

Transcript
DNS Security and Stability Analysis Working Group (DSSA WG)
16 February 2012 at 14:00 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 16 February 2012 at 14:00 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-20120216-en.mp3>

Presentation will be posted shortly on:

<http://gnso.icann.org/calendar/#feb>

Attendees on the call:

At Large Members

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

ccNSO Members

- . Takayasu Matsuura, .jp
- . Katrina Sasaki, .lv
- . Jörg Schweiger, .de (co-chair)

NRO Members

- .Mark Kosters (ARIN); (co-chair)

GNSO Members

- . Mikey O'Connor - (CBUC) (co-chair)
- . Greg Aaron - (RySG)
- .Scott McCormick (IPC)
- .Don Blumenthal – (RySG)
- .Rosella Mattioli (NCSG)

SSAC Members

- . Jim Galvin (SSAC)

Experts

- . Scott Algeier

ICANN Staff:
Julie Hedlund
Bart Boswinkel
Patrick Jones
Nathalie Peregrine

Apology: none

Coordinator: Thank you and welcome. I'd like to remind all participants today's conference is being recorded. If you have any objections you may disconnect at this time. You may begin.

Woman: Excuse me. (Scott Algire) joined.

Nathalie Peregrine: Thank you very much. Good morning. Good afternoon. Good evening. This is the (unintelligible) call on the 16th of February 2012. On the call today we have Mike O'Connor, Olivier Crepin-LeBond, Scott McCormick, Andre Thompson, James Galvin, Cheryl Langdon-Orr, Don Blumenthal, Greg Aaron, (Scott Algire), Takayasu Matsuura, Joerg Schweiger and Katrina Sataki. Oh and Mark Kusters has joined the call. From staff we have Julie Hedlund, Bart Boswinkel, Patrick Jones and myself Nathalie Peregrine. We have no policy for today's call. I would like remind you all to please state your names before speaking for transcription purposes. Thank you very much and over to you.

Mikey O'Connor: Thanks Nathalie. As always a fabulous job getting this all going. Just take that moment for people to indicate that there's a change to their statement of interest. All right. You will see in the agenda that we have a radical departure and Greg I apologize. I meant to write you a note to let you know that you're off the hook today on the DNS sec conversation that you probably came to the call thinking you were going to lead us through. But save that thought. It's not that we're going to depart forever but we had a pretty lively conversation on the Ops Group call on Monday and thought that we would take a moment to essentially do a course correction.

And so we're going to spend the first part of the call talking about what we talked about and get your sense as to whether or not we're on the right track. And then if we are we're going to dive into the methodology. What we've in a nutshell concluded is that we need to step back. And so the first part of this call right now what's on your screen I would like you members of the DSSA to meet the forest and forest meet the DSSA. Then DSSA I would like you to meet the trees and trees I'd like you to meet the DSSA. And essentially the initial message that I think we want to convey is that we've gotten ourselves a little bit lost in the trees and lost sight of the forest and I would to introduce you to the person that took you there. That would be me.

So what we want to do is take you through a conversation that we had on Monday to sort of explain where we're coming from. And then the conclusion that we arrived at just to give you a preview of that it's not that we're going to throw out the methodology or any piece of it but that we want to tailor it a bit and we want to start at a high level and then work our way down.

So here's what went on basically on Monday. It's that we hit that slide of the slide tech that we're using to update people in Costa Rica and said wait a minute our goal isn't necessarily to go at the super granular level that we're going. Our goal is to very quickly get to an answer to those major questions and really answer the, you know, ultimately the question that triggered this whole thing is the statement that came in god I can't even remember -- I think it was Brussels -- that implied that maybe that DNS is on the verge of collapse and, you know, urgent action is required. And with urgency like that we have to pay attention to the speed with which we're getting done our work.

One of the other things that we sort of zeroed in on is this notion that we really need to quickly focus in on the most likely and most dangerous wide impact threats. And the way we're going about it isn't - it's not that we're going to miss them. It's going to take us a very long time to get to that. And there is a pretty strong sense of urgency in the Ops Group that we need to pick up the pace a bit.

And so then we got on to a conversation about well what could we do to improve this situation and what you're in the middle of right now is the conclusion that we arrived at which is we need to take a step back on this call and take a look at where we are and where we're going and where this methodology is at and so on and so forth.

