ICANN74 | Policy Forum – Tech Day (3 of 3)
Monday, June 13, 2022 – 16:30 to 17:30 AMS

KATHY SCHNITT: The session will now begin. This is Tech Day Part 3. Go ahead, Eberhard.

EBERHARD LISSE: Thank you. Again, as before, please report if you want to ask questions via the chat or by raising your hand or by trying to unmute yourself. We will call on you. Normally, I would admonish the people in the room to get sat down, but since we have got a large number of remote participants, I am giving the floor to Brett Carr from Nominet. Please go ahead.

BRETT CARR: Good afternoon, everybody. I'm very grateful to get to speak to you this afternoon, mainly because I get to take this mask off for 10 minutes, which is nice.

So this afternoon, I'm going to talk to you about transitioning DNS for gTLDs between providers. I'm going to run through an introduction to Nominet's DNS team and what we do, talk about Nominet's involvement with gTLDs, and then go through the

various stages of transition, and then finish off with a couple of other things that we've got on the plate in this area as well.

Nominet DNS. We're a team of six and support the following services. We run both UK on gTLD authoritative DNS from 12 global sites. We run .uk and several second-level domains and 67 gTLDs currently.

We also run something called the UK Protective DNS Service. This is a recursive service that uses RPZ for filtering, which is used by a large volume of the public sector in the UK. And this is from four sites, all in the UK.

We also run something called the UK PSN DNS Service. The PSN stands for the Public Services Network, which is an internal public sector private network in the UK.

And then, obviously, we run Nominet's corporate DNS services, so internal resolvers, authoritative for Nominet.uk and many, many other corporate domains that marketing seems to invent every 10 minutes.

And then more and more these days, we get involved in DNS engagement protocols, policies, operational practice, and things like that, which is one of the reasons I'm here.

So currently, Nominet is the registry provider for two gTLDs—.wales and .cymru—and we are the backend provider for 67 gTLDs, most of which have been transitioned in at some stage

from other providers rather than delegated directly from ICANN in the first round.

We're also what's called a EBERO provider. EBERO is Emergency Back-End Registry Operator. So, we run on standby all of the time for ICANN with two other organizations to provide DNS and registry services at short notice. So we have to be able to onboard the gTLD within four hours, 24/7/365.

So I want to run through the several stages of gTLDs. And I want to shift over to the left a little bit because I can't read my own slides, which is a problem.

UNIDENTIFIED MALE: [inaudible].

BRETT CARR: No, no. That's fine. I can see it now. So I'm going to run through the various different stages of gTLD transitions. This usually starts with us being notified by a registry colleague of a new gTLD or our registry team colleagues are looking for business where we can run gTLDs through other organizations.

And when they run that business, we get a notification from them that we've got a new contract. We then do some assessment of the size and impact of that TLD or those groups of TLDs to decide where we're going to play some of them in our platform.

99% of the time, they're placed in a standard way in our platform because 99% of gTLDs are small with a small amount of delegations and a small amount of traffic. But now and again, we get something out of the ordinary or we get a large group of gTLDs in one go from one provider. So we have to do that assessment, basically.

So when we've made that assessment and then we made a decision whereabouts in the infrastructure and how we're going to deploy them—which, most of the time, is a standard way—we build some test infrastructure and some production infrastructure which looks identical.

And then every stage of that build is documented and is defined as code in Ansible. Ansible, for those who don't know, is a configuration management tool for defining infrastructure as code. We use Ansible extensively to automate some of this stuff.

When we get to this stage, the build is made up of a database schema, some provisioning infrastructure for the TLD, some DNSSEC infrastructure. And a lot of this stuff is shared, so it's not like we build a new one for every TLD. But it's obviously new configuration on existing infrastructure sometimes.

So it's a provision in DNSSEC distribution. The configuration on our two different types of DNS sites being global and cloud based. And then we create the initial zone files which, at this stage, are

dummies zone files with some records in it that are required for testing purposes.

We add some monitoring so that we know during this subsequent phase if anything's gone wrong. The level of monitoring at this stage is different to the level of monitoring when we're actually in production.

Then when we've got all of that infrastructure up and running, we hand it over to our QA colleagues to make sure that we haven't ... A lot of this stuff is automated, but there's still some human elements. So whenever you involve humans, there are always mistakes. Right? So hopefully, QA will back again and say everything's fine.

