
ICANN75 | AGM – Tech Day (2 of 4)
Monday, September 19, 2022 – 13:15 to 14:30 KUL

TED BARTLES: Hello, everyone. Thank you very much for coming. If we can go to the next slide, please. As Eberhard just mentioned, my name is Ted Bartles. I'm with the ICANN Meeting Team. And I'm specifically the meetings production manager. So I manage everything from audiovisual, room layouts, venue, logistics, and basically all the in-room aspects of the meeting. And then, my colleague.

SARA CAPLIS: Hi, everyone. My name is Sara Caplis. I'm the manager of the Meetings Technical Services Team, which falls under the engineering and IT function at ICANN. My team is responsible for the ICANN Meetings network and then managing the remote participation services.

TED BARTLES: So essentially, I guess we've been asked here to talk a little bit about our hybrid plan for The Hague, the last meeting in June—how we planned it, how we built it, and how everything worked out. To us, that's a bit of a difficult question because, basically, what we did was really just to expand on what ICANN had. So we're going to give you a little bit of an idea of what we did expand on, and what we did, and what we asked you to do. And then, I think more so, we're going to ask you for your questions because maybe you guys have a little bit more detail you want. So if we can go to the next slide, please.

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.

What we did, we broke it down into three different groups to enhance and to improve upon for the hybrid experience. So we had the onsite meeting experience, which was essentially my position, the meeting rooms, the meeting room tech, and the AV setups. We had the virtual experience, which Sara managed. And then, we also had what we think of as the most important part of it all was the participant experience and the coordination of the participant activities and how they would participate. Next slide, please.

From the onsite experience, it was pretty simple. I already have what you can see in the rooms, a pretty simple setup. We've got screens. We've got microphones. We've got projectors. And we have connections to the virtual world. This is something that we've had since I started with ICANN in 2009. So it wasn't a matter of adding something. It wasn't a matter of coming up with a new plan. It was more a question of how can we make everything better and how can we make sure that we have equitable participation between both the virtual audience and the in-room audience.

So the simple improvements from my side were we added one microphone per person. This was for health reasons as well as camera angle reasons. In the old days, we had a mic between two people so you end up with a lot of bad microphone habits. Who's got the microphone and who's actually in the camera shot?

The other thing we did is we added an additional screen to each room. This is the screen that, as you can see over there, shows the virtual participants and the in-room participants. In some of the rooms, we just show virtual. But for the most part, this guarantees that everybody in

the room and outside of the room can see both parties with no issues. We've always had, for the longest time, cameras in all rooms. But it's something that we made sure we improved and we worked on. And the good news of it is now we've also bought new cameras, which you'll start to see throughout the next few meetings.

We added additional time between sessions. I think all of you remember, during the virtual meetings, things got a little late nights, time zones. So this was mostly to respect the virtual participants and give them a chance to step away from their laptops for a few minutes. And also, it gave us a chance to really spend some extra time between sessions testing the virtual side, and the in-person side, and making sure everybody was all set. People from afar could be heard and seen.

We did a lot of work updating the participation manager kind of roles. We've always had a remote participation manager, someone who could either be virtual or in-person. So we decided to add one more person to that team, specifically to help direct the in-room, and to make sure everybody was in their seats on time, and do everything they could to make sure that the virtual audience was getting the support they needed in line with the remote participation manager.

Then one of the last things we did, which as you'll see, has changed a little bit, was we did some physical distancing of the seats. Obviously, this was for health reasons. In The Hague, we had about 1.5 to 1.8 meters between seats. Here, we've reduced that to around one meter. And that was just something to give everybody a chance to have a little space, keep the microphones a little more distant, allow the cameras to

work a little bit better, and obviously, for health and safety. Next slide, please.

So this is just a very quick, simple setup. This was actually the GNSO room in The Hague. Just an idea of this is how we think about it. We've always had rooms set up, always go to the venue before we get there. So limited capacity. Minimal area for observers to queue up, which was always a problem. That's why we gave everybody extra microphones. And then we also gave the extra space for the tent cards in the AV and asked for tent cards on all the tables. Next slide, please.

Just another quick little description. We broke up in terms of how can we support everybody the most. As I mentioned, we added the onsite participation manager. This is a member of the Org assigned to the physical meeting room as well as managing Zoom. They help facilitate the meeting through content sharing and queue management.

We had our virtual participation manager, which is someone we've always had. And this is a specific Org member assigned to just support the session for the remote participants and also help facilitate the queues with the participation manager.

Then we have the session facilitator or leader. This is the overseer of all of this. This is, again, an Org member assigned. They're always in the physical meeting room and Zoom. And they oversee and make sure everybody is present and doing their best to respect both sides of the equation. Next slide, please. I'll hand it over to Sara.

SARA CAPLIS:

Thanks. My team, as I mentioned, manages the virtual or remote participation side of things. We were fortunate enough that we had already made the switch to Zoom back in 2018. They were a company that really responded quickly to the pandemic. So when we were at the very early stages of doing virtual meetings, we were able to include some new features into Zoom, which we really wanted to retain when we came back to the hybrid setting.