So what we talked about is -- let me just get all that on your screen -- is, you know, take a quick look at where we are and where we're going. And one of the things that emerged is that there are other groups that are in a way dependent on our work. They're certainly looking hard at what we're doing and that we may be doing them more of a favor if we can get through this fairly quickly because they may have more resources and more time to expand and refine some of this stuff and clearly many of us would probably be participating in that. And so it's not as though we are the only ones that are going to be looking at this but we are doing really useful work for them and they could probably use this work sooner rather than later.

So what we came to is to take a look at the methodology slide in that deck that's there and sort of ask the question how would we ask each of these questions? And I'll get to that in a minute but I'll give you a preview of it just so you know what I'm talking about. This is the series of questions that we're talking about. What I did in this slide is I summarized the methodology pretty much on one page and then started asking this series of questions on the side. What we want to do is take a look at this and get your opinions and thoughts about how we should ask those questions.

And I've added something in to this which is I think that these questions are being asked in the wrong sequence. This is a thought that Scott provided us several calls ago and which I still am pretty convinced on. I'm also pretty convinced now having looked at the methodology that there are actually some errors in the methodology that we need to correct. But I think the big thing is that we need to stay at a higher level than the level we've been at.

This super granular analysis is just - it's going to take us too long to get to a useful result.

And so one of the things we got into and we will get into later in more detail on the call is that the range of impact question we're asking right now needs to be refined. This notion that we're splitting things in terms of major or minor TLDs isn't really the right approach. We really need to focus on who is affected before we go into the range of impact question. And again we've got some ideas about that that we'll go into when we get into the details. I just want to sort of replay the conversation that we had so you know the context of why we're doing what we're doing.

And so with that we also came up with a bit of a list of things not to do. One is this deep granular list of branches that we are working on right now. And then one that I tossed in after we had the call is this notion that we've combined a couple layers and in so doing we've created a lot of permutations for ourselves to get through in our analysis.

So that's the setup for this call. That's the reason that we're going to take a quick departure. And before I go into the actual tailoring of the methodology which we'll get to in a second I just want to take a moment and let other Ops folks chime in, let others of you ask questions just to clarify really where we're at going into the rest of this call. So I'll stop my soliloquy for a minute and let you talk. Anybody want to chime in at this point? Okay.

So let's take a look at the methodology at a really high level and I'm going to shrink this down just a little bit so we can fit it all in one screen. There it's all in one screen. And one of the observations that I want to make is that if you look on the bottom of the Adobe connect room you will see our two scales. You'll see threat sources and we evaluate the range of effects. And you'll see threat defense and there we're evaluating relevance. And I have concluded that those are the wrong scales to evaluate, that part of what is causing us

trouble is that what we really want to do and what we keep trying to do is evaluating the range of impact for the threat events not the source.

And this is what I mean by an error in the methodology. I went back into the methodology because I thought that maybe this was just a Mikey error which would actually cheer me up a little bit but it's not. It's embedded in the methodology this way. And I think that that's part of the reason why we're in so much trouble. And so, you know, I'm a methods guy. I actually helped develop the methodology that Anderson Consulting and now Accenture use and so I have a healthy disrespect for methodologies being perfectly right.

And I'm quite comfortable with the idea that when people build methodologies they make mistakes and those mistakes can cause people trouble and that the solution to that is to fix the methodology. That's really what we're starting today is to fix this methodology to work better for us.

And the first fix that I've already embedded into this is that just to show the way the methodology was originally written. The impact step is quite far down in the methodology the way it's originally written and we are struggling because we're up here. We really want to evaluate the range of impact for threat events when we're looking at threat events. And the methodology makes us wade. It makes us wade, you know, through quite a number of things. In fact if I reconstruct the methodology the way it's written the methodology starts with sources of threats first and then does the events themselves. And again I think we've been struggling because that's not the sequence that makes sense.

And so that's what struck me as I was going through the preparation for this call was that in addition to asking better questions I think we need to take a look at the order that we ask the questions in and that's what's on the table for today and maybe this will take us a little longer than one call. We'll see.

