
MONTREAL – ccNSO: TLD Ops Standing Committee
Sunday, November 3, 2019 – 10:30 to 12:00 EDT
ICANN66 | Montréal, Canada

JACQUES LATOUR: Welcome to our, I wish I knew which TLD-OPS number meeting we're at; 42? No, I don't think so. Welcome to our TLD-OPS Standing Committee meeting in Montreal. Just a second, I need to -- okay, so today's a big day. So the agenda today is, we'll go through a couple of action points that we have. We added DAAR update with John, a discussion. Still here? Oh my god, I need glasses or something.

And then, most of our time we're going to spend talking on the disaster recovery business continuity workshop that we're going to have this afternoon. So, we'll finalize all the logistics then, make sure we have all of our ducks lined up on what the role of the standing committee members are, the volunteers that we have, and make sure we have a good session. And that's the agenda for today. Well, the focus is not super great. So, I think Régis and I in our -- Régis, our co-chair; vice-chair, co-chair?

RÉGIS MASSÉ: Vice-chair.

JACQUES LATOUR: Vice-chair, in all of our grandiose expertise I guess, we've put off the deadline for the action item to the next meeting, ICANN67. The only

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

action item we have active for this meeting is to finish the planning and execute the workshop for today. There was a new action item that was added that we need to talk about at the next ICANN meeting, which is, we need to review the charter for TLD-OPS. I think we are to do that on a regular basis, so we'll take that on at the next meeting.

So, I think overall, we added one or two members to the group, so we can do a status update on where the mailing list and the repo is at the next meeting. But overall, the mailing list is on track and everything is working, so that's where we're at. So, DAAR. I don't have a slide for DAAR.

JOHN CRAIN:

I'm not going to use any slides for this. I am going to ask people who are interested in details of where we're going with DAAR to come to our session, which is on Wednesday at 3:30 in 514A. So, for this group, I think everybody's familiar with the abuse activity reporting call and the reports have been out there. We've, since Kyoto, which is two meetings ago, we've been having a lot of discussions with various community members about how to improve the system, how to change the system and make it more useful.

We've had a lot of discussions about cutting data different ways. At the moment, we're really cutting data between the pre-existing gTLDs and the last round of gTLDs, we've had a lot of requests about cutting it by TLD type. You know, is it brand, is it something different, so we're

looking into that and we'll talk about that in the session later in the week.

But we've also had a lot of requests from country code operators about how can they participate. Most of that internal discussion is about what kind of paperwork do we need, how do we actually get the data we need, what data do we need, things like that. So as of actually yesterday, I have permission from legal etc. to start accepting data from ccTLDs who want to participate and to make it extremely simple. Basically, we don't need paperwork. If the ccTLD wants some kind of MoU paperwork, we have a very basic document that we could do, but we don't need it.

And the process will be to drop a zone file once a day, which is what we need for DAAR, as we need the actual zone file into an automated system, which will then go into the system and be used to do counts and do duplications, etc. from DAAR, and obviously it's voluntary. For the gTLDs, it's very straight-forward. We have something called 'The Centralized Zone Data Access System,' the CDZS. ccTLDs don't participate in that and we're not going to ask them to participate in that, we're going to do something that is very specific, only for this use.

And that's why I think some ccTLDs may want paperwork to say that and we have a draft we can share with them that basically says this is the only thing we will use it for, etc. We're not expecting to get all ccTLDs to participate, obviously that would be the perfect scenario,

because we have the entire ecosystem. But we have had, you know, more than a dozen requests and some of them are extremely keen.

The other side of the discussion that we still need to have is, what kind of data we get out to the ccTLD. So, for the generics, for example, we have a, think you're aware of, called MoSAPI, and I'm not 100% familiar with the system because I don't run it. And the generics can get their daily numbers back through that system. So, the question is, "Can we do it for the ccTLDs? Can we do it for others?"

So, there's still some work to be done about how the interaction works. But I have the go-ahead, so if anybody in a ccTLD wants to be included in those reports or to be included in the system, please come talk to me this week or send me an email, because we are now, go for, I think we have removed all the barriers we had, internally, about understanding how we were going to do that.

And then the other request to ccTLDs is, especially if you're going to start being included in the system, get involved in the discussions, come to that discussion on Wednesday where we've launched a mailing list, although we did a soft launch, we didn't launch it officially. But we're going to have an open mailing list where we're going to discuss the technical elements of DAAR and what improvements are needed or what new useful things that could be for the registry and the registrar operators, and the community in general. So, it's going to be open to everybody.