We then work with our customer and ICANN to schedule something called the RST. For those who don't know, RST is Registry Systems Testing. And this is done by engineers at ICANN to make sure that the registry systems and DNS are fit for purpose, basically, to make sure that we're ready for live production. For those of you who've been around a long time, this was previously known as pre-delegation testing or PDT.

The ICANN people test that DNS servers are responding properly, the zones are properly DNSSEC signed, and that EPP, WHOIS, RDAP, etc., are all working properly; and where the registry agreement has references to IDNs, that the IDNs are set up properly.

If they find any issues at this stage, they raise a ticket. We have to fix this pretty quickly. RST comes with a charge, which the registry operator has to pay for. And if you fail RST, it has to be rescheduled in a timescale that suits ICANN and paid for again.

But if you fix the problems as they find them quickly, they don't fail RST for you. So it's in our interest to fix problems ... I mean, we've been doing this all quite a while now, so we don't generally have any problems. But not every gTLD is the same, so sometimes things crop up.

RST usually takes about two to three weeks at this stage, depending on the complexity.

We then move into Phase 3, which is the actual DNS/DNSSEC transition stage. At this stage, the current state is that the incumbent is producing and distributing the zones onto their own infrastructure, and their infrastructure is serving the DNS. Nominet are producing and distributing a dummy zone onto our infrastructure.

The change that happens at this stage is that Nominet works with the incumbent provider to ingest the current zones from the incumbent. So we pull a copy of those zones in on a regular basis. They send [inaudible] to an ingestion server.

We have some code that we've written that unsigns those zones or removes any DNSSEC information from the incoming zone and

resigns it with our newly created DNSSEC keys, but either adds or doesn't remove in the first place the existing public keys.

That, again, gets distributed to our new infrastructure. At this stage we're now distributing a zone that has got our DNSSEC information in it. So it's signed by our keys, but with the old key still present for post-publishing purposes.

We continue in Phase 3 by generating what we call several NS Change files, numbered 1-6. These have clear dates and instructions to the incumbent provider. Again, we've got code that produces these files for us. So we run that code based on giving it some inputs, and it produces these files. And I'll quickly go through what's contained in each of the files.

So in the first one, it has instructions on adding the Nominet DNSSEC keys to the existing zone in order for them to be pre published. It also add the DS to the TLD for nic.tld.

And then in Stage 1, they add the DS records to the root for the new keys. In Stage 3 we ask them to add the NS records for the first half of the new name servers and remove the NS entries for the first half of the new NSes. So in this stage, half of the NSes are removed and half of the NSes are added. And at this stage, the Nominet name servers become live, so we start getting traffic. Traffic for DNS is now split between the old provider and the new provider.

In Stage 4 we add the NS entries for half of the new entries to the root using the Root Zone Management platform, and we remove the first half of the old DS records.

And then in Stage 5 and 6, we repeat this effort again, adding and removing the second half of the NS record and DS records.

A final step at this stage is that we often have to change the tech and admin contacts for the TLD in the Root Zone Management portal. This is not always the same because some registry providers have different requirements [to all of those], but the majority of the TLD stays as admin contact and Nominet become the tech contact.

So now at Phase 3, the current state is that the Nominet name servers are now in the zone and the root, and so are fully serving all of the traffic. DNSSEC has been signed by us, but we're post publishing the old keys. But the zone is still being provisioned by the incumbent provider, so we're still pulling the zone in from them—unsigning it, re-signing it, and publishing it. So at this stage, the operations of the TLD are dependent on both providers.

Now I have to say, there's probably people in this room or remotely who we've worked with on this in the past, and I think kudos goes to the way people operate in this area because my experience over the years is that even though people have lost business—so we've taken business from them ...

And let me be clear. This happens the other way around. People take business from us as well, so it works both ways. People work in a very professional manner, and this has worked extremely well over the years. So we don't get people being awkward because we've taken business from them, which is really good. Obviously, it's important for the stability of the TLD at this stage that we work together.

And then the final—not quite the final part of the jigsaw—we do the Registry Transition. So this is where the registry itself, the database itself, is transitioned from the incumbent provider to Nominet. So we work with the incumbent providers to set a date for this, and we—well, "they" probably at this point—publish downtime. This has to be notified to ICANN for SLA purposes, but also to any of the registrars so that they know that during that downtime period they can't register and change domain names.