So anyway the proposition that I'd like to put on the table is that we move threat events first, that we just do threat events. We don't try and tie these two together; that we just look at the events. That we evaluate the range of impact and that because we're doing that we move the range of impact part of the conversation either right up as the next step or maybe even within this step so that we can really quickly focus on the threat events that have significant impacts for the internet as a whole and, you know, just to take us back to our goal which is here.

You know, we're really trying to find the actual level of frequency and severity of threats to the DNS. That's our goal. And so let's put that first. Let's not discard everything else but let's focus on that because if we can get through this that helps some of the downstream groups and it also helps focus the work that we're going to do.

I'm not convinced that threat sources are actually the next thing that we need to do. One of the things that I'd like to lobby you on is that we may want to do vulnerabilities next. And I didn't move that one because I wanted to have a conversation about that but that's a possibility. I'm not convinced that the predisposing conditions is that important for our work. We can take a moment to just zoom over. Predisposing conditions seems to be more aimed at an organization that is doing a more broadly defined project.

So this is a sample table of what these predisposing conditions are and I really think that the one that applies to us is that we're complying with technical standards, in our case, you know, the whole RFC environment, the whole DNS set of standards. So I'm not ready to throw this one out yet but I'm not that excited about it. Just that, you know, a lot of this is now personal evaluation not anything but that.

Clearly likelihood is something that we need to evaluate at some point because eventually we do need to get to the summary of all this which is risk. But at this point I'm not sure what sequence to do this in and that's part of

what I want to put on the table for us to talk about. So now I'm really going to stop talking and let you all react a bit to what I've been saying and offer ideas as to the sequence and the question that we're really trying to answer, you know, really just start to dig into this a bit and get ourselves comfortable that we're on the right track by stepping back from this bit, going a little higher in the level of detail, etcetera, etcetera. So please react to what I've been saying.

Man: (Unintelligible) Cheryl.

Mikey O'Connor: Go ahead.

Cheryl Langdon-Orr: Thanks Mikey. It's Cheryl Langdon-Orr for the transcript record. First I was a little nervous there for a while. You've calmed me down. Even I managed to put my hand up or I should say put a little red teacup in agreement at that shuffling some of the things about. I certainly think the opportunity of doing the subset that you proposed in the order you proposed is a good thing and I would be very supportive leaving the probabilities up in order of priority. I too am not going to lose much sleep over the predisposing conditions. I think it's one of those things that actually does have to be there. So we can kind of just leave it as a bit of peel or the vegetable we don't like and to the corner table.

I appreciate the fact that there is other components up or bits of the worlds that are relying on, waiting for, poised to watch or in some way, shape or form affected by the work here and I think shuffling allows us to give them a reasonably clear set of understanding to mull over and opportunity to perhaps make some course corrections and get started on work that is complimentary our field, to that that we're doing in our field.

So but the only hesitation I have is that we as a cross community workgroup need to be really clear that this work however has to be done in a totally well tested, authenticated, formulated, properly sanitized and then shared and

accessible manner. And that's a pretty rare beast if it starts getting carved up too much. So if you promise we're not going to end up walking out and trying to put some sort of Frankenstein on stage then yes. You've got more positives out of me than I thought you would. Thank you.

Mikey O'Connor: A giant sigh of relief. When I passed the Cheryl Langdon-Orr test I know I'm on the right path. I suddenly realized that I wasn't taking notes so Cheryl I just butchered your last point and feel free to expand on that. Does anybody else have some observations that they'd like to share?

Jim Galvin: Mikey this is Jim.

Mikey O'Connor: Go ahead Jim.

Jim Galvin: So this is Jim Galvin. I had a question for you. We talked a bit about those who are affected in user communities with the phrase we were using on Monday. Are you going to say a little more about that here before we move on?

Mikey O'Connor: Yes. For the rest of you I should explain that. We had a - the very lively call we had on Monday was really triggered by the notion that one of the things that's left out of the methodology is a question who and which users of the DNS are affected by a given threat event. And we started to come up with a taxonomy that said that the scale that we'd like to use would say that ten, the highest impact scale, is a threat event that affects the global user community, that everybody that uses the DNS would be affected by that. And then tails downward that an eight might be - and some of these are just be way of example rather than done because I was doing my notes from that call. I realized that we didn't really finish it.