And the reason we did that is, one of the feedbacks we got is that, yes, we've been very transparent about all our methodology and how we do things, but not as much in some of the design discussions, so we want to just open that up to the engineers. And we understand that people will go off-track and we'll get people trying to have political discussions and stuff like that, but we'll try and kill that and keep it very technical.

So, that's really the DAAR update. Please come on Wednesday if you're interested. If you want to participate as a ccTLD, or you want to know more about that, come talk to me, otherwise talk to Samaneh, who's my colleague, who isn't here, but will be around all week and she is kind of the lead on that project at the moment. So that's really the update of where we are. We have quite a few changes that we're thinking about, about how we're displaying cut data, etc. but we really want to do that in combination with the community that's actually looking at the reports and using the system. So, we really want to have the mailing list and our discussions drive that rather than just us do it on our own. I think that's it, really. I'm happy to take questions.

JACQUES LATOUR: I think we're going to get a few.

BRETT CARR: I have a comment.

JACQUES LATOUR: Brett?

BRETT CARR: So, as a registry operator, ccTLD, and a lot of gTLDs, I think some synergy in the way you take data and produce data between gTLDs and ccTLDs would be really, really beneficial from our perspective.

JOHN CRAIN: So, our intention for taking data is to use the, what we call the 7 file access system, which is what is used for the -- so we are going to use exactly the same systems, we're just not going to push it through to the centralize zone data for others to access, unless the ccTLD community turns around and says, "Yeah, we really want to do that," but that's for you guys. So, we are going to try and make it work exactly the same.

BRETT CARR: Thanks.

UNKNOWN SPEAKER: Hello, this is Michael from .ch and [inaudible]. I think that's great news that we can participate. One question, why do you need the zone file? As far as I see, to calculate the at-use call, you just use total numbers of domains in the zone file.

JOHN CRAIN:

So, the main reason for asking for the zone file, there's two reasons. One is, that's the way we have our systems built to do it, to take the counts from the zone file. But the other questions that are ongoing is, how do you, for example, take names that are in the reputation feeds and flag them as no longer resolving? So, there is data, for example, a list of all the names would be just as good a zone file, but a zone file is data that exists. So, there is more study into the status of the data that we can do and that we tend to do where we need the actual domain names. It's just the zone file's a convenient mechanism for that.

JACQUES LATOUR:

So, I think one of the discussions we had internally was, for ccTLD to opt-in DAAR to TLD-OPS was one way. Because I think the reach that we have are all the security contacts and it's the right place to do it. So, I think that would be a good mechanism. So, do you have reporting today without zone files? Can ccTLDs opt-in and get some reports without, and then if they see value they can opt-in give the zone file and start that way?

JOHN CRAIN:

I'd have to talk to our contractor and our data scientist about that. So, I'm not saying no, because I don't know. I'd have to think about how they would make sure that they're doing uniform measurement. What we don't want is different mechanisms from all different people, right? So, if we have a slightly different mechanism for ccTLDs that doesn't

include the zone file, that might be a possibility. But that means we then wouldn't be able to do the measurements against the lists that count our names. So, I'm not saying no, but I'd have to go back and find out about that.

JACQUES LATOUR: Because there's probably an easier way to get more CC's to join if it's, you can opt-in, see what you have, and then -- is the Wednesday session after the DNSSEC workshop? Or the same time?

UNKNOWN SPEAKER: Same time.

JOHN CRAIN: I don't know, yeah.

JACQUES LATOUR: Same time?

JOHN CRAIN: We have a habit of collisions, with these meetings, so that doesn't surprise me. Unfortunately, it sounds like it's at the same time.

JACQUES LATOUR: Okay.

JOHN CRAIN:

For sort of reference, at the moment we've got about two dozen ccTLDs that are actually begging us to get their zone files so they can get into the systems. So, there are ccTLDs that are completely fine with it, but I understand the issues around zone files and the history around that, etc. So, if we needed to do something else, we'll have to look into that. It is going to complicate the system if we have to use a different mechanism for CC's versus gTLDs.

JACQUES LATOUR:

Any other questions? Alright. Thanks, John. Alright. So, today we're going to talk about the status of the document, the DR/BC plan. We're going to talk about the workshop and make sure we have all of our ducks lined up. So Dirk got unanimously approved as our fearless team leader for writing that DR/BCP work and he did a fabulous job on that. We had a team effort, but you were leading the thing, so that's really good. We had many, many discussions on how to take the science and make it an art, or vice versa.