One of the things that we implemented was the raise hand feature in Zoom for all participants, regardless of whether they're onsite or virtual. This just contributes to equitable participation so that no matter where you are, you have a chance to get in line to speak. We worked with our language services team and our partner, which is Congress Rental. They're the ones managing the onsite interpretation as well as the remote simultaneous interpretation sessions.

Starting, I believe it was at ICANN 68, we started to integrate their system into ours so that participants could listen to the interpretation within Zoom rather than just mp3s. So we have retained that for onsite participants as well as the virtual participants going forward.

We've really encouraged our virtual participants, rather than just being listen-only, to take part in sessions. That's obviously through the raising hand but also to share their video so that we can see them on the secondary screen in the room. It contributes to their presence in the session.

As we mentioned, in addition to the timings between sessions, we do have a somewhat reduced schedule this time, as we did in The Hague.

That just allows our virtual participants to deal with the time zone a bit better and take part in sessions around their regular scheduled life.

In addition to the onsite participation manager training, we also created various guides for participants, depending on where they were located. Then we've done several trainings of our staff who support these sessions, whether it's policy support staff, the GSE members. We do a regular training with them to explain their role in supporting these sessions.

Then finally, we did implement the Zoom Waiting Room video, which explains how people can participate regardless of where they're located. Our main priority was not to lose features that we had implemented during the pandemic, and take advantage of them, and leverage them for the return to hybrid meetings. Next slide, please.

TED BARTLES:

Back to me. This comes to the third area in all this improvement. This is the area that we really, in the end, decided, this is the important part. I know everybody wants to know, "How technical was it? What did you do, technically." As I said, we didn't really do a lot technically. We improved what we had.

But what we really did, after that, is we decided it was about the participants, and us, and everybody together that really needed to work to keep it as equitable as possible. What we did is we really asked you to help us, and to follow some simple rules, and to work together, and to respect each other. And as a group, we were the ones that all made this hybrid work.

If you want to break it down, we broke this down into three kinds of participants. You've got your session chair and/or moderator, and your virtual participant, and your onsite participant. So we looked at it from those three sides and we said, "What does each group need to feel like they belong, to feel like they can work together? And how are they going to work together?"

So this is where we just started to add a few other simple things. We said, "Why doesn't everybody get in Zoom? That way we have one single queue." We asked everybody in the rooms to remember that people were virtual and to avoid side conversations. Again, another reason we gave everybody a microphone, so that people didn't start to go back and forth with each other and leave the virtual participants out.

We told everybody that we were going to put everyone on camera. And we also said, "We're going to bring all the virtual people into the room." So in essence, it was about as simple as you can get. But it was just a matter of some easy planning. Next slide please, I believe.

It was a question of how can we all work together? In the end it was just a mutual agreement between all the parties to remember everybody and to work together. That's our simple explanation of it. I think maybe you have more questions, though. I know Eberhard asked us here because I know there's some interest in this. So for our next remaining time, I want to leave it up to you guys to ask any questions you might have.

EBERHARD LISSE:

Are there any questions from the floor? While we're looking at that, I was interested in the technical setup last time—how many containers and so on. This is how many kilometers of wires and so on.

When first I was a contact then I contracted COVID around The Hague meeting, I couldn't come, so I chaired Tech Day virtually. And I must say, that went like no problem, really. It was really impressive. I also vice-chaired a meeting of a PDP working group. That all worked out very nicely. So now that I'm online, I notice that I tend to forget to ask for remote comments because I'm local again. But this is important because we can increase the audience, not only from on-site but also the virtual participants. We have had two presentations from onsite. And the last one in this block will be virtual.

So this hybrid concept is something that I think we should keep on for the foreseeable ... We should keep it on permanently because we will increase the participation, not just in Tech Day but in all our meetings. It will be better for people who can't travel. From my continent, it's often very difficult to travel.

But Zoom works. I shudder if I think about Adobe Connect, which we had been using a while back until Puerto Rico. I couldn't imagine to do this in that system. At least Microsoft Teams, which I am not very fond of, which also works. I'm not complaining about it. It's just that I don't personally like it so much. And I don't really think an open-source solution has yet to come up that does it very well.

What I liked about the IETF meeting that was in Vienna is that when you enter a room, you can check in with an app. Then you are counted in

that room and you are able, on your app, to raise your hand as a participant. Here, as a participant on-site, you have to still log into Zoom to raise your hand. Since everybody does their e-mail during all these meetings anyway, they can just run Zoom without being too much interrupted during their work.

But I like this idea to be able to check in by your phone. Then, whether you're using your laptop or not, you are able to. That was something that we should maybe look at if this is something that the audience thinks or your users, your members think we should do. Is there anything else? You're getting off easy. Anyway, that's it, then. Thank you very much.

SARA CAPLIS: Thank you.

TED BARTLES: Thank you very much. And I can say, just in response to your comment, the hybrid is here to stay. The second screen is here. The microphones are here. The cameras are here. We don't intend to remove any of this once things get better.

