ICANN
Transcript
DNS Security and Stability Analysis Working Group (DSSA WG)
13 September 2012 at 13:00 UTC

Note: The following is the output of transcribing from an audio recording of the DNS Security and Stability Analysis Working Group (DSSA WG) teleconference on 13 September 2012 at 13:00 UTC. Although the transcription is largely accurate, in some cases it is incomplete or inaccurate due to inaudible passages or transcription errors. It is posted as an aid to understanding the proceedings at the meeting, but should not be treated as an authoritative record. The audio is also available at: http://audio.icann.org/gnso/gnso-dssa-20120913-en.mp3
on page
http://gnso.icann.org/en/calendar/#sep

Attendees on the call:

At Large Members
Olivier Crépin-Leblond (ALAC) (co-chair)
Cheryl Langdon-Orr (ALAC)
Andre Thompson (At-Large)

ccNSO Members
Takayasu Matsuura. Jp
Rick Koeller, .ca (CIRA)

GNSO Members
Mikey O'Connor - (CBUC) (co-chair)
Don Blumenthal - – (RySG)
George Asare-Sakyi – (NCSG)
Keith Drazek - RySG

SSAC members:
Jim Galvin (SSAC Vice Chair)
Warren Kumari

ICANN Staff:
Julie Hedlund
Bart Boswinkel

Nathalie Peregrine

Apologies:
Katrina Sataki.lv
Jacques Latour.ca
Julie Hammer (ALAC)
Mark Kosters (ARIN)
Joerg Schweiger (cCNSO co-chair)
Scott Algeier

Coordinator: All recordings are now started. This conference is being recorded. If you have any objections you may disconnect. Please go ahead.

Nathalie Peregrine: Thank you very much, (Louise). Good morning, good afternoon, good evening. This is the DSSA call on the 13th of September, 2012. On the call today we have Mikey O'Connor, Cheryl Langdon-Orr, Olivier Crépin-LeBlond, George Asare Sakyi, Andre Thompson, Rick Koeller, Don Blumenthal, Takayasu Matsuura, Keith Drasek, Warren Kumari and Jim Galvin.

We have apologies from Julie Hammer, Jacques Latour, Jörg Schweiger, Mark Kosters, Katrina Sataki and Scott Algeier. And from staff we have Julie Hedlund, Bart Boswinkel and myself, Nathalie Peregrine.

I'd like to remind all participants to please state their names before speaking for transcription purposes. Thank you and very much and over to you, Mikey.

Mikey O'Connor: Thanks a million, Nathalie. As always perfect; everything's perfect. And welcome, everybody, to the call. Pretty short agenda. Going to keep working on this spreadsheet today. So we'll just take a pause and you can either chime in with changes to the agenda or if you have an update to your statement of interest you can speak now. All righty.

What you see on your screen is version 6 of that spreadsheet. I think last week were on version 2 so I've been sort of tinkering with it. And just to sort of sketch out the hour in advance here's what I thought we could do. The

ISPs have now gone through this once on a call. And so we've beaten up this set of answers a little bit more than my first draft.

And we learned a lot just by going through those answers. And so I thought that what I would do is take you through these in a little bit of detail partly because it dawned on me - I can't remember whether it was the last call where Mark was, at the beginning of the call, a bit concerned that what we were going to be doing is revealing or at least opening up the possibility of revealing information about the security arrangements of Registries and root server operators.

And I think by the end he was a bit calmer about that because this is not getting at what people do; it's whether they do it. And so it's pretty general stuff. But I kind of want to take everybody through this again to sort of refine our understanding.

And my thinking is - and I'd like your reactions to this idea - is that this week we would sort of learn it and then next week we would each take a copy of it and we'd work through it together and try and just fill out from our own perspective our answers to these questions just to sort of test this method.

And so, you know, that's sort of my goal is to get to the point at the end of the call where we sort of go yeah, this is something that next week we could kind of go through together on the call and at least take a first cut from each of our perspectives - the results.

I mean, you know, this is - we're going to have to surround this with loud draft, you know, capital letters, don't panic anybody; this isn't policy, this is just information gathering. We'll have to do all the right caveats but that's sort of the direction I'm going.