But the focus was to try to make it applicable to small ccTLDs, I think we did that. And now, this afternoon, we're going to use the documentation, we're going to have templates that we built, extracted out of this document, and we're going to do a tabletop simulation and see if the whole thing works. So, I think it's a pretty ambitious activity for a bunch of volunteers to do, but last minute we,

I think we've managed to line up everything together. This is the old slide, right? Oh, yes.

UNKNOWN SPEAKER: Sorry. When we speak about small ccTLD, what involves that? Number of domains?

JACQUES LATOUR: Oh, smaller than Germany. Good answer?

UNKNOWN SPEAKER: Very clear, very clear.

JACQUES LATOUR: The goal was, initially, we wanted, we assumed, we made the assumption that the small ccTLD, with five people overall staff, would not have a DR/BC strategy, and if they were going to be stuck in a disaster situation, they wouldn't have the documentation or the experience to handle a disaster. So, that was our focus, to make it small. The documentation, I think, applies to all ccTLDs. If you're a cc and you don't have anything, this will be useful, irrelevant of size. So, I think we've managed to cover all the major points.

But the main thing with the DR/BC is you need to test it, you need to live it, you need to have experience, and if you don't do that all you have is a pile of paper that is useless. If you know how to live the paper, that's where it's useful. So, doing a tabletop exercise and on a

regular basis is important. But I think here we're going to show how to do the tabletop exercise using the documentation. I think that's where the value is; that's what the team decided and then we said, "Why not? Let's do a workshop." More stickers.

Did we want to go through the document here, or? We don't have to, I think we're...

RÉGIS MASSÉ:

I think we don't have to go to the document this morning. If we go there, there will be no surprise for this afternoon. As we sent the document two weeks ago, to the whole community, I hope people will be at the workshop this afternoon, have read it to be prepared to have a good exercise this afternoon.

JACQUES LATOUR:

So, we should send an email out to remind people to read the document again.

RÉGIS MASSÉ:

Yeah.

JACQUES LATOUR:

How many here read it? Okay, so we need to -- some skimmed it, some read it, so it'd be useful if you have a good skim, at least a good read on the document. But the point here, we have a document, it's a draft, and it's in pretty good shape. This afternoon we're going to

simulate the workshop; we're going to have the registry compromise exercise, so we have a good agenda, that's done. What we need to do is, we haven't talked a lot about observation, what the role of the standing committee is. We should be walking around and when we see people are jammed or stuck, we should document, take notes of where it doesn't work, I think. And then that will be useful in the end for some sort of lessons learned.

There are two CIRA employees coming in, project managers, that are in charge of the DR/BC plan, they'll be there. And they're more diligent than us at taking notes, so maybe we can leverage them to go table by table and take notes. Because in the end what I would like to do is have the list of things that would make the documentation better and then we incorporate that, and then publish that as our V1 for the plan.

RÉGIS MASSÉ:

Perhaps we can just add that the goal of the workshop, as Jacques said, to test the draft, the draft playbook for the moment, after getting feedback from the different teams this afternoon incorporated, the feedbacks and some change in the document before publishing it. And we will translate it, like the first playbook, in the seven languages to be the most useful for the community as it can be.

JACQUES LATOUR:

Question?

JOHN CRAIN: I think that's a very useful approach because that is part of the mantra of disaster recovery planning is to go through things, test them through exercises, adapt, and keep updating that way. So, I think that will be a good lesson for everybody as well, to learn how to actually do that.

BRETT CARR: And it should be repeated over and over again, because things change all the time; not necessarily repeated by us, but we need to drill into people that if you put this plan together then you need to look at it again next year, and the year after, and the year after, don't let it set gathering dust in a cupboard.

JACQUES LATOUR: So, this is where we're going to spend most of the time. So, logistically-wise, we have, we're going to have five tables, so we have around 50 registrations, excluding the steering committee member and the volunteers. So, we're going to do five tables of 10. We have yellow stickers, numbered one through five. People are assigned, randomly, one through five, so people have a table assigned with their name. Sorry, you guys might get broken up, or you might be together, I don't know. So, it is done random. I think the agenda -- do we want to share the agenda? No? We need to look at it, or?

RÉGIS MASSÉ: No.

JACQUES LATOUR:

No. So, we have an agenda. I think timing-wise, the workshop works. The first part is, Dirk is going to go through the document and overview the document. We'll go through the sections, why we have these sections. So, timewise, we have about half an hour, around that, I think. Yeah. So, everybody's going to have the DR/BC print out, we've got a bunch of those printed. So, we're going to have a copy people can take note on.