But let me show you what I came up with so far in a new version of the scale. So the ten would be users of the global DNS. An eight would be users of a global TLD. So for example .com. A five might be users of a specialized

regional or country TLD. And then the three and the one I left as question marks. We need to fill out. And the point that I think Jim is trying to get me to make is that we really need to do this first and then once we've figured out who the user community is then we can do the range of impact to that user community.

So that range of impact really becomes at two step process rather than a one step process. And that it was this missing step that was causing us a lot of trouble on previous calls because we kept getting caught by the fact that, you know, people would say well it depends. It depends on who this threat event is happening to and how broad that is. And so this is our first try at least at coming up with another step that precedes the range of impact and then reworking the scales that we would use to do that.

Does that cover the bases for you Jim? Jim is not on the Adobe room so I have to sort of check back with you and see whether you're okay with that Jim.

Jim Galvin: Yes. Thank you Mikey. Actually just to add to your scale so the suggestion I guess at the time that I had made and the way that I was thinking about it -- so let's just check this out on the group here -- is when thinking about the community that's affected. A scale of one and three would refer to those who are part of the TLD that is affected or inside it. So a concrete example, you know, if it's something related to a country code it would be those who are inside the region as directly affected. A scale of five and eight as you go up increases the user community so five and eight would be a way of expressing the regional or those who are outside of the affected area but their use of it is affected.

So for a TLD anyone else on the internet who might be using anything that's inside that TLD and inside that domain name. And then ten of course would just be global internet users, you know, the entire internet and the effect that it has on the entire internet. That's it for me. Thanks.

Mikey O'Connor: Thanks Jim. I typed those who are part of the TLD are part of the region into the scale so that we recorded that. Because I knew that that's what we were trying to do and I was having a little trouble reconstructing that. Joerg go ahead.

Joerg Schweiger: Yes thanks Mikey. It's Joerg for the transcript. I think we might want to invest a little bit of time to better specify what we really mean by this (unintelligible). For example if you say under age users of global TLD I think you for a purpose did not say generic but probably mean and I think generics I think we do have to be very precise. For example a global TLD might very well be let's say .pro and I wonder if anybody would really feel that the range of impact would be severe for the DNS if .pro would be affected in any way. So we have to be a little bit more specific and precise giving explanations to the numbers. Thanks.

Mikey O'Connor: Thanks Joerg. I think that the way I would have answered that is to say that .pro is a specialized TLD so that it would be more in the five category. Because I agree with you we certainly wouldn't want to put .pro in the global TLD thing. How could we reword that so that it was a little bit more accurate and get to the sense? Because basically what we're talking about is .com and maybe .net. We had a little debate about that on the call. I mean one way we could do it is we could specify those two generics and.

Joerg Schweiger: We might get back -- Joerg for transpose again for a transcript -- to what we already have been using like major and minor for example.

Jim Galvin: So this is Jim, Mikey.

Mikey O'Connor: Go ahead Jim. Jim gets to break in because he's not on the Adobe room.

Jim Galvin: Sorry thank you.

Mikey O'Connor: Don't worry.

Jim Galvin: I actually - I specifically don't like the phrases major and minor and in fact Mikey I'm not real fond of the phrase global TLD. I mean in my mind a TLD is global. Its effect or its range or its usage might not be global. So .com becomes the concrete special case of something which really does affect the global internet if something bad happens there. And then you have a lot of smaller country codes and even some small GTLDs where if something were to happen there, you know, the effect would be much less dramatic. It would matter a great deal inside the country or matter a great deal to the users of a GTLD but for the global internet it would have much less effect.

So I really am very - I much prefer just calling a TLD a TLD and then we should be talking about the effected user communities when we want to judge the effect and what it means. Thank you.

Mikey O'Connor: So I think what I'm hearing is I think we're all pretty agreed that ten is the users of the global DNS. You know, maybe what we're getting in trouble on is that we are trying to specify two things in one scale. We're trying to specify the range of the users but we're also trying to specify the thing that they're using, a TLD or a DNS or whatever. And maybe what we need to do is just step back and say the 10 is global users. Anybody got any ideas on that?

I mean, this is the - I tell you what I'd like to do for a second is maybe spend another two or three minutes on this and then indicate that we're not done; that we need, you know, that this needs work. Throw it out to the list for some conversation; let people think about it for a bit. And then go back to this sort of very top level view now that we've clarified what we're talking about when we're talking about range of impacts.

