to coffee shops and think they are going to the dot corp mail server and they're going to end up going to some new gTLD they didn't expect. You're going to have a lot of those kind of problems. And that's different than the -- the certificates that are issued for these. Because we can -- CAs can kind of take care of the certificate issue, but what I want to know is what you've done to reach out to these networks that have configured their -- that have centered their entire operations around these internal networks to get them to reconfigure so that they're not going to create all these problems. And I'm not sure if any outreach has been done yet there. And you guys did identify it in the report but you've reached out to us, the CAs, but there needs to be some kind of reach out towards them. Have you done anything yet on that?

JEFF MOSS: Does anybody want to take that? Patrik, and then I'll go.

PATRIK FALTSTROM: I think what you're pointing out is very valid and even if it is the case that one of those -- one of those domains are really covered internally, just like you're saying, it might be the case, of course, that the result would be different, depending on whether the solution of the domain it happens on the inside or outside of that internal area, which is not only whether you're on your corporate network outside or also, for example, whether your VPN connection happened to be up or down which is something that might be a little bit flaky sometimes. So that is something that -- that we're looking at, of course, in SSAC at the moment, but it is not -- as I said earlier, it's not a work item that we have picked out. But given the discussion that we have here and it might very well be the case that we get one of those triggers that actually will create some kind of work for us.

STEVE SHENG: And actually I believe one of the recommendations in SAC 45 was actually outreach to those that might be affected, and I think that's still an open item.

>> I just wanted to say too up until 2011 the documentation was showing that you should actually use these internal server names for various things like BlackBoard and Exchange and so asking people -- the reason that 2016 was picked by the CA/B forum as the deprecation date, the complete deprecation for these things is we had a lot of customers that came to us, especially in the education space, that said that they couldn't make that change until then because they needed the time to get the money to upgrade their networks, to get their network

operators trained and things like that. So I like your idea of making it not a on or off type thing but saying yes, you can be delegated but you'll be delegated in 2014 or 2015 when we -- when the -- when we have a chance to clear out these internal server names certificates and we know that there's -- CAs stopped issuing on it for two years, you know there's not going to be a problem there anymore. You can even have the browser say yes, anything issued before this date is not to be trusted. And that gives everybody time to change over to make those -- and make that transition.

When the customer comes to you and says, "My cert is going to expire in two months. I'd like to buy a new one," they're, essentially, doing the outreach saying, "I'm sorry, sir. You're not going to be able to purchase that any more. You need to come up with a plan."

>> That was adopted in 2011, that requirement that all CAs have to try to mitigate that at the time. But, usually, they are coming to you about two days before the cert expires. And there, like, oh, crap. I've got to get a new cert on my server or else my entire operations are going to be insecure. So we give them a year's cert or something like that.

JEFF MOSS: So you've been communicating with them all along?

>> Yes. As of 2011 there's a requirement that all CAs have to make that outreach to try to get them away from that name.

JEFF MOSS: Good. Warren.

WARREN KUMARI: I should also mention, though, that getting somebody to change what they need the server name for to be fully qualified is much easier than changing all of their other infrastructure to rename it. So making machine be mail.corp, and then have a cert for mail.corp to food.com is different than changing dot Corp. everywhere. So that's a subset of the problem.

>> Actually, the cap forum did a study on why people are using this that maybe I can share with you. I'll see if I can get permission. Their results are interesting. A lot of people think that these internal server names are required. They don't realize they could use an FTD, and they don't know how.

JEFF MOSS: We have time for one more question.

CHRIS: Chris again. Perhaps I'm a little bit dense. But I'm wondering if you guys can summarize for me what are the actions that you want the community to take as a result of this? And what are the actions ICANN is going to take? Fadi came out and said he will stop the program, if there are security and stability concerns. And, obviously, there are a lot of people in the room that are concerned about that and don't want the program to be stopped. So how do we progress these things? How is ICANN going to progress these things, and what do you want the community to do to help?

JEFF MOSS: Anybody have a comment, or I'll take it? Yes, I had a feeling you would say that. So, like I mentioned earlier we're tracking all the risks that have been raised or there may be risks that we identified internally that didn't come from the community. And we constantly look at how we're going to mitigate them. If we come across an issue like the internal server issuer, cert issue, that has real-world, large-scale impacts, that would give us pause.

For example, if we could not have worked with the CA/B browser form, if we could not have mitigated it, that would have been a serious enough issue that it would have caused us to seriously consider whether we'd have to modify the program. So we need to take it on a case-by-case basis.