On the day of the Registry Transition, the incumbent makes their registry read-only and then produces a database escrow file. For those who don't know already, database escrow files are produced every day anywhere by all gTLD providers and deposited in an escrow provider so if the backend provider disappears overnight, somebody can resurrect the TLD without too much hassle. This is a standard process, anyway.

But they produce a specific database escrow file during the downtime period and then supply that to Nominet. Nominet then

populates our database, which you'll remember we created back in Phase 1 with dummy data. We replace that dummy data with the data coming in from that incumbent provider. And then we generate a new copy of the zone file. That zone file then goes through various checks. And if that everything checks out all right, it's loaded into our provisioning system.

And then the DNSSEC and Distribution infrastructure is switched over from using the incoming zone from the incumbent to using the new zone that's coming for our database. And then at stage, the registry has [completed the transition].

Now, I'm skipping over lots of stuff that's happening from WHOIS and EPP and RDAP, because this presentation is focused on DNS rather than registry services. But our other colleagues do lots of work to make sure that there's as little downtime as possible to those registration services as well.

And then we move into a cleanup, Phase 5, when the requisite time period has happened. We remove the old DNSSEC keys from the zone that we're post publishing that belonged to the incumbent provider. And we remove the old DNS records from the parents as well.

And when we're sure everything's running fine, we also remove the incumbent zone from our infrastructure so we're not pulling that zone [from them] anymore. And then we're pretty much done.

There's just a couple of other points at the end of the presentation that might be interesting. As I mentioned earlier, we do transition outs as well. Unfortunately, we do lose business from time to time. I'm happy to say we've done less transition outs than we have ins, which is good news. Our position with transition outs is that it's out opinion that the transition should be led by the gaining provider. And so when we have meetings with the provider, we make this clear to them that it's their position to drive the transition.

However, we encourage but don't mandate the use of a similar process to the one we use because we've honed this process over quite a few years now. We find it works. And therefore, it's often useful for other people to use it as well. But again, we let them drive the process. So if they wanted, they could do something slightly different. We'll go with that [inaudible], but we take care to point anything out that we think might cause a problem.

We also do something slightly different where we're doing a transition in for ... The majority of gTLDs run their DNS servers within nic.tld. But some people run them out of bailiwick, so they might run them in their own domain name, for instance. And in those cases, the domain name where the name servers are doesn't have to be DNSSEC signed.

And that causes a slight problem with our process, so we do something slightly different where the name servers are all out of

bailiwick. But it's really just an additional step, basically, that we follow.

As regards to future plans, as I mentioned there's a lot of infrastructure code and automation here. But there's still an awful lot of manual steps as well. And a lot of that is because, again, like I mentioned earlier not every gTLD is the same. And so things can be different on each gTLD. But we do have plans to make this a lot more automated than it is and to iron out some of the bits that involve human beings.

And in particular, we will definitely be looking at a lot of automation if we ever get to round two of gTLDs. Obviously, whenever you do a gTLD from scratch that's delegated from scratch via ICANN, that's a not so easily repeatable process. So we'll be looking to add a lot of automation there. But automation is something that we're very keen on doing. And we've done a lot of that already, but we intend to do more.

And I've missed the requisite question slide at the end, but if anybody's got any questions, I'm happy to answer them.

EBERHARD LISSE: Thank you very much. We are not strapped for time, so please come up with questions. I like handbooks. Yeah. I even saw a little. Clearly we did a handbook. We have a handbook. And when we, a few or a number of Tech Days ago, reported on how to

switch keys, how to switch infrastructure, and how to switch name servers in one rush, we developed a [inaudible] engineering plan.

I did [inaudible] presentation so that I could always look at the [inaudible]. And we very carefully designed each step and the outcome, and did the same thing. That's the only way. DNSSEC is complicated, but easy. Yeah? When you do it right, it all falls into place.

Recently, we learned of an island in the Pacific who forgot to change the DS records in the root, and it was a bit of a problem because they [phoned][inaudible], but not a 24/7 number for IANA to help them. So having a book with each step and the expected outcomes to compare to that is very helpful.

BRETT CARR: Yeah. Completely agree, Eberhard. One of the things I was very encouraged about this morning was Kim's mention that the RZM has got some API facilities coming soon because a lot of the manual steps that we do in our stuff for the moment is interactive with the RZM. And if we can do that in an automated fashion, it'll get rid of a lot of manual steps.