EBERHARD LISSE: Anyway, we have noticed a tendency from ICANN to inflate rather than to contract. Okay. I'm going to release you with the final barb that everybody remembers that the next [African] meeting is going to be Washington.

That leaves me, then, to announce Brett Carr. He's a new member of the Technical Working Group. He's also a member of the CSC Effectiveness Review Committee. And he and Donna Austin, who I don't see—maybe she's remote—they are going to give us ... They were willing and able to get in on short notice because one of our presenters absconded and the other one got bumped off his flight. So, Brett, you have the floor.

BRETT CARR:

Thank you, Eberhard. I don't know if Donna is in the room or remote. If she's not, then I can do the whole presentation, I guess. I wasn't planning to. But if Donna's not here, then I will crack on. Next slide, please.

For those who don't know, the CSC stands for the Customer Standing Committee. This is a group of individuals that represent the customers of the naming function of IANA. It was born out of the IANA transition away from the US government, however many years ago that was. And basically, it exists to monitor the SLAs that were put in place for the PTI/IANA to make sure that they're performing correctly and the naming customers are getting the correct service levels from them. Next slide, please.

This graphic just shows the structure of the CSC and what it's there for. There is a contract. There's a legal contract between ICANN and IANA/the PTI for them to provide naming services. The CSC are there to monitor the performance of them. And then, every five years, there's something called the IANA Function Review that checks that IANA are also performing correctly according to the contract and has a look to

see if there's any changes needed. There's also something called the CSC effectiveness review that happens every three or four years. I can't remember the exact timescale. That's what, effectively, this presentation is about. Next slide, please.

This slide just talks about the various different reviews that happen in the CSC. Obviously, the CSC has a charter that details what its responsibilities are. Then there is, as I said previously, a review of the effectiveness of the CSC. The first one was started in October 2018 and was completed a little while later. Then some recommendations came out of that report, which were implemented in 2019.

Then there was a second review started in October '21 which is now out for public comment. We would encourage anybody who's interested in the room to have a read-through. The review is not very long. It won't take very long to review it. And we would like to hear people's comments on it, pro or against, obviously.

Then we also have something called the IANA Function Review, which happened in October 2018. Took a little while longer. That finished in March 2021. There was four recommendations which were all admin in nature. They were to provide clarity around the naming function review triggered by non-satisfactory conclusion of remedial action procedures, which we'll talk about a little bit later, I think. Next slide, please.

The second Effectiveness Review, which is finished to the extent that it's out for public comment, concluded that the CSC is operating effectively. There is a high level of effectiveness due the commitment,

knowledge, and expertise of the people appointed to the CSC, which obviously makes me feel very proud because I've been involved with them for a few years now.

There is sometimes some problems about attendance of the meetings. The CSC is quite a small group so there's only actually four members of the team and then various liaisons to different other parts of the ICANN communities. So one of the things that has been suggested in the review is that the members of the team that have been appointed by the various supporting organizations also look for alternates so if they're not able to attend the meeting, they can send somebody else, and then to inform the appointing organizations.

One of the other things that they came up with in their review was concerns that will the CSC be able to attract capable volunteers. One thing I would encourage people, if this subject interests you at all, do some digging into it because we do look for people. We're not looking for anybody right now because we've just appointed two new people. But in a year or two's time, we'll definitely be looking for additional people as people's terms expire.

There is no expansion of the limited role of the CSC. And the CSC is in close cooperation with the PTI to develop a framework for records and reviews of the IANA Naming Services SLAs. So that's something we're looking at doing, potentially, in the future. But again, we're looking for public comments on the findings. Next slide, please.

There are various SLAs that the CSC monitors and reports on. PTI is compliant with them. These metrics are put in place as part of the IANA

Transition many years ago. There are currently 64 metrics within seven different groups for monitoring various different things like IANA technical checks and the time it takes staff to process things to make sure that changes get done, gTLDs get created, name server changes get made, contacts get updated, etc. All of those things that gTLD and ccTLD registries might ask IANA to do. Those metrics are monitored to make sure they're happening in a timely manner.

Those metrics are directly related to the Naming Function contract. The SLAs, if anybody's interested in them, are listed on the PTI's website under the SLA table. Next slide, please. But none of those SLA's would be a massive surprise to anybody, I don't think.

The PTI publishes a monthly report with measurements based on performance percentages of each of those SLAs. And that report is made available to the CSC before their monthly call so that the CSC members can review their report and discuss any issues with PTI, if there is any issues. Obviously, we want to know what caused it and if it's something of concern or something that they're already addressing, etc. Next slide, please.

This table shows a summary of the CSC findings over the last six years. You can see over at least the last three years, almost every month there's been 100% compliance and they've hit the SLAs. There are a few indications that you can see where the performance has dropped to 98.4% or 98.5%. That indicates one missed SLA. I can't remember the specifics of what they missed but I can tell you that it was explained satisfactorily in every single case. So the CSC have currently got no

concerns about the performance of PTI whatsoever, based on their performance, as you can see in the table. Next slide, please.

