LOS ANGELES – DNS Risk Framework
Wednesday, October 15, 2014 – 15:00 to 16:15
ICANN – Los Angeles, USA

JACKS KHAWAJA: Hey, guys. Are we ready back there to start? Hello, everyone. For those of you that don't know what's currently being held in this room, it's the DNS Risk Framework Update. I think there was a slight change in the schedule, and we were just made aware of it maybe about ten minutes ago. If you're in the wrong room, you may want to jump over. I don't know where the other session was that was supposed to be in this particular room, but this is the DNS Risk Framework.

My name is Jacks Khawaja. I'm the Enterprise Risk Director for ICANN. On the panel, we've got John Crain. Do you want to introduce yourself, John? Maybe not.

JOHN CRAIN: Not without a microphone. Okay, yeah. I'm John Crain.

JACKS KHAWAJA: Okay. Today we're going to go over the DNS risks that were identified through a risk assessment that was conducted early this year. For those of you that don't know, about a year and a half ago the board and ICANN staff issued an RFP to identify a vendor to conduct a DNS risk assessment or actually developed the framework for a risk assessment. That was Westlake Governance.

Once Westlake Governance developed that framework and shared it with the ICANN Board, it was adopted. It was actually published for

public comment. Comments were received. Once the comments were received, we modified the framework. In early 2014, we conducted a DNS risk assessment. The risk assessment concluded in April. Those risks were shared in London with the community at an ICANN 50 session in London. If we can go to the next slide, please. Who's driving the slides? Oh. I could use it if it actually works. Cool.

A brief history as of today where we're at. From that framework, 23 risks were identified. They're across-the-board. They hit a variety of assets – from the networks to the hardware to different components of the DNS; from the equipment, the information assets, to the program itself. Because of that, those 23 risks were not reviewed in detail yet. That's what this session is about.

We're going to introduce the activities that we would like to conduct with the community in order to attempt to mitigate those risks. Again, this is going to be a collaborative process. As many of you know, there are a lot of asset owners – those that own those risks within the DNS – that simply are not within ICANN's remit or we have some sort of relationship with or that we can work with to mitigate those specific assets. When I say assets, I'm talking about the hardware, the software, the ISPs, the other vendors that put out software that may have some sort of impact on the DNS.

As you can see up on the chart – it may be a little bit hard to read – but within the calendar invite on the meetings.icann.org, you'll be able to find this model. This model was shared in ICANN 50 in detail. Part of that paper that was written for the DNS risk framework explained how the risks were derived. We reached out to RSSAC and SSAC. We had

technical experts come into ICANN where we identified these risks. They're certainly not in detail here in this framework, but that's intentional. That was by design.

A paper was issued in ICANN 50. This was shared with the community, and it was published in ICANN 50's DNS Risk Framework meeting from that particular session. We just wanted to go over the history so that everyone knows what we started with and where we're at now.

It was a consultation paper. That consultation paper was the basis for the community session that we held in London with the community. What we did was we went over the framework. We went over the needs of that session, which were to review the framework and insure that everyone was comfortable with what we had identified in terms of risks and how we approached the whole risk assessment itself.

If you'd like to read more about the paper and the risk assessment and the risks themselves, you can find that on either ICANN's website for ICANN 50 under DNS Risks or you can go to this actual presentation that sits in this meeting's a calendar invite.

The very first risk that we're going to tackle, and this is something that is probably very fresh on many people's minds, is DDoS: Distributed Denial of Service. I'm not going to go through the whole description, but we're going to scrutinize this risk and we're going to break it down and we're going to discuss that in detail.

The purpose of bringing this up, this right here has the description and any additional explanations of that particular risk. Why did I bring that up? Many people have different interpretations of what DDoS is.

Ultimately, it's dragging down a piece of hardware or service. That's essentially what it is. It could be multiple attacks. It could be any type of threat that drags down the Internet's bandwidth or server CPU and so on.

This was a good risk to start with, we felt. The risk that we showed you on the previous screen is the first risk on the list at the very top in dark blue: Denial of Service. This isn't prioritized in any order, but it would seem like an applicable risk to tackle up front.

As I described the background, it was presented at ICANN 48, the actual framework, then we went into the details from there with the community.

What I'd like to do now is hand it over to John Crain. We'll get into the details of DDoS and how we are going to approach the risk assessment and how we really want to collaborate with you the community.