EBERHARD LISSE: Yeah. For us, a small ccTLD, occasional contact RZM is working nicely the way it is. But if you do this on a regular basis, you want

something like [inaudible] or whatever that is well described that you can adapt your systems to that you push a button and it automatically goes to the other [inaudible]. And you get to [inaudible] so that you don't have to, other than monitoring, sit there.

I don't see any questions. I don't see any hands. I don't see any unmuted names. So I am releasing you. Thank you very much. It was a very interesting presentation. I hope we will meet in person next time.

And now, without further ado, Jordi Iparraguirre. [inaudible]? Just let me see. Is he on?

UNIDENTIFIED MALE: Yep.

EBERHARD LISSE: Oh, there you are. Okay, can you please share his screen or share his presentation, please?

JORDI IPARRAGUIRRE: Okay, here we go.

EBERHARD LISSE: There you go. We can see it and I can hear you. Go ahead. You have the floor, Jordi.

JORDI IPARRAGUIRRE: Okay. Thank you very much. Hello, everyone. Well, as we are getting to the end of the day, more than a presentation that's going to be a story. It's a story that we've been working together with DNS Belgium (.be), DK Hostmaster (.dk), and EURid (.eu). It's a presentation that was shown a couple of weeks ago at CENTR meeting. So if you were already there, it's basically the same.

The idea is to share it here to inform you about where we are going and what we do, and also to collect feedback from you about ... Maybe you are doing similar things in your region. Maybe you heard about other projects similar to that one. Or maybe, well, you just want to provide some input. That would be also very, very interesting.

So basically, a couple of years ago we started the project and endeavor that, well, had a kind of doubtful outcome that was to try to find out if it could be possible to, in the framework of GDPR, share personal identifiable information to try to find abusive domains on different TLDs.

So, why? Why that need? Basically because one of our colleagues received a phishing e-mail that was asking that person to enter bank data to verify his own identity. Here you have an SMS that that very same person also received. But instead of being an SMS, that was an e-mail.

So there was a campaign of spam and phishing. And said, “Okay, it's interesting.” Because that domain that was sent, the one that person received that you have here, basically is Dutch for ING—which is a Dutch bank or European bank—payment request. That very same domain appeared on .edu three and then four days later, basically the same with a typo there to catch people.

So if it were just focusing on the registration data, that was perfectly valid at the first sight. A valid name, a valid address. Nothing suspicious. Except, of course, the domain name itself.

So we had the risk of missing that domain name and then letting that get into the route and harm people. So, well, hopefully we are looking for these kinds of domains. But, well, nevertheless we’re aware [inaudible] interested in knowing about, do you registry-whatever have received, have you seen strange domains that could be interesting to share with the community?

So then, .be, .dk, and .eu entered in a challenging trip beyond the explored limits of GDPR because we knew that we really had to take care of that. We cannot [believe] that. So nevertheless, we wanted to explore the limits.

So we started to try to manage through different ways and through approaches to help each other trying to fight reducing abuse. Or if you want, instead of attacking abuses, okay let's only take care about the WHOIS quality. Is that WHOIS quality that

we're under, the WHOIS data we are managing, is that good enough? So let's try if we can help each other.

In case you have any doubt, we are Gandalf. We are not the warlock. So, trying to prevent the abusive registrations to pass.

So besides that, we also had in mind a kind of, okay, we have to move. Because in our legislation—in our case, the European Commission and the others' respective countries—we have some pressure to act against abuse. So we'd better act before we are told what to do.

[inaudible] will sell because we know the business. We know the industry. And we know what can be done much better than something that wasn't designed.

So between Ljubljana in February 2020 and Jamboree or ICANN today, we had to overcome different problems, basically while understanding how the GDPR works, the legal framework of all of that, the infrastructure, the bid of .eu., the COVID problems, and so. But finally, we managed. Right?

And then, as mentioned, we started DNS Belgium. Maarten, at the bottom, is the member that's representing DNS Belgium. Dk Hostmaster. [Erwin] and myself started to move that forward. We are presenting that here at ICANN. And we did it at CENTR. And, well, in some months we will try to explain how it went. Now we are almost finished.

Basically, we discussed and we defined the goals and the scope of the program. We studied how GDPR was impacting what we wanted to do—data sharing, retention, removing data, etc., and so on, who was the owner and processor.