Rick, the ISPs that took a look at this with me were from Latin America, Japan and Europe. These weren't ISPs yet; these are members of the ISPCP so these tend to be policy people who work for ISPs rather than technical people.

We had a pretty lively conversation yesterday about this and moved the ball forward. We're going to circle back to the whole ISPCP in the GNSO and take a look at it together again.

And then I think once the DSSA sort of irons out how we actually gather this data and whether this is the right instrument to do it then at least the ISPCP is going to go back to their respective membership and run this by them. It's pretty iterative because we're, you know, kind of blazing new territory here. But we had a really good call and so I think we're going to learn a lot by doing this.

So with that let me start - unless anybody just thinks this is totally nuts and I'll just dive right in. Remember we, at this point, have about 11 big chunks of activity that we've identified in that circle diagram. And this is just a table of those chunks trying to identify what in this case ISPs do.

But, you know, you should be thinking about this from your own perspective whether it's as a representative of, you know, an organization or a constituency it doesn't really matter. It's just a way to gather thoughts about this at this point.

And clearly the goal is to try and find out if there are overlaps, which hopefully there are actually, and then even more importantly if there are some gaps in this so that then we can carry those gaps forward maybe towards Toronto and have a conversation about that.

So what we, ISPs, are starting to evolve towards is that we have sort of two flavors. We have regular ISPs and then we have what we call connectivity providers which are often telephone companies or very large ISPs.

And we may have to fill out two of these sheets; right now it's combined. Because you'll see that if you look in Row 16 down here in this cell here we're starting to hedge and sort of say well in terms of steering the research and analysis with regard to the DNS, i.e. the DNS that we're working on in the DSSA, ISPs are probably happy to leave the steering of that to others.

However where we have local Internet Registry and telephone number mapping kinds of functions we may be behaving more like part of the DNS. And we haven't really gotten our analysis all done on that. But we do have a very - a somewhat more DNS-like activity in that area. And it's basically what sits underneath the Dot(ARON) or the DotArpa portion of the DNS. And so we're still chewing on that.

But with, you know, we are really interested. You know, we are - we do a lot of this ourselves but we also are very interested in the research and analysis of the DNS because it impacts our customers and it impacts our upstream and downstream partners. And we do a lot of the sharing kind of stuff right now and we'd like to do more if we can figure out how.

So we - that's sort of our current view. And, you know, I'm now starting to use the "we" but it's really pretty preliminary yet. I wouldn't take this and, you know, I wouldn't take this to the bank.

Now the next two, standards, tools and techniques is the one, you know, we're looking at now. Again we participate in that but we also have a somewhat different role when it comes to this stuff under DotArpa and the e-num - the telephone number mapping stuff. And we are pretty dependent on the standards, tools and techniques really working right.

Same sort of thing except in this particular case we sort of bumped ourselves up a notch and said that we probably would want to participate in the steering. We don't want to lead it because it's, you know, we don't feel that it's appropriate for us to be leading the standards, tools and techniques development when it comes to the DNS.

Now if you step back and talk about DNS in a more general sense than I think ISPs would probably be even more engaged. But when it comes to the DNS in the narrow sense that we've described it here in the DSSA we just want to be there; we don't necessary want to be the leaders.

And then finally in the last sort of knowledge-sharing part of our functions we actually do a lot of this. We train people in how the DNS works both within our organizations and for our end users. You know, we also train them in the lower case DNS more broadly.

But when it comes to the root servers and DotArpa and all that kind of stuff we teach people how to do that. And so we gave ourselves the sort of highest level of participation here. And said, you know, look, we do this kind of stuff.

Another thing that emerged is that we do a lot of training about the DNS in our respective communities whether that be a country or a region or a city. And we don't just train other ISPs; we often get involved with training governments and other companies so we're pretty active in the education part of this.

And so - and again, you know, we do a lot of, you know, we participate a lot in all sorts of (glue) layer things ranging from the ISPCP all the way to the IETF and the Internet Society and on and on, (ARON), you know, etcetera.

But again we don't feel like we need to lead the effort to get that stuff put together as long as it's good with our standard caveat that when it comes to (LIRR) and telephone number mapping we probably need to be more active.