JOHN CRAIN: I'm John Crain, Chief SSR Officer at ICANN. We're not going to get too much into trying to solve any kind of risks here because we're not prepared for that. I certainly am not. I do want to talk to you a little bit, quite a bit, about how we can actually move forward.

Typically when you examine risks and you're trying to mitigate risks inside of a typical organization, one of the things you do is you define assets and you define processes. You look at the risks to those. What are the threats against them, and how can we mitigate them?

When we did this risk assessment, we didn't really do that. We sort of went from the opposite end of the stick. We have a list of 23 risks, and we don't really have an asset list. It's kind of deliberate. We're looking at the risks to the ecosystem, but the ecosystem is rather large. If you try to define every single asset in the ecosystem, we'll be doing this until the next 500 years and we'll probably get no progress on actually trying to mitigate risks.

What I'm going to put in front of you is a thought I've been having about how we can do this. You can feel free to tell me it's the wrong way. But if you do that, you'd better have a better way because this is something we've got to tackle as a community.

There are 23 risks, and our initial thought is that my team (the SSR team) will take each of these risks and look at them based on sphere of influence. I work for the ICANN secretariat, the ICANN staff, so I looked at this a little bit from my own viewpoint. There are things that I can implement. We have a network; we run a network. If there's a risk caused by software on a router and I run that router, then obviously I can do something to mitigate it.

There are things that we have as an organization and as a community we have some direct influence over, then there's all the rest of it. I don't know if you were in the previous discussion where there was a lot of reporters and stuff. People seem to think that ICANN runs the Internet. Of course, we don't. We don't own those assets. We often don't own those risks. Therefore, it's really hard for us to even consider how we would implement any kind of mitigation.

Assets that we directly control that relate to the DNS, and this is just a first list. This isn't the white paper that I want to write up, but these are some of the thoughts of what should be thrown in there. We have all our suppliers. There's risks to our supply chain, just like every other business. We own some corporate infrastructure. For example, if we decided a mitigation to a specific risk was some form of authentication like DNSSEC, then obviously we can implement that on our infrastructure.

We also run some infrastructure that is used by the Internet itself and by the users. For example, ICANN runs one of the 13 root servers. If you take DDoS as an example, there are things we can do there and I'll come back to that. There's probably lots of other stuff that actually run. We actually manage a couple of TLDs: .int and .arpa are things that we have some of the infrastructure for. We can do things about that.

If we can identify risks to those assets, those are things that we as ICANN staff and our engineers and executives by spending money could actually influence. This is very typical corporate risk. Most of these will actually end up in our ERM process that Jacks runs.

That's fairly straight forward. That's pretty typical risk management at organizations. You look at what assets you have and you control and what assets you can actually do something about. It gets more interesting when you look at the stuff we can influence.

I'm kind of shifting a little bit with my definition of ICANN here because the first one is really the staff, but in some ways this could also be the community. We have contracts with operators of infrastructure, registries. In those contracts, we have things like SLAs. SLAs are a form

of mitigation of risk. ICANN's not going to implement those SLAs. Who here works for a registry? We can't go and implement things on your infrastructure. We can't go and tell you how you'd mitigate this risk. That's your job.

Through things like SLAs and sometimes policy processes, we can influence how those risks are mitigated. The same can be said for the registrars, but you have to remember it's their remit, their risk, and their mitigation to actually fix these things or to mitigate against these risks.

We can see how we can influence this sphere. We can do it through contracts, contract negotiations, we can do it through the PDP, we can do this through working with the community on best practices and peer pressure. There's lots of ways to mitigate risks when you can influence somebody directly.

Then there's a third area. I'm going to say this again – we don't run the Internet, we don't manage the infrastructure. The Internet is, by design, a network of networks. To many of you, this may seem very obvious, but you'll be surprised how many people out there actually think that ICANN can go and implement something on a name server of a registry somewhere or that we can go and do something inside ISPs.

People just don't understand that, and they need to start understanding that this is a network of networks. The owners of those assets, whoever they are, are the ones who can actually influence them.

We can do that through outreach, education, best practices – and by "we," I mean the community. We have some excellent examples of this.

There's an organization called ISOC that has a program called Deploy360, which is helping to get DNSSEC and V6 and things like that out there. That's probably not really our job at ICANN. We can help; we can participate.