And what we're asking the community to do, if you think about the nature of the problems we're talking about, the internal cert issue is a symptom of the larger problem of expanding the rout-able address space. And I'm sure everyone on this panel cannot fantasize of every possible problem. But people in the community who work on this all the time, maybe you come up with a situation that we need to be aware of. So what I'm asking for, what we're asking for is a call to the community to tell us help identify any other issues that you may be aware of. I don't want somebody to say, "Oh, yeah. I've known about that for five years" and tell us two minutes before we're expected to go live.

So part of it is, you know, if you need to use our coordinated disclosure process, please do. If you want to call me privately and use a fake name, do that. But we're not going to turn away anything and dismiss anything unless we investigate it.

CHRIS: Understood. So, based on the language you just used then, am I to take away an understanding that ICANN feels that these issues up here have been adequately addressed?

JEFF MOSS: I'd say currently there's no issue that is a show stopper, because they're all being addressed right now.

DANNY McPHERSON: I would, as an individual and an operator, disagree with that, actually. I think that there are some significant residual risks that are being transferred unilaterally to users and consumers of the Internet. And we need to evaluate the implications of that. I mean, if you step back, one of the promises of the new gTLD program is that, hey, we can have gTLDs that are more secure that operate in a more secure manner. This is the inverse of that. My request is that ICANN look at that aspect of it. And I know we've moved leaps and bounds already from where we were. But I do not believe that what's been stated so far and what's been done adequately addresses these vulnerabilities.

JEFF MOSS: So which -- just for the community, which one is the show stopper?

DANNY McPHERSON: It's any of the individual ones and certainly the aggregate. None of these issues have been addressed. Having -- I'll let others on the panel -- I know I've made this statement, but we had a 90-minute discussion about revocation doesn't work. It could happen after this. You know, for other name spaces, what about these implications? And, ultimately, we can't simply, as a community and obligations and responsibilities to the DNS and the dependent system of the user DNS, unilaterally transfer that risk to those consumers.

JEFF MOSS: So your OSCP example, if the browsers start enabling OSCP revocation checking, if the servers enable it by default, that's going to be a fundamentally different situation than what we have today. Right? So it's not that there's no mitigations that exist. There are mitigations. Our challenge is to get the servers to make that a default behavior, not a, you know -- Steve?

STEVE SHENG: Danny, for the record, I disagree with you that revocation does not work. Thanks.

WARREN KUMARI: So I don't think we're going to be able to remove all the risk from this. The DNS is a large complicated interrelated system. And, when you make changes to it, something's going to go bad somewhere. We've seen things like that from roll over and die. We've seen things like that from when you blacklisted an entire list of name servers because somebody's blocked specific names. What we need to decide is how much risk are we willing to accept from doing this and who actually takes that risk? Who are we handing that risk off to? And are they ready to accept it, and are they even in the right position to accept the risk?

JEFF MOSS: Anybody else have any comments? If not, we're going to close on that. Elise? John? No? All right.

This our inaugural session of this. If the community -- if you think this was useful, I'd be happy to do this at every future ICANN meeting. So I'm just curious. Raise of hands, who thought this was useful? Obviously, we can refine it and maybe take questions in advance so we can more specifically address your concerns. Great. So I think we can look forward to doing this with you guys again in Durban. Hopefully, with some updates. Okay, Mikey, since we love you.

MIKEY O'CONNOR: I know. I'm waiting. There it goes. I'm not going to be like him and start whacking -- I know that drives you guys crazy on your headphones. See?

Here's a different perspective on this. In a way what the new gTLD program is a new product being offered by a bunch of vendors. And I -- I am an ISP. I'm on the pointy stick end of all this. Because, when this things breaks, they don't call ICANN. They don't call donuts. They call me. And, as an ISP, I'm requesting you to make sure that your product works well. Because the last time your product didn't work well, I got blamed for the fact that it didn't work. And I got to spend a lot of money on customer support calls. So this time around it would be nice if the product worked better. Now, the people who are offering this product are pretty keen to get a lot of cool revenue going in the DNS. And I think revenue is a great thing.

But this product doesn't feel quite ready to me. And the examples that Bill and I have been pounding on about corp.com are about one of many. So I'm in the camp that says let's not be too blase about this, and let's put some of the onus on the people who offer this product to help get it fixed.

If you talk about who should go talk to the browser vendors, well, I shouldn't have to go talk to them. ICANN shouldn't have to go talk to them. It should be the people whose product this is that should be trying to make their product a good one out in the marketplace. That's just a rant. Sorry about that.

JEFF MOSS: With that, thank you very much.