There has been some changes to the SLAs over time—very little, I have to say. I think that when these SLAs were put in place originally, the team that worked on that part of the IANA transition did a really good job of putting those SLAs in place. We found very little reason to change them over time. But in 2008, the process for amending the IANA Naming Function Service Level agreements was done. And the procedure for modifying the process—for amending the service level agreements—was also put in place.

Those processes became effective with the amendment of the contract in March 2019. And since then, we made three changes to SLAs. Oh, hello, Donna. I'll let you make comments at the end when I've finished. So there've been three changes to date—three technical checks that are completely operational; some new SLAs for internationalized domain name label generation tables, which are also completed; and one for ccTLD creation and transfer, which is also in place now as well. Next slide, please.

One of the things that the Effectiveness Team Review came up with in their review—and I'm sure Donna will talk about this a little bit in a second—is that there is an idea of maybe doing some kind of regular review of the SLAs. This is not something we've done in the past, generally. But it's probably a good idea to ensure the SLAs remain current and relevant. At the moment, no one's responsible for initiating that review process but it's probably something we should do.

So there's a question mark about whether the CSC might undertake that regular review. And it's this kind of thing that we would like input from the community on. So please, if you've got a view on this, leave us a public comment. Next slide, please. That's questions. So I'll hand over to Donna and see if she's got any other comments.

DONNA AUSTIN: Thanks, Brett. Apologies for the lateness of my arrival. Have you been through the full—

BRETT CARR: I've done [inaudible]. I've done the whole [inaudible].

DONNA AUSTIN: I'd love to. All right. So I guess Brett's got it all well-covered. The only thing I will say, as a member of the CSC Review Team, is to thank the CSC for the work that they do. This is the second Effectiveness Review that I've been part of. And I think the success of the CSC is unfortunately one of their downfalls as well, in that it's particularly boring, I think is some of the language that we hear.

So one of the things that we were conscious of, I think, in doing the CSC review was how do we address the possibility that it becomes difficult to find volunteers to fill the roles in the future? We've been very fortunate to date that we've been able to get good people in the roles. So we've got a few recommendations in there that are intended to address that possibility.

I don't know whether it's still obvious to people but the CSC fulfills a really important role. And unfortunately, we don't know how important that role is until things start to go wrong with IANA. They're there to make sure that everything is working in good order. And so far, they're doing a great job.

But to the public comment forum, please. We're interested in your comments. We'd welcome your comments. And as Brett has gone through with the SLAs, that's something that we really would appreciate your input on because if there's enough support for the question we're asking, then that's something that we can finalize as part of a recommendation that we think there is value in the CSC working with IANA to review the SLAs. Thanks, Brett.

BRETT CARR:

Just one final note from me. If this has perked in any interest in you at all, there is a more comprehensive presentation of what the CSC does in tomorrow's CCNSO meeting. And if you've got any questions now, I'm happy to answer them. Or speak to me afterwards.

SUHAYB ALGHUTAYMIL:

Suhayb Alghutaymil from SaudiNIC. My question, how CSC is gathering this report? Do they have access to root servers and generate the report, or asking by e-mail, or how?

BRETT CARR:

It's the IANA/PTI who actually generate the report. So they've got monitoring systems in place to check that different things they're doing

are being monitored correctly. So they produce the report with automated monitoring systems and then we check the report that they produced.

EBERHARD LISSE:

Okay. Are there any other questions from the floor? Good. Thank you very much. Thank you very much in particular for being able and willing, on such short notice, to step in to assist us with filling a slot. Thank you very much from the floor as well.

So that leaves us a little bit early with the presentation from ICANN Technical Engagement in this block. Champika is on remotely, from what I understand.

CHAMPIKA WIJAYATUNGA:

Yes. I am presenting.

EBERHARD LISSE:

Then go ahead. You have the floor. Thank you.

CHAMPIKA WIJAYATUNGA:

Thank you very much. Welcome, all. This presentation is all about ICANN Technical Engagement. My name is Champika and I will be presenting remotely. I also have my colleagues onsite, Yazid and Nicolas. They are onsite as well. So if you have any questions or if you want to discuss anything further, feel free to discuss with them. Next slide, please.

First, let me just try to introduce our team. ICANN Technical Engagement is a function within the Office of the CTO. So you can see who we are. Our vice president is Adiel Akplogan. Adiel is also actually attending the meeting physically. So feel free to talk to him as well.

Usually, when you think of technical engagement, we have different regional leads. The requirements that you have in one region may not be the same as in other regions. This is also why we try to actually provide these technical engagement matters on a regional needs basis. So we have a regional team. You can see who we are. And again, as I said, Nicolas and Yazid are also actually onsite. Adiel is also onsite. We also have Carlos, who is also going to take care of the trust and public safety-related engagement, especially with the law enforcement agencies. So Carlos will cross the regions in this case. Next slide.

Now let me just go through, actually, what we normally do. I'm talking from the Technical Engagement Team's perspective. We deal with, basically, disseminating the ICANN technical work to the community. So this is actually mostly in the form of capacity building, outreach, and also actually working with the community on various technical support matters.