And we are probably going to go off and assemble a gaggle of LIR telephone number mapping type members of the ISPCP and ask them about how active because none of us on the call were real comfortable making this up for them.

You know, the ISPs I represent don't do much of this except, you know, to the extent that they're LIRs but they're relatively small in the scope of what they do; they're not like Comcast or AT&T.

Anyway that's sort of the first chunk of the work. Then the more technical operational part we drop back a bit again with our usual caveat about the LIR and telephone number stuff.

But we - and so what we started saying is look, we do a lot of DNS stuff, lower case DNS, we do hardware, software, vulnerability, DNSSEC, all sorts of things. But we don't do it for the DNS; we don't generally do it for the top level that the DSSA is looking at. So we put ourselves in the consume category here rather than the participate or do. And the same went for the upstream and downstream kind of stuff. We sort of backed off on this and again except for LIR.

But when it comes to the (glue) layer we're really interested in that because that's a place where we can share our experiences with the lower case DNS with those who are doing the upper case DNS and learn from each other. And again we kept ourselves at a pretty low level on this. We want to help steer but we sort of backed off from actually participating on this again with our LIR stuff is the exception.

Pretty similar reaction on the operational and technical practices. You know, if you sort of take these two together this is sort of the hard - the first one is sort of the hardware and software and tools and the second part is more focused on the human organizational side of the thing. They are very intertwined. And so our reaction was very similar that we participate at sort of the same level in both of these.

And then we get to the last half of this sort of frontline stuff. And here we feel like we're right at the front lines. You know, if the DNS - there are certain kinds of SSR events that take place in the DNS that we're right on the front lines of. Often our customers call us first. So, you know, we kind of need to know where we can go for really good information about what's going on.

And on the flip side it's often our customers that are the source of the attacks. And so we felt that even at the DNS level we do stuff when it comes to incident response and event monitoring.

Again in terms of the DNS we don't need to lead the (glue) or the core but we want to participate in both because, you know, this is so close to the core of what we do. And so we'd like to participate in the sharing, steering part of this. But again we don't feel like we're the people that need to lead the charge.

And then the last part, which is sort of the managerial, you know, if you think of audits, (cobit), all that stuff, this is sort of the risk assessment stuff. And this is, you know, where the DSSA plays.

We do a lot of the same things, again, internally. We have all sorts of similar internally defined, internal DSSR controls. But we don't do that for the DNS. We're consumers of that, you know, we're hopeful that those things are done very well and that they yield great results but we don't do this.

And again our - this is why we're thinking we might fill out another sheet. We need the - we stuck this caveat about the Local Internet Registry and phone number mapping stuff because we may actually be part of the DNS in that role.

And that's also a conversation that we might want to have with this group because, you know, it's too bad Mark's not on the call because some of you folks who are more familiar with (ARON) maybe it will help us figure out whether we're part of the DNS on that or not. But that's a topic for another day.

So this profile is pretty consistent through all of these. You know, we do very similar things so we certainly cheerfully participate in the sharing layer so that we can trade our best practices and learn from others. But the rest of it we feel like we're consumers.

And I think that that's pretty similar for risk assessment of the DNS. Again, you know, this is the DSSA job in a way and now the newly emerging DNSRMF project is going to inform sort of how some of this unfolds. ISPs are happy to participate and certainly consume it but we don't lead any of this.

Same goes for risk planning, you know, the - just to remind you a risk assessment is okay well what are the risks and risk planning is well what are we going to do about those risks. And again we do this for ourselves and we may have a fair amount to contribute to the conversation. But we don't do it for the DNS.

And finally same thing in terms of managing all this stuff in an organization exactly the same profile. So that's kind of our story at least the second draft. I'm pretty sure that this is going to go through several more iterations.

But I thought it might be useful for the rest of you to hear that story just to sort of get the context of what we're trying to do with this data gathering thing. And I guess at this point I kind of want to switch back to conversation mode and ask the question how comfortable would you all be individually if we were to, you know, I mean, I've now added a tab to this spreadsheet that's exactly the same as the one that we just went through but it's all blanks.