And let me torture your eyeballs for just a minute. One of the things that the methodology does have, which might give us some hints, is they've got a range of effects - I found - this is one of the errors that I found in the

methodology is that this range of impact thing appears several times in the methodology. And the scales are almost the same. It's almost as though it wasn't - it wasn't edited well.

And so this scale is back to the sort of scale that we had for range of effects in the lower part of your screen. And instead of calling it cyber resources maybe what we do is we call it users of the DNS. And so if we substitute that wording what we wind up with the scale kind of like this one. Jim, I apologize but I'm back to a reworked version of the scale that came next.

What I did is I rewrote those to say sweeping involving almost all of the users of the DNS; extensive involving most of the users; wide ranging involving a significant portion; limited involving some of the users and minimal involving few if any of the users.

What that would do is it would remove the dimension of DNS TLD, you know, all of that stuff would tend to get removed from the discussion and thus wouldn't be, you know, it wouldn't present the problem of putting these two things together.

Would that side-step the issue that you're raising when we're talking about major/minor TLDs, Jim? You may be muted or you may be puzzling.

Jim Galvin: No, sorry. This is Jim. It just takes a moment for me to get off mute.

Mikey O'Connor: Yeah.

Jim Galvin: But, you know, I guess I'm okay with that for now because I really don't have a better suggestion. We struggled with this a lot on Monday when we were talking. You know, I mean, the problem that I always have is, as by way of example, is a small TLD is, you know, it's a crisis for even 1000 people who use it.

You know, and even though it's a minor effect on the Internet as a whole and, you know, you have to balance that against the dotCom which is a crisis for the global Internet as well as everybody inside it. And that's the distinction that I think we're struggling and trying to find a way to make.

So I'm okay with anything which is a little different than what we're doing and let's just try it and we'll see how it goes.

Mikey O'Connor: Well actually the interesting thing about - and, Cheryl, well a queue is building and so I'll stop and let the queue play out. Cheryl, go ahead.

Cheryl Langdon-Orr: Thanks, Mikey. Cheryl for the transcript record. I'm comforted very much where we're heading with these new cells because I think working with a consensus agreed nomenclature is very, very important and it needs to be a nomenclature that is therefore not only agreed upon by a current group which is large enough to be more broadly representative of people's ideas and concepts so we can narrow it down with Occam's Razor to the simplest terms.

But it also therefore gets to be repeatable in the future and in future analysis it would be incredibly useful for us to have got this right. The - or as good as it's going to perhaps, Jim.

I still think - my feeling at the moment is that the use of the sweeping, extensive, wide-ranging limited involving some more than limited involving fewer users is a good one and probably more useful of the two because otherwise we get into a value judgment on - where we have to the purpose of an intention of use.

And then we get into the areas of unknown and not well known intentions of use in TLDs. And that may or may not be size of TLD related. And so I think I'm happier along this way.

But I want to put in a little rider and that is that it may - if you're doing it this way, Mikey, may mean we now need to do another paired analysis which is not just range of impact but we also need to pair it with an analysis that not of relevance but of the severity which is in that taxonomy in that Table 4 you just put up which is the severe effect, the minimal effects, etcetera, etcetera.

So - and we might actually need to pair off with that - those things that we've actually managed to avoid to date. Okay. Thank you. Enough prattling from me.

Mikey O'Connor: I think that's lovely because it actually feeds right into - what the methodology does is it in fact splits the breadth of impact or the range. I had to come up with a new word when I realized that this was a repeated part of the methodology.

But they do call out severity. And here's the starting point table for that. Oh, God, this is going to be hard to display. Hang on a minute.

Cheryl Langdon-Orr: Yeah, Mikey, I think what I'm saying is if we redo - and I think I'm happy to redo the, you know, the extensiveness of how much of the DNS global or not is affected I think that that range of the nomenclature and that new taxonomy is going to work. But it's going to be more useful if it's worked paired with one of these more detailed, you know, exactly what degree of severity can we also put to it.

I mean, something that affects absolutely every user in a, you know, in a vastly - a vast TLD which has a lot of geopolitical or global reach but actually isn't a terribly bad thing. It affects it. It might be a security risk but it doesn't actually resolve in a security consequence just a risk of that consequence wouldn't be picked up otherwise.