As I mentioned to you earlier, we have regional focus. So there are different languages that we need to take care as well. So because of that, even our contents, materials, we also try to cater for different languages. So at the moment, we also offer these technical engagements, the capacity building, in English, French, and Spanish as well.

Our focus is mainly on the DNS, DNS security, and DNSSEC. If you think of the capacity building, actually this is one of the very important tasks that we do as a team. So we offer various hands-on training, webinars, and also actually outreach activities as well. We also support various other teams in the organization, especially the Stakeholder Engagement Team and so on, in terms of providing these activities to the community.

We have a lot of collaboration involved with the community. Something, actually, what we do closely is also to work with country code top-level domain registries, ccTLDs. Currently, we have been working with a number of ccTLD registries in terms of improving their operational processes, especially in relation to their technical details and configuration issues, especially in terms of DNSSEC. Also DNSSEC signing and preparation, readiness assessments and so on. Now, it doesn't mean that we go and try to do it for them. But we try to actually work with them, guide them, and provide assistance so that they can try to perform those activities.

We also work very closely with the network operators, the ISPs and so on. Especially with regard to this, we work very closely with the network operator groups, especially in the regional organizations. And there, of course, these NOGs are being scattered throughout the world. So we work very closely with them.

Then also, other organizations such as TLD organizations like LACTLD, or AfTLD, or APTLD and so on. We have had various collaborative capacity building events with them and also provided a lot of training to ccTLDs and the related communities.

Also, one of the other functions that I mentioned to you is about the law enforcement agencies, working with them. In different regions, again, there are different regional leads as well. We work closely with bodies like INTERPOL, and Europol, and so on, and also various CERTs, and provide these supportive activities, especially in terms of training and capacity building.

One of the other things is, actually, we also have an online learning platform called ICANN Learn—probably most of you know this—especially for newcomers who want to learn about ICANN and various content related to ICANN. So we do get involved. So we do get involved with developing these contents as well. And we try to actually develop courses related to what we do.

In the recent past, we have also been working on a virtual lab platform as well and we hope to integrate this into our ICANN Learn platform so that people who want to do some hands-on type exercises and so on, they could try that through ICANN Learn as well.

Again, depending on the specific needs, specific requirements, and so on, in different regions, we try to take some regional approaches in providing these supportive activities. Next slide.

In terms of various training and capacity building, we offer a number of training courses. These training courses are all published in the ICANN public website. You can see that. This is just a snapshot of that public website. There are 15 training courses currently on offer, again, starting from very basic, to intermediate, to advanced-level courses as well. And

depending on the need from the community, these courses are being offered.

So if you think that, in your organization, in your region, or in your country and so on, if you need some capacity building activities, you can always send us a request. You can submit a request through this form. One of the things is also, actually, you may need to have some lead time for us to actually do the planning and so on. So that's something for you to consider.

Some of these training courses were offered prior to the pandemic. But then, especially during the pandemic, we tried to cater for online, remote-delivery options as well. So we have catered some different courses on face-to-face basis as well as on a remote basis as well. So we could accommodate both of those. Next slide.

One of the important things, I was also explaining about the online lab platform. Basically, this is a cloud-based platform. Previously, before the pandemic, of course we used to go to different places, different events, different countries and so on. And we tried to offer these courses. But then it was quite challenging.

One thing is that we had to carry this equipment, especially the hardware. Or the products and so on, we had to carry. The other thing is, if we are not carrying, then we have to ask the local host or the local organization who is organizing that training, to provide this infrastructure for us. So this involved a fair bit of administrative work. And also it wasn't easy all the time.

Also, the other challenge we found was, especially in some countries, especially in some of the developing economies, developing countries, we found that connectivity was not really very good. So we found that the connectivity was dropping. And sometimes we could not perform quite well. So we had different challenges, of course, in different places as well.

So this is also one reason why we thought, as an option, in addition to the face-to-face option. Also, at the same time, we had this pandemic came, the COVID-19. So during that time, we could not travel anyway. So we started actually deploying this cloud-based lab infrastructure. So during the last two years or so, we have been working on this. Of course we have progressed quite a lot in this regard. Currently, we have a full-fledged lab that we could offer on a remote basis. Later on, I'll go through that as well. So this is quite very efficient.

Also, one of the other efficiencies of that, we can accommodate quite a lot of participants. Sometimes, when doing onsite, we had to limit the number of participants depending on the room size and things like that. And also, the other thing is we had, with regard to the connectivity issues, because now it's all cloud-based, we could actually pick the nearest, best place so that we could get a better response in terms of the response time. So due to these reasons, found that a cloud-based infrastructure actually works pretty well.

So this is hardwareless. We don't need to actually carry the stuff and go through with all those risks and so on. And this is also very much scalable. We could run parallel training because sometimes we also found that if we are doing some training in LAC region, or at the same

time if we have another training request in the Africa region, or Asia Pacific, or wherever it is, we could run these training courses parallelly as well because we can run many of these labs independently. So these are various scalable options that the cloud-based platform has given to us.

We also plan to integrate this into ICANN learning platform, as I said earlier, so that if some participants are keen, they could also actually try to test this out independently by themselves as well. Next slide, please.