Then we'll have Q and A on the handbook to see how that goes, get feedback. We're going to set up the team and then we're going to spend a good 45 minutes to fill in the form, so we'll have blank forms. And the idea is that we're going to fill in the form in the context of a registry hack. So, everybody's going to think about how or what we need to do to think about what our risk is, what the impact, and all of that. So, everybody's going to participate in that. And after that we're going to talk about the forms, how they were useful, if there's any question.

And then we're going to spend an hour on doing the actual simulation exercise, the tabletop exercise. And we're going to have pretty cool collateral to do that. We went out of our way to make it interesting; I think. The only thing we don't have is a flying dragon. I think the rest we got. And then we have a break and then we'll do a debrief. We'll go through the observation, the lessons learned. And then we have about an hour reserved to have a discussion if we need to.

So, I think, overall, we've got three blocks, blocks three, four, and five lined up. And it should be interactive, it should be interesting, it should be; it's an experiment, we'll see how it goes. I think it's the first time anybody has tried this at an ICANN meeting, but we'll see how it goes.

Questions? Do we have all of our ducks lined up? Anything? Yeah. The only thing is that all documentation is on a train right now from Ottawa to Montreal. So, if the train is late, you know.

ERWIN LANSING: What could possibly go wrong?

JACQUES LATOUR: Nothing. Let me text, "Are you on the train?"

RÉGIS MASSÉ: Do you have a disaster plan for that?

JACQUES LATOUR: Yes, it's called online.

BRETT CARR: The first item on the plan is, 'Don't leave your plans on the train,' right?

JACQUES LATOUR: Yeah. Don't forget the stuff. But I split up, there's three people, and I split all the stuff in three people. So, we have resiliency.

JOHN CRAIN: Are they all on the same train?

JACQUES LATOUR: Of course. Alright. There's a lunch?

KIMBERLY CARLSON: Yeah, there is a TLD-OPS, it's a closed lunch for those who registered. It's on 720 at 12:15 and a special thanks to Nominet for sponsoring it.

JACQUES LATOUR: Ooh, thank you. Thanks, Brett.

BRETT CARR: Wait until you eat it before you thank me.

JACQUES LATOUR: Any other business? I think we're ready, so we'll see how it goes.

BRETT CARR: It's a record, the quickest meeting we ever had.

JACQUES LATOUR: Still an hour, I think. No? 30? 20? We can talk about the action item in detail if you want. Yep, I think all the work was done before this meeting.

UNKNOWN SPEAKER: I want to ask, besides your organization, of this group. Who has built or is trying to build DRP or BCP for your ccTLDs? Who is trying to build DRP, disaster recovery plans, or business continuity plans for his own TLD? Can...?

JACQUES LATOUR: I think pretty much everybody should be building their own. The level of maturity for where they're at could be different.

JOHN CRAIN: And this isn't just something for ccTLDs, this is something that every business should be doing and disaster recovery should be part of the standard business process, it just isn't always.

JACQUES LATOUR: But I think that the main reason we are doing this is that, I think in order, the maturity level of your plan, sometimes you have to bring in a consultant to your business to make it better. So, I think a lot of the bigger ccTLDs, they have the resources to bring in the expert to be better. And what we're doing is, I think, we're taking that experience that we have and trying to bring it to the smaller ones that can't afford

or don't have the expertise to do that in-house. And I think that is the experiment that we're doing.

So, certainly for Alejandra, like in Guatemala, this is useful for her organization, for the ccTLD to operate, because her BCP plan is 'call Fred.' So, you know, I think it's going to be useful for that. We also had some requests from a couple of CC's to do the session in Cancun, similar to this, depending on how it goes. So, I think it's Guadalupe, or something like that. So, some of these smaller CC's might benefit from this.

RÉGIS MASSÉ:

But I think there are two things important here. First of all, follow the disaster recovery plan, that you must prepared for that, and I will hope, hopefully we hope, that the workshop and the playbook will help you to have the tools, the materials to prepare them.

The second thing that's important to test, as Brett says. For example, at AFNIC .fr, we are making once a year an exercise. We designed an exercise with a few groups and train the team; so one day, one morning, we know that they have their exercise to be prepared. We start the exercise, sometimes it's DDoS mitigation on DNS, sometimes is the publication of false information in the DNS, technical things. And sometimes we try to see all the parts, crisis, human resources, what's impacted the company, and we have all this and I hope to train the whole company once or two times a year.

UNKNOWN SPEAKER: I think I suggested, I forget the word, best practice twice a year. I think it's an answer to your first question about breaking us. I prefer to split our group, maybe breaking us could be a little bit hard and could be useful because we're working on DRP and BCP for many months, and maybe we are not having an [inaudible] from our problems.