So 23 risks, some of them a little nebulous. We could spend at least six months arguing about which one is the more important, which one to tackle first, and that's what we'd spend all our energy at. What I did is I decided let's just take the first three or four in the list and do an initial pass at them to find where the assets are to see if we can actually get more specific on some of the risks and whether or not there are mitigations that could be made or, in many cases, are already in place.

These are not new risks. DDoS is not a new risk. We've been aware of the problem of DDoS for years. We've done things like Anycast. Anycast is one of the mitigation methods, if done correctly, against DDoS.

It's not like all of this is stuff that's doom and gloom. What we find as we look at these, there's a lot of work already in place. Maybe all we need to do for some of these is document that, assess whether or not there is enough mitigation in place and the correct mitigation. If so, then fine. We'll just come and reassess it again maybe in a year's time. If not, then maybe we make adjustments.

These are actually quite large tasks. I think if we do the first sweep across the 23 over the next months, it's actually something that we can bite off and we can actually do and then we can come back to the community.

There's a description. I didn't write the description. I'm not going to waste my energies arguing about whether something's an "and" or "or" or anything like that. We're taking the descriptions as we have them through the process that we've already been through. These have already been in front of the community, they've been talked about to death, they are what they are.

What's the actual risk? In layman's terms, it talks about the probability that parts of the DNS will be either taken offline or made less available – slowed down, whatever – by DDoS, by large packet floods.

To figure out what the effect of that is, you have to actually understand where these assets are, what they're used for, and of course who owns them. That's what we're trying to do with some of the spheres.

We keep trying to look at the DNS assets from both sides of the DNS. The authoritative data service like the registries, etc. but also from the other side, which is the people asking the questions. The recursive servers are out there. DNS is not just registries and registrars. People actually use it for stuff. That's the stuff we really worry about.

When I did a first glance at this, the assets that we can actually directly do something about are really those things that affect the global DNS. Those things like the TLDs we operate and the root servers we operate. Obviously our corporate infrastructure needs to be up and running for us to do our day-to-day job, and we need to do it well.

We run our own recursives. On the recursive side, we run recursives for our stuff. We don't run a recursive service for the public DNS. It's not ICANN's job; it's not what we do. Not really much we can do directly

there. Obviously, with our own things, we'll make sure they're in order. The office recursive server in ICANN's LA office is hardly going to really affect the overall stability of the Internet.

When we get into the areas that we have the direct influence over, that's interesting. We have a few ways we can influence this through the policy and the advisory committees. We have SSAC, we have RSSAC, they do advisories. RSSAC is about to release an advisory on service level recommendations for the root servers. That's a form of indirect influence. We don't have contracts with the root servers, so we can't just go and say, "You must absolutely do this." We do have ways of advising, and this is one of them. It's going to be called RSSAC 002, and it should be out soon.

We have contracts in place with many of the registries and registrars. Not all of them. Another thing that people forget is that ICANN doesn't necessarily have contracts with ccTLDs. There are registrars out there in that realm that are not part of the ICANN world. They really fall into the next realm, which is where we can indirectly influence. These are where they are.

In the recursive realm, we don't have contracts or anything like that with recursive operators either. There's not really much we can do there even through sort of indirect influence like this.

The stuff that's really outside our realm or the things that we can only influence indirectly are the registrants have authoritative name servers. They run them in businesses. They could run an authoritative name server from their home if they want, and people do. There's not much we can do there apart from educate.

The recursive, that's the same. We're not really involved in the other two spheres with those. We can talk to them. We can educate them. By "we," I mean we as a community. Those advisories from places like SSAC and RSSAC are, of course, one of those forms of communications.

The other question when I look at a risk, beyond mitigating the outcomes of this risk occurring, are there other things we can do to tackle root causes? DDoS is a good example of this. Things like BCP 38, source address validation, open recursive resolvers – these kind of things are part of the cause, if you like, or part of the enablers of DDoS. Botnet is another one of these.

Should we be involved in working on these? Are there things we should do? It's important to remember that there are already a lot of things going on here. There are projects to look at open recursive resolvers. There are projects looking at trying to push BCP 38. Is there a role for ICANN staff – i.e. my group, the SSR group – or for parts of the ICANN community to actually have an actual role in this? That's an interesting question. I don't know the answer.

All of this is going to need to be written up more succinctly and with more data points so that we can have a look at this. We also have to decide whether or not this is a high likelihood and a high damage, if you like, kind of risk. We also have that kind of work to do here.