We tried to define a very easy and decentralized data sharing infrastructure so that it is not a centralized system that someone has to take care of and maintain. Each registry participating will have its own, and it's very easy. We've made it very easy to maintain and to install.

And we're really ready right now to start to share data. As soon as we start, we will keep some metrics on that. We will study the results. And that's why, in some months, we will report back.

So, how does it work? Well, the idea is that participant registries will do that in a totally voluntary basis. As I mentioned, it's a decentralized system. It's basically taking advantage of the collective intelligence of each one of the registries by sharing data.

The guidelines. Again, GDPR is there. So, privacy by design. Privacy by default.

And then [two kinds of rules] that we tried to [inaudible] the registry that participates feel really, really comfortable on that. Basically, each registry offers what it thinks may be useful to other registries to find abusers.

And then each registry takes whatever they think that the other side is offering is really beneficial for them. So if you are not interested, you just delete it. No problem. If you don't want to pick up the files, you don’t pick up the files.

So if you want a little bit of paraphrasing and changing the sentence on the GDPR from Postel, well, “Be cautious on what you share because of GDPR, but be conservative in what you accept from others.”

We also took into account different things about security, encryption, accounts, logs, etc. And then, well, as mentioned also, you are not obliged to take data. You just share what you want and you take whatever you think is important for you. And the rest, you remove it. And then, nevertheless, you are forced to remove data from your systems at least once the data is one year old.

And then on that journey, we moved a little bit from abuse to data accuracy. Okay? As our goal is really to detect abuse as early as possible on the life of a domain name, at the very beginning, in the first hours, in the first days, when that domain is not yet active or has not been detected by experts, it's very difficult to say, “Yeah, this is this kind of abuse. This is phishing. This is spam. This is malware. This is botnet.” But on the other side for us as registries, it is much more easier to say, “Yeah, that’s an abusive registration.” Why? Well, maybe because of the domain name

itself. Maybe because the registrant data does not make sense. There are inconsistencies, whatever. The connection data, and so on. So we move on that path.

Because if you detect abuse based on content—usually counterfeit, pharma, phishing, etc.—okay, that's fine. But it's already there. You are late. So that's why we want to really detect them as soon as possible.

And the abuse that's not visible—botnet, spam, and malware—it's very difficult for us as a registry to classify. So it's not about the issue. It's about the registration.

And then, in finalizing right now, as we are focusing on the registration data accuracy and risk prevention, there is an effect. The primary one is that we fulfill our mission to have a very good, or as valid as possible, WHOIS data. At the end of the day, the registries [are for that].

As a secondary effect, as we are trying to ensure that the registrant is really who it says it is, we will have less abuse. Nevertheless, it is a delayed effect there. Right? As we are putting more barriers to the abusers, well the ones that really want to use that domain, that TLD, to abuse will adapt with new strategies.

So that's going to be a never-ending race, a never-ending marathon between the abusers and the registries trying to

identify malicious registration. So that's why we think that it's worth to cooperate there.

Then how does it work? And I'm going just to explain how are we going to implement it in .eu. But while discussing with .dk and .be, the idea is basically going to be the same. So it would work for any registry there.

The idea is, as I mentioned, we want to detect the abusers as soon as possible in the life of the domain name. So, try to detect suspicious registrations before the delegation to be able to delay them and stop them.

So another point would be the post-delegation checks. We know that we're not going to be able to really catch all of those at the pre-delegation. So nevertheless, we are going to crawl and analyze our zone to find out if there is something strange there. So we are going to have the domain name. We're going to analyze registry data, crawling, whatever. And then all that data science, things you've been seeing in previous presentations and so on.

We're also going to have human review of the reports, just to have a human double check the decisions of the computers. And also to learn. To learn because the behaviors of the abusers change and we have to adapt to that.

And finally, well, the registries are going to put up some WHOIS Quality processes to verify the real identity of the registrant.

So after that, the point is to share data with partners, what we call alleged suspicious domain names or weird registrations. Something that really looks strange by the domain name or the registration data itself.

And today, for instance, .eu is just sharing domain names only. The domain names themselves, but also domain names on the name servers, the mail servers, redirection, and others. Right?

So here's the picture, a little bit, how it works. This the .eu. A little bit flow of the whole process. APEWS is the first one, the pre-delegation one. Then there's a decision. That pre-delegation system thinks, "Is that correct so it can go to the zone? Or are we going to ask an identity check?" If the identity check fails, then the domain is going to be suspended because that person could not verify their identity.