I've been going through quite a lot of capacity building training stuff. One of the other important functions of the technical engagement team is also actually work on different projects and so on. I'm not going to go through all the projects here.

But one of the things I wanted to highlight was especially here, which is actually a guidebook on DNSSEC deployment. So this actually aims at the ccTLD registry operators because they can basically use this as a guide, and try to understand, and try to deploy DNSSEC. So that is the objective of this guide.

OCTO, the Office of the CTO, we publish various technical publications. So this guidebook is also actually published as an OCTO publication, the OCTO 29. So if you refer to this, especially if you want to go through in detail, as a guidebook, as a guideline, how to deploy DNSSEC, especially from a registry operator's point of view, this explains well in detail. It talks about all the focuses on registry manager stuff, registries, registrars, and so on, and basically anyone who wants to deploy

DNSSEC. I think it will be quite useful. So we would encourage, in case if someone wanted to deploy DNSSEC, to actually refer to this guide and have a look at it as well.

We are also seeing, actually, some ccTLDs, some registries, they have started adopting this guidebook. So it is actually a good sign. And hopefully, in the future, we get more and more registries, actually, who follow this as well. So there is the link published, in case if someone wants to have a look. You can see it in the slide. Next slide.

Some of the observations, actually, we have seen over this period, nowadays, we see there is great awareness, especially in terms of improving security in organizations. This is also why we have been getting constant requests for various capacity building activities, training activities, and so on, which is also a good sign. So we see that organizations try to spend more resources on improving their security.

And also, actually, we do get requests for translated [inaudible] as well. This is also why, actually, we have a regionalized approach again, especially in terms of translated [inaudible]. And also we can deliver in a few different languages as well, in this case. So this includes, of course, the hands-on labs as well.

We also found that some ccTLDs, they are a little bit hesitant to sign their zones. I think it is mainly because there is a little bit of fear or doubt that they might break the things. In case they try to deploy DNSSEC, especially one thing, of course, is when you get to DNSSEC, they might think it can be a little bit complicated dealing with all the signatures, dealing with the keys and things like that. And also, actually,

this might also break their DNS, basically, in case the signatures get invalid or whatever the reason or the performance issues. These are general understandings that we found that many of the organizations, they have.

But of course, during training and so on, we try to go through in detail about all these aspects, especially the signing aspects, the key rollovers, and all that we do in detail so that the participants actually get to try this out in a lab environment so that they can understand how they could perform that. Initially, we tell them to do it in a test environment aspect and then to gradually deploy that in their production environments.

We also try to provide some technical assistance, as I said, support. We do get various queries and so on and we try to support them and help them. But then still, there is quite a lot that we have to do in terms of supporting various organizations. Not only, of course, ccTLDs but also the other operators as well and whoever in our parties who are running the DNS infrastructure, and second levels, and so on. So we try to do this gradually.

And also now, we try to work closely with the network operators as well, especially in terms of the validation aspect. We also see that, in all the regions, we are working with a lot of operators and try to improve the validation aspects as well. Next slide, please.

Let's try to have a very quick overview of the lab infrastructure that I was talking about. At the moment, we have a set of individual virtual machines that we run. We run LXC containers in these. The platform

that we are going is AWS. So initially, we went with LightSail and then later we started with EC2 as well, mainly to cater to larger demand in terms of memory, the capacity, and so on.

Now, in fact, we can actually run even 350 containers or so. We have done that in some events because every lab user would get several containers running different servers, like authoritative servers, the recursive servers, and so on. So they would configure all that in a real environment. So one would run many containers. So we can, in fact, depending on the number of participants, number of attendees, we could cater a large number of containers, actually. So this is a very good thing comparing to offering something in a physical environment. So the cloud-based lab supports that.

We also do support IPv4 as well as IPv6. This is another feature that we support. And also, we have done a lot of scripting in the lab to make it easier, especially in terms of deploying the lab initially and setting up the lab, the initial configs, and things like that. So there's a lot of scripting that has been done so that things are much more automated there.

Generally, we would run multiple recursive name servers. One of the other things we wanted to actually do was to have some sort of variants with the different name servers as well. This is also why we try to go with multiple name servers. So we have BIND, we have UNBOUND to run the recursive name servers. And then for authoritative name servers, we also have NSD, BIND, and also we have OpenDNSSEC as well. We also have routers. We do FRR over there.

And in some cases, together with some of the collaborative organizations that we work with, we also actually run some labs on RPKI or Anycast related labs as well. So in these cases, we also actually run RPKI validators and the routers for the Anycast labs as well.

So in all these cases, the lab has been very scalable. We could dynamically deploy the infrastructure quite easily and also cater to many participants as well. We also try to, depending on our upcoming needs, future needs, and so on, we also try to even expand the lab as well to cater for those demands. Next slide, please.

This is just for you to get a snapshot of the lab overview that I was talking about. This is actually every participant who joined the lab would get this kind of containers so that they can run their different servers within themselves. So they belong to a certain group. Within this group, they have a number of containers—each participant, basically.