And my thought was that maybe what we could do is think about it for the week, maybe even take some notes and then next week just sort of step through this not comparing answers, not - I'm not - I think if we were to say, well okay let's all tell each other which answer we gave that would get pretty slow.

But if we went through it and asked questions and said well what does this mean again? I could tell the story again and we could hammer on it, write some notes and use this as a way to essentially simultaneously or in parallel do that desk research that we were talking about doing. Again not as final but just as a first draft at - sort of collect a lot of first tries at this stuff and see how it goes.

Does that seem like a reasonable approach? And I'll stop. I've talked straight for a half an hour. That either means yeah that'll work fine or total shock and despair. Rick, go ahead.

Rick Koeller: So what I'd like to - sorry, Rick Koeller for the transcript. What I'd like to understand, Mike, is so from a ccTLD perspective would you - you met with a committee of ISP representatives. Would you look for a group of the ccTLDs to - and/or other TLDs to provide a collaborative answer or individual organizations to provide their answers?

Mikey O'Connor: Here's sort of what I imagine is that we as individual representatives of either our own organization or as a group of organizations like ccTLDs, we as individual people on the next call, will do what I did about a week ago and just make up a first draft.

And, you know, make it really clear that it's super-duper draft; it's not official. It's just Mikey's first try or Rick's first try. That then, you know, let's say that you and Jörg and - I can't think of who else is a - who else is a ccTLD?

Rick Koeller: Louis...

Mikey O'Connor: Yeah and Louis, you all do a - on your own. And, you know, save those files. And then either trade them amongst yourselves or trade them with us or both and use that as a starting point for a conversation that eventually sort of ripples out first within the ccNSO and then maybe out even beyond that to talking to other ccNSOs - or other ccTLDs.

You know, my thought is that this is the beginning of a conversation that probably never really has an end but we can sort of take snapshots of it at various points in time and use this as a way to trigger other conversations in the community.

But, you know, I'm sort of making this up as I go. You know, I think that one of the key points that's emerged for me is that at least at this stage this isn't policy; this is just good guesses. This is just sort of we who have been asked to represent a bunch of different organizations writing down our initial thoughts and then using that as a way to go back for verification and further refinement of both our answers and the model itself.

Because, you know, people may say well that business about education training and awareness that's not right; that's - there's really two pieces to that. And so it's research, you know, it's research on a conversation but where it ends up I'm not exactly sure except to know that if we had half a dozen answers from different perspectives on this that we could array on that circle diagram we would be well on our way towards being able to start to answer the question do we have any gaps or overlaps.

Because ultimately that's the - that's the goal of this section of our work is to see if we've got any gaps and if so highlight those. And then, you know, if we find gaps that's the place that I think we really want to come back and ask people well do you, you know, let's say that, you know, we have the ISPs here and this set of answers and we're all totally relaxed about steering - well that's a really - well we'll leave it there because it's a ludicrous example so it's a good one in a way.

So, you know, should ISPs be in the core of that function. You know, and so the group comes back to me and says, Mikey, why don't you go back to the ISPs and ask them. And I go, okay, and I do.

You know, I'm seeing this as the beginning of a conversation about who does what that's fairly open-ended. And I guess what it does is it sort of takes me back to the roots of why were formed in the first place is that by having a conversation this way we avoid the issue that sort of started the DNS where the decisions about who does what came from sort of the top down.

You know, if we have a conversation in this kind of context we might collectively say, you know, the ISPs really ought to be in the core of this. And, you know, the ISPs might even say I'm not sure that's right. And the rest of you might say no, you go back to those ISPs because here's why we think it would be a good idea to have them in there.

And by doing that change the dynamic of the conversation about who does what. That's sort of the endgame that I think we're aiming at with this. There's a long answer. Did that get me close, Rick?

Rick Koeller: Yeah, it did. Thanks very much.

Mikey O'Connor: Cool. Anybody else? Again I'm really looking for some feedback here because as you have experienced I can get out in front of you and I don't want to get so far out in front of you that I've lost you. And, you know, it may be easier...

((Crosstalk))

Mikey O'Connor: Go ahead.

Rick Koeller: This is Rick Koeller once again. I'm just looking through our Phase 1 report and I'm trying to find if we've sort of defined what we mean by edge, (glue) and core so that as we go through this we've got some level of definition around each of those...