Nevertheless, in parallel, we inform third parties—cybercrime companies, law enforcement, whatever—expecting them to really act on those domains that we deem as suspicious/strange so they can investigate on one side and also block them at other levels. While we do the identity check, they can be blocked at the DNS level, at firewall level, at the browser level, or whatever.

So the question amongst us was, okay, which other data points beyond the domain names could be relevant? Can we share them with other partners, too? And again, how the GDPR applies here.

Well, we came out that it seems that we can really share all of the information, e-mail provider and username, for instance, the registrar/registration hour, and all the domains involved in name servers or IP addresses, mail servers, autonomous system numbers, and so on.

Hopefully, well, the idea is that we are going to be able to feed our detection systems at the pre-delegation and at post-delegation based on information shared by other registries. So if the domain or an e-mail of a registrar, whatever, appeared in TLD #1, if this is shared, then TLD #2 can feed that into the detection systems and detect that issue. If possible, a pre-delegation. Otherwise, a post-delegation.

This is feasible because, well, we had also all of our lawyers working on that. All of the three different legal departments were going through the GDPR. We're getting, also, external counsel to see if that could be feasible. And in the framework of GDPR and with the argument that we are going to provide—what's the word—public benefit. Right? Stopping abuse. So, yes, we are covered and we can share this kind of information.

So finally, the next step. Well, setting up the infrastructure that has been already defined and just a matter of doing it. And it's pretty easy. Start exchanges. Measure the impact of that. Is that really worth? We think it is, but we really want to have some metrics there. We know that starting just three TLDs is going not

to be a lot of data there, but we know that the more will be, the more interesting data points we're going to collect.

Then with that experience, propose and implement improvement. And then, of course, welcome other ccTLDs that also abide to the GDPR in Europe to work on that.

So, that's all on my side. Thank you very much. And if you have questions, we will answer that. Thanks.

EBERHARD LISSE: Thank you very much. Very interesting. There's one question in the chat from Jacques Latour from .ca. “Once the registration information is deemed false or fake, failed ID check, can you share the entire details of the fake registrant information? Name, email, address, etc...”

JORDI IPARRAGUIRRE: Sure, I can take that. Quick answer is yes. And that's what the lawyers worked on for some time to find out that what we are doing as ccTLD registries and preventing abuse, or at least having a clean database, is in the public interest, which is a clear reason for doing this. Also, sharing is in the public interest.

So this is not based on consent. This is not based on changing [inaudible] service or anything else. We do this in the public

interest. And that's it. We might have to update some privacy policies on our website, but that's about it.

EBERHARD LISSE: Then Cristian Hesselman mentioned that they have a similar initiative in the Netherlands to share information about DDoS attacks. So you might be able to learn from each other.

Mark Elkins asked, "So should the EPP system be extended to include an 'Identity Number' as a standard for the registrant?"

Did you get that?

JORDI IPARRAGUIRRE: Yeah. This is something that also somehow appeared in the CENTR meeting a couple of weeks ago. That is, yeah, we can detect abuses and we can share data about potential abusive domain names. What about on the other side? The registrants that have been verified with a proper digital identification on any other system? Can we share information about those so if they have a dot-whatever, we do not bother them again if they are going to register domains in another TLD?

The idea was, yes, it would be nice. The problem is, as was mentioned, which ID are we going to use to really identify that very same person or that very same registrant into different registries? Are we going to use the ID number? The national

identification number, for instance? All the countries can use the same one.

So we have not yet explored that. We've not yet thought into that. It's something that's on the radar, but we've not yet gotten to address it because we want to address first on abuse and then think about how could we expand that into verified registrants.

[ERWIN LANSING]: I can add some more to that, Eberhard.

EBERHARD LISSE: Go ahead.

[ERWIN LANSING]: There are several other projects going on. One of them is also looking into once a registrant is verified in one registry, can that ID then be shared with other registries so the registrant does not have to we reidentify for another registration in another registry? That is not this project. I know some other people are looking into that. That would be really interesting.

Another thing here, when we're talking about EPP extensions, that's, again, another project we're looking into to make these processes a little bit more similar for the registrant, and especially the registrar. Because right now we're all doing it slightly different. And there, also, EPP might need to be extended so we

can share the status of the ID checking while we do ID checking so to register can follow along and also inform the registrant of what [inaudible].