So I think if we do the pairing right this is a good way to go. And I'll stop because I notice I'm blocking Jörg, sorry. Sorry, Jörg.

Mikey O'Connor: Jörg, go ahead.

Jörg Schweiger: Okay yeah, thanks. Jörg for the transcript. I'm not so sure whether I'm not I'm available - or I'm capable of picking up what Cheryl really said. And I'm still a little bit confused with major, minor and so forth.

I think we do have to differentiate between severity and that might be evaluated in its very specific scale. On the other hand on the range of impact why do we have to specify something like - it may be major or minor which - where there are a couple imposing that or a specific TLD I think that we are not even interested in whether it's Com or not or it's dotOrg or dotPro I think we are only interested in the sheer number.

So why don't we try to just give the number of users that are affected and work with certain boundaries to say if, for example, 100 million users would be impacted then this clearly is kind of a 10, just to give it a number, or a representation where we can easily interpret - interpret what we really do mean by 10, 8, 5 and so forth.

So just put it in boundaries of number of users that are affected and then on the second step we can say in which severity level they are affected by Table - what it is it - I can't read it, D5(a) or something, sorry.

Mikey O'Connor: I'm typing in Jim, for your blind participating. I'm actually typing percentage possibilities in. And we can refine these clearly. So let me just put question marks so that - because I kind of like the idea of enumerating this and essentially skipping the whole - clearly there's pretty strong enthusiasm for the idea of getting rid of the major/minor TLD notion and really focus in on the breadth of the impact.

And then pairing it with the severity so that if we have - as Cheryl said, you know, if we have an impact to everybody in the Internet but it's a minor

impact that's not a big deal as opposed to an impact to almost everybody in the Internet and it's a major impact; that's a much more important thing to look at.

So I think that what I'm hearing is that we're sort of homing in on only one range of impact scale not the two-step process that we came up with on Monday but really a one-step process that really focuses on the number of users that are affected and then paired with the severity of the impact to those users so that we've got, you know, both of those combined in the same analysis.

And it seems to me that that's pretty straightforward to do. You know, what we could do is we could...

((Crosstalk))

Mikey O'Connor: ...basically take all of this stuff and put it into the range of impact discussion for threat events. And have those two evaluations going on just like we've been doing, you know, two choices that we each need to make when we're doing our voting threat event by threat event.

And in fact we could probably do that in one call because if we just stay with threat events - and let me show you the list of what I'm talking about, Jim - we really at this point only have about 10 threat events. This is the major zone file, lesser zone file, root zone, IANA zone, DNS SEC list.

And it seems to me that we could come through that list in one call and make these choices that we're talking about. Now we may, having zeroed in on this, say well wait a minute there are some more threat events that we want to add to the list but even if we doubled the length of the list I think we could get through this part of the discussion very, very quickly because we're not creating the permutations of who's the threat source for the threat event we're simply describing the impact of the event. So...

Jim Galvin: Mikey, can I get in the queue?

Mikey O'Connor: ...I'm pretty perky about that. Jim, do you want to - I'm hearing somebody trying to get in; is that you, Jim?

Jim Galvin: Yeah, just wanted to get in the queue.

Mikey O'Connor: Go ahead; I'm done.

Jim Galvin: Okay. Here's the concern that I have with number - let me try a different approach to describe what's on my mind here. And it goes to, again, the question of major versus minor and even trying to quantify by number of affected users and using that as the severity of a problem.

For me the issue is still about context. It's still about which user community you're talking about because in some ways the size doesn't matter. And I guess I'll just - I feel like I'm saying this again and I apologize for that. But I'll do it with that preamble let me use this example again.

If I am a registrant in a small TLD even if there's only 100 of us who have registrations in that TLD it that TLD is down it's a crisis for me. I mean, it just is; it's a 10 event in that point of view.

On the other hand from the point of view of the Internet and the users at large now it just depends on how many people actually use that domain and that's a number that we may not actually be able to speak to. So the severity of it, you know, we can't even really accurately speak to.

So what bothers me about the vote-taking is it's the fact that we seem to be focused on size and we're trying different ways to quantify that size and I just - I'm uncomfortable with that. I feel like some of these issues can be handled

in text when we get to writing prose because we can talk about the different perspectives of looking at an event and describe it in that way.