For example, if you compare, I went to a business continuity conference in London last year and I saw other resources like mining and banking; they speak their own language, and it's not useful for us. An experience can be a good reference, but in order to adapt it to a TLD or ccTLD, it can be hard with evaluation. Expect the people from ICANN, ccTLDs are like [inaudible], are unique in their own country, except the US maybe, and that's more difficult to evaluate and make comparisons with our DPRs and BCPs.

RÉGIS MASSÉ: Yeah, I understand that. Just remember that the initial goal of the workshop and the playbook was to prepare ccTLDs for a natural disaster, disaster and recovery plan. And we have to think that we can give advices, we can give things in this document, guidelines. Every ccTLD has to specify what is its own context, when you speak about that, [inaudible] for example, it's different in Japan, it's different in the islands of Africa, it's different in Guatemala, or Puerto Rico for example. So we started from there and after that we just wanted to make a general document for helping the community about that.

Just one thing more. We don't have many discussions on the mailing list, that was the initial point of TLD-OPS group, and when some ccTLDs have problems, we don't have many issues. So, one of the things we can all share is these kinds of things. Okay, we are making playbooks to help the community, but I think when there are problems, sharing with the TLD-OPS group, that is secure, that we know the charter is very important, too.

DIRK JUMPERTZ:

Dirk Jumpertz from EURid. I think, especially with regard to this exercise that we did, we wanted to create a playbook that is very practical, where you can basically grab the pieces that are relevant to the ccTLD business, or the TLD business in general, because there's not that much difference from a disaster recover perspective for a ccTLD or gTLD. And it's basically, has a modus operandi that would help anyone.

And I do agree that if you look at other businesses, like the mining industry, financial business, they very often have risk officers and people that are used to dealing with risk assessments. And we tried, in the book, to basically get everything out of that, and make it very, very practical so that you can basically grab the different tables that are in there and start filling them in and then have a go at it. That has been the setup of the entire book.

JACQUES LATOUR: We made it relevant for ccTLDs, so we pretty much all run the same operation. We have registry, we have DNS, we have people, so.

JOHN CRAIN: A few years ago we actually ran a series of trainings on disaster recovery from ICANN for ccTLDs and this was what was asked of them, “What would you, as a community, what would be useful?”, and this kind of book, where you can pick stuff out and start building your own thing was one of the asks out of that, that many ccTLDs asked for. I think there is a need.

UNKNOWN SPEAKER: Was it the ACRP?

JOHN CRAIN: Yeah, I think that’s what it was called. It was a long time ago. But that was one of the asks out of there that we often got was, “Could we have some type of playbook that we can pick and choose from?” so I think there is a need for this.

JACQUES LATOUR: When I joined CIRA, I had that ACRP book on my desk, it’s like a hundred pages of training. And when we were tasked to build the DR, I actually went there to see what was presented. But it was all the science behind the DR/BCP, it was super hard to take this and actually

implement it. And then I took it, threw it in the garbage, and I said, “We need to...”

UNKNOWN SPEAKER: You need something practical.

JACQUES LATOUR: Yeah, something practical. And I think that’s what -- tonight the feedback is, we want to know if it’s practical and useable. And I think it should be.

DIRK JUMPERTZ: I hope so. That’s basically the goal. As a matter of fact, most of what is in there is highly inspired by our own disaster recovery plans and business continuity plans. So, it is something that we are already using and my experience with it is, it’s practical enough so that everybody understands it. You don’t have to have some kind of high-level documents that are telling you what you should be doing. This is different, these are practical instructions.

JACQUES LATOUR: Anything else?

RÉGIS MASSÉ: Just one thing about that. As I said, the goal of this document is to give guidelines and help for the ccTLDs. I just remembered, during the

first workshop we did in Puerto Rico, in ICANN Puerto Rico, we did a workshop to just get feedback and brainstorm about this topic before writing the document, and one person from the ccTLD in Africa, I don't remember which one, said, "We are only two people in the registry. So, if one person is not here, it's a disaster for the mechanism of the registry." "Okay, we can't give clues for that, but the document can help to..." Okay, having all their information, that's to keep in mind that if the people are not there, there is something written on the paper, the names of the contact, things like that. So, it's a good practice.

JACQUES LATOUR:

Anything else? Kim, you want to say something for sure, no? Alright. That's it. So, we have lunch at 1:30. Thank you. Meeting closed.

[END OF TRANSCRIPTION]