What do we want to do going forward? This is as much a question to you in the community as it is any kind of statement. We're thinking the first piece is easy. We'll take these 23 risks and we'll take a first glance at each one and write this up so that the community members that are interested can read them and help us with those. We'll hopefully do

that fairly soon. That'll be about documenting the assets, identifying any existing mitigations that we think are out there, and also identifying if we think there's other things we can do.

There's a lot of people in this community and specifically in this room that have expertise that, once we've gone through the first drafting, might want to help some of this. You might not. You might be too busy with your day jobs.

There's a question there that I have. Considering that we don't own all the assets, we don't own all the mitigation, is there a way that we can work with you in the community who are interested to actually start first documenting some of these issues more clearly, deciding which ones we really think are priorities?

Not all 23 of these risks that we found here are going to be the high priority ones. We can't deal with all 23 of them in one go. Maybe there're risks we've completely missed, and that would be a great thing to tackle as well.

How do we want to involve this expertise or how can we involve this expertise in the discussion? We have a small group of about six people, and we have a thousand other things to do as well. We're definitely going to put a lot of resources into this, but what I don't want is to produce something and just throw it to the community and say, "Look. It's done."

We need you to actually help us with this. We need your help. Make sure that we as a community have the right documentation of this, the right understanding. Frankly, we're going to need some help in actually

if we decide that some of these are ours to try to mitigate through indirect action and education, etc., to help us with that. (I don't care about that.)

That's where I came to. That's the question I came to this group with. Is this a good way to go forward? Are you comfortable with the staff taking the first lead on some of this? Have we got it all completely wrong? Some of you are risk experts, right? If we've got it all completely wrong, let us know. But then also let us know how to do it right.

I'd like to have a discussion about that. Do we have any questions in the audience or anybody who has any comments? Are we completely insane, or are we going down the right path? Is anybody awake? I see stunned silence. We're going to get to go for coffee early or to the bar. I'm not going to go to any of the other sessions because half of them have gone already.

That was basically the message. This is an update. I know some of you are going to be interested in some of this, so I'm going to hound you. Even if you don't get up to the microphone and say you're willing to help, I am going to come and find you and hound you and ask you for some help in this area. Somebody behind Steve. There we go.

[MATT STAFFBURG]: [Matt Staffburg], .se. What is the time plan?

JOHN CRAIN: We went through the 23 risks. If we're just going to do this on our own, it's different depending on which way we decide to do it. If we just go

with this, I wanted to get the first – we had three risks at the top, one of which is actually a combination of various risks that seem to be correlated. I want to get the white papers on those out before Marrakech.

The question, then, is if we can get those out before Marrakech, is there room to do some kind of working session? Should we maybe come together to tackle one in more depth and actually have a discussion about whether we got it right and whether there are things we missed, or do we just keep publishing the next?

I'm thinking I actually want to have discussions with people about how we're actually going to turn these risks into mitigations or not. I'm perfectly okay with the answer to some of these being, "Well, there's nothing ICANN should do here." I want to make that decision. Does that answer the question?

UNIDENTIFIED MALE: I have two questions. One is: are you going to focus on any technical issues like, for example, DDoS? There are some solutions out there, but how are you going to look at it? Are you going to [inaudible] document them as a risk assessment framework is what you are looking at? That's my first question.

JOHN CRAIN: Being as we've not done it yet, it's hard to say. I'm a technical person. I have to say there's a risk that I will focus too much there. Then when we publish the document, you will turn around and say, "Well you missed

these other implications." I think that's okay if that gets the conversation going.

The word technical is a little strange. It's never really clearly defined. I think most of these risks are very technical, some of them are not. I will be talking to people like Jacks to help me go through these. These guys have more risk experience than I do. I'm just sort of a geek. I don't know if you want to speak to that or are you good? You looked like you were looking.

JACKS KHAWAJA: For any of the risks that are attacked in this exercise, our goal is to really ensure that the risks are reduced to a level which we're satisfied with. Ultimately, we can never completely remove a risk. It's very difficult to do. In some cases, you can get insurance to insure against it – not really as a DNS – but you can get insurance to insure against it or in some cases you can avoid the risk entirely.