You know, if you're inside the event it's a crisis; if you're a user of the domain inside the event, you know, it's a crisis. If you're outside of it maybe it means nothing or maybe it means something if you're a dotCom because it's a large number of users outside of the event that are affected.

So I don't know if that helps any but just a slightly different perspective on what's been bothering me about these votes that we take about how important something is. Thank you.

Mikey O'Connor: I wish that you were on the Adobe room, Jim, because what I was doing is as you were speaking I was toggling back and forth between the breadth scale and the severity scale because I think we all agree on the point that you're making which is we need both scales in order to really describe these impacts. We need the scale of how widely is this felt but we also need the scale of how severe is it for the people who are affected.

And I think that we cover that by doing the severity discussion and the breadth discussion at the same time. Because then what we can say is...

Cheryl Langdon-Orr: Cheryl here, Mikey.

Mikey O'Connor: ...the breadth of this impact is limited to a single TLD or a single region or, you know, a very small group of users but it's catastrophic for them and accommodate both of the points that you're trying to make.

Cheryl and then Jörg, go ahead.

Cheryl Langdon-Orr: Thanks - thanks, Mikey. Cheryl for the transcript record. Look, exactly, Jim, I think we're all in - well certainly I am and I'm seeing no chatter in the

space - in the Adobe room space to the contrary. I think we're all in strong agreement with that.

And I think the right pairing, which is what Mikey was just saying, will allow us to bring that out of the analysis and that avoids using numbers and even I'd tend to avoid using percentages, Mikey, just if that - if there's any influence at all.

The other thing I think that is good about the right pairing with the right nomenclature in this analysis is it allows the whole testing to be set up to be robust as a useful methodology for a right-now analysis with what we need to be thinking of as the legacy TLD world and allows it to be far more useful in the post-expanded TLD world which of course is, you know, the precipice upon which we are all standing.

If we get it right now it'll be far more useful for future and repeat analysis to be done. And it would be, for example, then I'll stop talking, a very tiny, narrow extremely limited in terms of number - raw number - of user, future new gTLD could be absolutely disastrous not only to the users in that TLD space if there was a serious failure.

But should that be a critical resource management space have end user affects which go beyond registrants and simple integers. So I think it's essential that we go down this path right. Thank you. And keep away from minor, majors and percentage type numbers if we can.

Mikey O'Connor: Jörg, I'm going to get to you for a minute but - in a minute - but I want to point out one more dimension that the methodology - the methodology actually has three scales that it uses to evaluate impact. It's got the breadth one that we've talked and the severity one but it's also got one that talks about the kind of impact.

And this is a list that we will need to refine a bit but let me just rattle them off for Jim. The scale of the type of impact is harms to operations, harms to assets, harms to individuals, harms to other organizations and harms - I've edited it to read the world.

And, you know, it may be that what we want to do is evaluate all three of these things at the same time to sort of capture the notion that Cheryl was getting in in that last moment.

Jörg, go ahead.

Jörg Schweiger: Okay, yeah, thanks Mikey. Jörg Schweiger for the transcript. I fear I'm not in sync with what Jim said or at least the way you summarized it, Mikey, because I think we are interested in numbers and nothing but numbers if we're talking about how many users will be affected.

And I do not see your point in the differentiation between whether the affected user is in a small or in a large community and if it's something terrible for the user that is being concerned or not.

Because we are looking at the DNS level and there's no such thing as a connection between a catastrophic fall down of a certain TLD and its impact for this TLD. Yeah, this is a major crash for this very TLD. But nevertheless it doesn't affect the - the DNS and that's what we are looking at.

So we're not talking about severity, I'm not talking about severity for this very specific user that is being affected but I'm talking about what is going to happen to the DNS - what are the severity and the impacts of the DNS.

And there it might be that on one hand, for example, we do have an impact on let's say 10 millions of users whereas the severity of this very impact is absolutely low. We see that something is happening to the DNS but, well, it doesn't mean that much to any of those 10 million users.

On the other hand we do have severity issues that might affect 10 million users and their - this affect would be disastrous because the DNS isn't running anymore overall. So in this case we do have the duality of both; a sheer number that is affected and that is - I'm thinking a word - distinctive on one hand and on the other hand we have an impact that is really catastrophic.