So you have a client machine so we could test a lot of queries and so on. Then we also have a hidden primary and also running multiple resolvers as well. And also, we have secondaries as well. Again, as I told you earlier, we run different name servers in different machines. So this is the general overview of that.

We have seen that this lab has been very effective in terms of getting the participants to really understand the concepts, basically implement the lab during our DNSSEC validation in their resolvers, do the signing, as I said, and also doing the key rollovers and so on. So in all these areas, we have seen participants do actually learn quite a lot. And we have seen the results quite effectively. Next slide.

This is something which I wanted to share, especially from a technical engagement point of view. Especially over the last two years, we have offered quite a lot of training courses remotely. Generally, we have been offering around five to six training courses every month on average. So that's the frequency. But this has been increasing as well.

But now again, with the travels, some places are opening up and so on. We also hope to go to different events and physically offer some of this training as well. But good thing is that with the lab, even if we go to a remote place, we could go to a place or event. We could still try to have the cloud-based lab running as well. So we have that scalability as well.

So the main message that we want to give is in case you want to try this out, these sort of activities for you to be familiar, especially with the DNS, DNS security, DNSSEC and so on. Even we provide this lab to law enforcement agencies as well because they also use some of these labs, not of course for DNSSEC type configurations but mostly for using RDAP type queries and things like that.

So in case, in your environments, if your organizations have some needs to do some capacity building, to do some training, even to actually use this infrastructure and so on, do talk to us. We are quite happy to work with you and provide you the necessary support.

As I told you earlier, my colleagues, Nicolas, and also Yazid, and Adiel, they are all onsite. And finally, I also actually want to give some dedication to our teammate, colleague who passed away recently, because he put a lot of effort to prepare this lab as well, along with the

team. So we want to remember him also at this stage. And thank you very much. That was about Paul Muchene, I was talking about.

So thank you very much. And if you have any questions, I can take now. Or even Nicolas. As I said, Nicolas and Yazid are in the room, physically. We can take those questions. Thank you very much.

EBERHARD LISSE:

Thank you very much. While I'm waiting for questions from the floor, how do you measure success? Number of ccTLD managers participating or how do you measure it?

CHAMPIKA WIJAYATUNGA:

Yeah. Actually, it's a good question. It is always something very hard to measure. Of course, we do have those numbers and things like that—the number of participants. We also try to do some post-training surveys and things like that.

But actually, we have also seen, over the years, a number of ... Say, if you are speaking of ccTLDs, we know the ccTLDs whom we have been working with. And also, we have seen them deploying DNSSEC. So that is a direct success or a direct measure that we can take. So if you take the past few years, we have seen a number of ccTLDs who have deployed DNSSEC, who have signed their ccTLD zones. And then directly, we have also worked with them. That's something we could easily measure.

And also, we do consider some of the DNS measurements—say, for example, things like the validation measurements and things like that.

So if we have worked with a large network operator in a particular country or in a region, and if we see that, in that country, the validation stats or the validation measurements have gone up, we could actually see some sort of correlation there as well. We are also actually working on some projects as well to further find out about these measurements, basically the DNS demographics. Hopefully we can improve these things.

EBERHARD LISSE:

I think this is very important to measure the number of signed domain names, not whether a ccTLD has implemented DNSSEC per se. I know of two ccTLDs who have participated in this program, who have had to un-sign, who've lost their keys. So I think measurement is extremely important but we should go into real measurement that tells us actually not just input, how many workshops conducted, how many ccTLDs have done something. But we must measure success. In any case, anybody else from the floor?

NURAH MUHAMMAD:

My name is Nurah from MYNIC. I would like to know whether ... Do you have any plan to organize a certified training where you can measure the competency for those attending the training? Do you have any plan for that?

CHAMPIKA WIJAYATUNGA:

At the moment, of course, we do not offer any certification or certified training as such. But of course, we have worked with some

collaborative organizations, especially some NOGs and so on, where they would actually, after the training, they would do some sort of assessment and provide some certifications through them. But not directly through us.

So this is something that we can discuss and see the practical possibilities of doing that. Thank you for bringing this up. To answer your question, at the moment, we do not have that sort of certified training.

EBERHARD LISSE:

I don't really understand the question. Other than putting a certificate on the wall, such a certificate doesn't mean, really, much in the actual operation. All of this is very well-described. All of this can be very well studied from RFCs and from other documents. If a simple gynecologist in a developing country can figure it out, hotshot masters and PhD degrees at competent registries should be able to do it even easier.

I find myself a little bit critical from measuring input as a success. I'm not criticizing the initiative but I find myself a little bit at issue with how not just the engagement team but how anybody measures success of something like this. You don't want to go and attend a workshop. You don't want to get a certificate. You want to figure out how this works.

You want to implement a system that is easy to maintain, that is secure, and that is not overly complicated. For example, using OpenDNSSEC, when you don't have an HSM, is something that I have never understood. It's extremely complicated. And big ccTLDs with millions of domain names, they assign their names with BIND or other inline

resolvers or servers. So you don't really need to use a complicated setup when you're starting out.