For instance, putting an office in the middle of a hostile location. You can put it there and put security controls around it like guards and so on, or you just don't put the location there. It's very difficult to do for the DNS. We really need to be practical about it and try to put the right controls in place and work with the right parties. I think that's the key is working with the right parties.

UNIDENTIFIED MALE: My second question was: how are you expecting that we can help you? Is there any discussion group or mailing list or something? That's my question.

JOHN CRAIN: If I understand, how can you get involved and interact?

UNIDENTIFIED MALE: Exactly.

JOHN CRAIN: We can do this multiple ways, and I guess the question goes back to the community as well. We could form a mailing list, done that a million times. We all need more mail, right? Yeah, exactly. We could form a working party. We could let staff go off and write the first things, and then we could come back and do a workshop. We can have a completely separate meeting somewhere where we go off and bash these out with whiteboards and papers.

Being an engineering type, I tend to be more favorable about doing some kind of workshop that's well documented if people are interested to do that than I am to have yet another mailing list. These things, of course, you can do multiple of. I don't know if we could get a sense in the room of whether people would want to have a mailing list to work on this or if they'd just prefer that we go away, do the work, and come back when there's something to work on.

One of the problems at the moment is that we're right at the beginning. There's going to be a lot of [graft] to get it to a stage where we can start thinking about things that we can operationalize and mitigate. I'm open to either.

I prefer to go down the workshop route if people are willing to put in the time. It costs people time and money to do these things. It's kind of one of the questions I have to the audience. Do you want to put some effort into this? We can go off and write all this stuff, and then it'll be just another document from ICANN. People who know me know that I'm much more of a fan of working with the community and having the community do all the work so I can sit at home and drink beer. Oh, no, that wasn't the reason.

Does that answer your question? We have a question online? I can't hear you.

UNIDENTIFIED MALE: The mic was not on.

JOHN CRAIN: You got it? Can we get the microphone on, please? Here you go. Just take this.

CARLOS ALVAREZ: Thanks, Steve. This is Carlos, ICANN SSR staff reading a comment submitted from a remote participant, Alejandro Pisanty. "Creating a collaboration is extremely important. Involving CERTs [inaudible], APWG (already involved), some ISOC and other professional society chapters is feasible. Another important point for this effort is to get closer to communities where risk can be mitigated in languages other than English. I am teaching about this stuff. I have already done it twice this semester, so happy to contribute. Would be glad to help set up

workshops and believe that a more active approach is best. John has it right."

JOHN CRAIN: Alejandro, thank you for that. I've actually already started talking to people in the CERT world a little bit about this as well. I agree there's other areas of expertise apart from these in this room. I'm glad you support the collaborative way of doing this because, obviously, as I've just said I prefer that as well. Mark?

MARK [UNKNOWN]: If you'll go to your slide number ten where you divided up into the parts, just the overview there – back, this one – it's quite obvious that the gray part is so much larger than the other parts. The blue part is very small, and the yellow one is a little bit larger. The grey one you should have made it very large. From that, it's quite obvious that involving the community increases the chance of also covering parts in the gray part.

JOHN CRAIN: Absolutely, 100% agree. My PowerPoint skills are atrocious, so I didn't know how to shrink the box. Well, actually I did, but yes, absolutely. I think the stuff in the blue area, once we've defined the risks more correctly and we've identified those assets, are going to be fairly straightforward. I think many of these risks, not beating our own drum here, but we've done a lot of work in these areas already for our internal stuff.

There's risk in there about IPv6 transition and risks of the complexities of dual-stack systems. We've done all that. Deploying Anycast into our root server, we've done a lot of that work. That stuff's the easy stuff because we have control. The stuff as you go down gets harder and harder, I agree. I think, and this is why I'm coming here and saying community. If it's not a community effort, we're not going to succeed.

Any more questions before we go to the bar? We good? Okay. I am going to chase you all, at least the people I know already. Thank you for coming and listening to us. I'm going to start working on these documents. I've already started working on a couple of these documents, by the way. I'm going to start going forward with these documents with our staff.

I will probably be reaching out to some of you for sanity checks. I will discuss with our meetings people and the people I see in this room that I know about the possibility of maybe doing a workshop in Marrakech unless people tell me that there's a better location. Maybe it's not an ICANN meeting. Maybe it's the first meeting. I'm open to all of that. I will reach out to people, and we'll figure this out. We don't have to do it in this room. With that, you've just got some of your life back. Thank you very much.

[END OF TRANSRIPTION]