And I think we do have to look at those categories and not at how a certain user might feel - who is affected might feel if something happens to his TLD and so forth.

Mikey O'Connor: I'm thinking that we're violently agreeing with each other and we're not quite describing this right yet. And so I think what I'd like to do is park this for now. I think we're actually quite close to the right language. And I'll make it clear in the way that I leave these documents that we're not done.

Let's - and I'll try and publish enough of this to the list so that people don't have to wade through giant mine maps to find the scales that we're working on. And let's push this out to the list for some conversation. And really I think it's bad word-smithing at this point. I think that the destination we're headed is pretty much a place that we all agree on.

What I'd like to do for the - well I guess I'm not going to do anything else on this call because it's three minutes until the top of the hour. I was all set to go jump into another giant pile of work but we're not going to do that.

So let me take this out to the list. Let's think about it over the next week. Let's revisit this on the next call. And let's spend the last two minutes just checking in. Do people feel like we're headed in a good direction here? Is this - I think that we are but I'm also running the call so it's always hard to tell.

Just last chance to sort of throw up your hand and say you ops people were crazy, this is all wrong, etcetera. Anything we can do to change course, get back from the precipice that are - Greg, go ahead.

Greg Aaron: How long is our working group been in existence now?

Mikey O'Connor: We are on our 40th status report so I would count us at about 45 weeks.

Greg Aaron: Okay. So, you know, by Costa Rica we will have been in existence for almost a year.

Mikey O'Connor: Yeah.

Greg Aaron: And at this point we don't have anything down on paper. And I'm in general not in favor of working groups that have indefinite endpoints. And we need to start thinking about putting something down on paper and figuring out when this group is going to deliver a work product.

Mikey O'Connor: I think you share the sense of the ops group. I would push back a little bit on the fact that we have nothing down on paper because we actually have quite a bit of stuff strewn back in the...

Greg Aaron: Well, I mean, what I'm talking about is progress towards a deliverable work product. We ultimately - every working group delivers a paper. We've got our charts and everything but we have not started the hard work which is actually something that is always difficult to do which is putting those words down on paper. So I think we need to get serious about a schedule.

Mikey O'Connor: I don't think you'll get any disagreement from the ops gang. Cheryl, go ahead.

Cheryl Langdon-Orr: Thanks, Mikey. Cheryl for the transcript record. Greg, a little bit of a push back on what you were saying and picking up on what Mikey said. Each of these analysis and - whatever the word is - the plural of analysis - that we've

done results in a scale, an aggregate, a number for the numbers people that we can quite quickly use as real data.

Real data, which is what is missing and the reason that we actually have this wonderful world of conjecture that we're trying to drag ourselves out of into actual measures. If it was easy it would have probably already been done. I think this is one of those rare occasions where it is inevitable that this is not going to be a short-term workgroup.

With, you know, its first report giving cold hard facts that people could all go off and start utilizing. I think this is definitely one of those workgroups where probably the toughest part is not actually word-smithing the report but in fact establishing the selection of methodologies and the refinement of the methodologies which is what we're sort of in the middle of now. I think that's all highly reportable.

But if the actual measures were easy to have picked out and just written a report about this group probably wouldn't even exist in the first place because it would have been material that we all had ready access to or reasonable access to at the right level.

I think the new tweaking of the direction will get us to that much wanted end point in a much more manageable (unintelligible) scale. But I for one am not overly concerned about the fact that we'll be looking at, yes, 12 months worth of work and probably meeting number 45 or 50 by the time we gather together and put out some more paper because I think this - the scale of this project deserves that. Thank you.

Mikey O'Connor: And with that we're two minutes after the hour and that's the final word.

Thanks, Cheryl. Thanks, all, for I think a very useful call. I'll push some stuff out to the list and we'll see you in a week. That's it for me.

Nathalie, you can...

Cheryl Langdon-Orr: Thanks everyone.

Mikey O'Connor: ...wrap us up.

((Crosstalk))

Nathalie Peregrine: Thank you, (Tonya), you may stop the recordings.

Coordinator: Thank you and thank you for joining today's conference. You may disconnect at this time.

Nathalie Peregrine: Thanks (Tonya).

END