In any case, there was another one from the floor. If there is anything from the remote, Kathy, please make sure that I know about it.

KATHY SCHNITT: We do have one remote.

EBERHARD LISSE: Then let's take the one from the remote first.

KATHY SCHNITT: Thank you. This is from Gaurav Kansal. "Do you have content in Hindi language also?"

CHAMPIKA WIJAYATUNGA: At the moment, we do not have in Hindi. So again, this is something that we will have to look into, especially work with some collaborative partners and see if we can produce that. At the moment, we do not have.

EBERHARD LISSE: I also have a very strong opinion about that. I think if you're doing this kind of work, DNS or DNSSEC, English is an absolute requirement. The RFCs are published in English. So if you're doing things on which a TLD depends upon ... I don't read it in German, which is my native language. I read it in English.

So while introductory material and material for your clients, it's good to have it in your native language, the implementers really need to focus on being able to understand enough technical English to be able to communicate with people who can help them. There is many situations on the same levels that have done this. And if you have questions, we are all a friendly audience here. There is lots of contacts that you can make at this Tech Day and other meetings that can help.

There was a question from the floor. Please identify yourself to the record.

NICOLÁS ANTONIELLO:

Nicolas Antoniello from ICANN Org. Or actually, I'm part of Champika's OCTO Team. Just to add to a side comment on some of the questions that have arisen, I wanted to say that one of the rationales—besides those already mentioned by Champ, one of the rationales behind the design of this kind of lab was to provide the community a means to practice in a real environment—in an environment that is as close to real life as possible. In theory, everything almost always works.

But then, when you go down to having an environment that you already have a router, already have a server, already have different network topologies, when you can try configuring your recursives, configuring your authoritatives, having a hidden authoritative, and configuring your zone transfer policies, and then trying out stuff like configuring routing in those authoritatives to use Anycast or in the recursives, we think it's a great tool for people. They just have to go and test it.

Actually, one of the things that Champ didn't mention is that once we build this virtual topology, the servers—the authoritative server for the lab, for instance—is a real server. It's a real DNS server with a real IP you can carry from anywhere in the world. So the only servers that are not 100% accessible from the world are the ones that are inside the network because they have private addresses. But from that, the whole thing is real stuff.

So one way of measuring success is because of the feedback we have after that and because of the people trying to do the things. In many cases, for example, we have programs with universities. We then evaluate them. So we make sure they complete all the training and they were able to configure everything according to the common best practices or whatever they are trying. Thanks.

EBERHARD LISSE:

Make no mistake. I am a big fan of this. It's like build it and they will come. It's a good opportunity that people become aware that there is such a situation where you can start playing with this without affecting your infrastructure. And there is a large number of small ccTLDs or TLDs, but especially ccTLDs in developing countries, that don't really know how to do these things. So having access to a playing ground or to an environment where a lot of things can go wrong but nothing can get damaged is actually a good idea. There is Yazid over there.

YAZID AKANHO:

Thank you, Dr. Lisse. I'm Yazid from Technical Engagement Team as well. I would like also to add a few comments to the presentation that Champika did and also to Nicolas' comments.

The first point is about how do we measure success? Indeed, it's something. Even internally, it's hard to answer. Let me give you a short example. Sometimes when we are delivering the trainings, even the participants themselves realize that they are not, maybe, following some of the best practices we promote during the training and then take the engagement to actually try to follow the best practices.

We have done trainings in areas where people do not even have a clear documentation of their infrastructure. And when we discuss with them, we show them the potential impact of not having a documentation. They themselves realize that it's something extremely important. Maybe one or two months after, they share with us a draft and they ask for input or for comments.

Those are things, for us, which are real success, regardless the number of people. Sometimes we do the trainings, maybe with two or three people. But for us, the number is not the only one success criteria.

Another, I would say, success criteria is indeed the language we use. Of course, I fully agree with you. English is key if we want to grow in this industry. But sometimes in the past, we have seen our trainers coming, only speaking English language while the participants maybe would like ... They agree to have the content or the system in English but they expect the trainer himself to speak in their own language, in their local language.

This is also something for which we receive positive feedback for the Technical Engagement to have been able to, at least now, deliver the training in Spanish or in French. And we look forward to also be able to deliver in other languages.

The last one, also we try to show the participants the diversity in terms of solution, of software which exists. We don't directly go to OpenDNSSEC, for example. We start, indeed, with the DNS software, whether it is BIND or NSD and then we move to OpenDNSSEC, just to show them the diversity which exists in the ecosystem. Thank you very much.

EBERHARD LISSE:

Thank you very much. Even more than capacity building, I like to do this from the bottom up. And that seems to be the case. I know Yazid was a Fellow at one of our previous meetings where he gave a presentation about Raspberry Pis being used, which I still recall relatively fondly. And it's great to see that this gets exported into the wild, so to say. All right. Thank you very much. We'll see each other in half an hour. Let's give them a big hand.

CHAMPIKA WIJAYATUNGA:

Thank you very much and hope you have a good time.

RAJIV KUMAR:

Thank you. Thank you very much.

[END OF TRANSCRIPTION]