((Crosstalk))

Mikey O'Connor: I’m going to - I should have done this before I started the call, I apologize. Let me open up what is probably the best definition. Oh and it's - I've got my screen sharing screwed up. Hold on a minute. I have to shut it all off to start it up again. Let me see if this works now that I've got my theory.

I've been having a very hard time getting this - the screen sharing to work correctly on this new computer. I think I figured out why.

And, Rick, I'll post a link to this on the - this particular file is sitting out on our wiki and right after the call send a link to it unless somebody can beat me to it and post it into the chat. It's in the - oh wait, maybe I can do this. Hold on. Yeah, I can do this. I think I can do this. Here's a link to it that you can use to actually get a copy of it.

This is about as close as we get to definitions right now. And it's all the little headings underneath those. And I think that that's part of this iterative process. That's part of what was helpful about having a conversation with the ISPs about those cells in the spreadsheet is that we started evolving the definitions as we had the conversation. And I think that we, the DSSA, will probably do that as well.

But this is sort of the latest iteration of a definition. And I think eventually what we need is paragraphs for each of those to make it clear what's going on. So there, that's my best try. I can make the words a little bigger. There.

Anything else that people think? How comfortable are you with the idea of thinking about this for a week and then we'll do one of these together - each of us produce one on next week's call.

And by the way I think it'd be just fine if Bart and Julie, you and - you two and Patrick, wanted to do one as well. I don't think that there can be too many of these because it's still so fluid; it's still so much in the learning stage that I think we all learn from each other on this.

But, you know, if people are comfortable with that what I'll do is sort of write that idea up and post it to the list and sort of invite people to think about it for a week and then we'll really hit it hard on next week's call. I think it's awfully close to the top of the hour to really get too far into it. But, you know, if people wanted to start now that's fine with me so either way.

I'm worried. I've got this prickly feeling in the middle of my back that I've left you all bewildered and dismayed. So some reassurance would be in order or not. Gulp.

Don Blumenthal: It's Don. Maybe a little bit bewildered but not dismayed.

Mikey O'Connor: Anything I can do to reduce the bewilderment quotient?

((Crosstalk))

Don Blumenthal: No, no it's just trying to step back and see how we might be able to get something done from the Registry side so...

Mikey O'Connor: Yeah...

Don Blumenthal: ...we understand what's going on I think.

Mikey O'Connor: You know, I wouldn't try to do it from anything except sort of, you know, Don and Keith, you're both from the Registry Stakeholder Group. I think for next week's call the goal could be just do them from your own individual perspective. Take your best guess and then, you know, up in the status just write super-duper draft like I've been doing.

Because just that guess is going to be a terrific starting point for what could turn out to be some pretty lively discussions. But I wouldn't try to go to the Registries Stakeholder Group yet...

Don Blumenthal: Oh no, no I didn't mean that...

Mikey O'Connor: ...until - yeah, okay.

((Crosstalk))

Don Blumenthal: I'm sorry.

Mikey O'Connor: Yeah, because...

Don Blumenthal: No I meant just to try to pull a couple of Registry people that's all; not the whole group.

Mikey O'Connor: Yeah, yeah. And again, you know, you may want to do one first and then poll them with your first guess. But I totally leave that up to you, don't want to get that prescriptive about it. My goal is just to have a pretty good pile of these things in a few weeks because I think it's having a good first try pile that we can use to start looking for those gaps and overlaps.

Okay...

((Crosstalk))

Mikey O'Connor: ...last chance. Other than that I think I'll kind of end it here unless there are other questions. And I'll go write a little note to the list and get it out before the top of the hour. No, I'm not hearing anybody else. Okay, Nathalie, I think we can call this one done. You can signal to end the recording and we'll call it a day and we'll see you next week. Bring your thinking caps and we'll punch a bunch of these out. Thanks all.

((Crosstalk))

Nathalie Peregrine: Thank you very much, (Louise). You may now stop the recordings. Thank you.

Coordinator: Thank you for participating in today's conference call. You may disconnect.

Mikey O'Connor: Thanks, Nathalie. I'm, as always, ever thankful.

END