EBERHARD LISSE: And while one person from ... Kristof said a hash would be an option because the ID number can't be used. That is not ... The problem is that in some countries, there are no ID numbers. Even in Europe it is all different. The Germans don't have an ID number and are very strict for 50 years. They are not going to have ID numbers in Germany. In the Netherlands, you have the Burgerservicenummer. In other countries you have [inaudible]. That's very difficult.

The next question that I have here is from Joel Karubiu. "Are there any possible partnerships with ccTLDs outside Europe that have mirrored GDPR guidelines?"

I would assume that's a matter of trust and agreement.

[ERWIN LANSING]: Yes. So basically, it's not out of the question. But not at this stage of the project. We're not there yet. But I'm sure we would be open for that at a later stage if this turns out to be successful.

EBERHARD LISSE: I read now, “Could the European ccTLDs (especially) connect laterally across the border to share these designs for good practices?”

In other words, is this open source?

[ERWIN LANSING]: Well, there is no code, but there's this presentation. So, yes.

EBERHARD LISSE: [inaudible]. But the point is, of course, that the more you divulge how you do things, the easier it becomes to circumvent it. Ongoing problem.

JORDI IPARRAGUIRRE: Eberhard, just a point. The complexity here is as much technical as legal. Technically, we easily agreed on a kind of data format and the process to share data, which is really easy at this stage because it’s also a prototype. The complexity came from the legal side. Okay? Finding the reason why that could be feasible, the kind of data could be shared, the framework of the GDPR-abiding registries and countries.

So that's what, initially, is really having us limited to explore how it goes. But in terms of code, it's just sharing data in the format you want. So, no. No complexity there.

EBERHARD LISSE: Okay. Jacques asked another question. “The Secure Domain Foundation was established to enable information sharing between ccTLDs but did not get traction. Are we looking at building something similar?”

[ERWIN LANSING]: Yeah. I remember the Secure Domain Foundation. [We were at talks] with them a couple years ago, Jacques. I’m unsure about the details. I think there will be overlap. I think there also will be differences. I’m not quite sure exactly what they were doing in the same sense. Yeah.

JACQUES LATOUR: Well, they spent a lot of time and effort to build a system where a ccTLD could submit data and then query data like ... Let's say you say this e-mail addresses bad and then you can put it in there to say it was associated to your domain. And then you can query in there to find what other ccTLDs had abuse with that e-mail address.

So if we agree that fake information or false information that fails, once it fails the ID check, if we can share all of our information in the same bucket, that becomes useful. I’m not saying we should ...

Maybe we should talk to them to see if it's an option. But they built something to facilitate this already.

EBERHARD LISSE: It's an interesting little discussion going on in the chat which I'm not going to paraphrase because there was no real question to the panel, as it were.

I'm not seeing any hands on the participant pod. That leads me to thank you guys very much. Interesting presentation.

Anybody wanting to communicate with them, the e-mail address, as I've said a few times, is listed in the agenda and clickable. So you're more than welcome, I'm sure, to communicate with Jordi who is spearheading this today

Of course, you have a standing invitation on one of our next Tech Days to report back on the outcome of that.

Let me just see. There is no more question. Thank you very much.

Before I asked Cristian Hesselman to wrap this up, that leads me to thank the ICANN staff—the technical staff and our ccNSO secretariat for their Zoom magic. I'm amazed how good this works.

I actually think this is a format that we should keep on doing. Maybe having two rooms is very helpful because we had, at the

most, 126 ICANN, if I recall correctly. We have now 109 participants listening. And the idea is to increase the reach.

So this format, I think, works quite well. So while we can make our distance a bit smaller socially, I think this works.

Cristian, you have the floor.

CRISTIAN HESSELMAN: I have no slides. Thank you, Eberhard. So Eberhard asked me to do a brief wrap-up of this session.

So my observation, personally, is that I think this is an excellent overview that we saw today of the recent developments in the DNS ecosystem, and then very often from ccTLDs.

So the five major topics that I identify today, in addition to the host presentation by [inaudible] Müller, were anti-abuse work, so to speak—anti-abuse activities of the various ccTLDs. We saw DNS infrastructure/engineering, DNS measurements, and DNS-tools. So I think that these were the four big chunks that we saw today.

So in terms of DNS-tools, we saw ZONEMD which Hugo discussed earlier on today about signing the zone files, let's say, in a situation where it's at rest, so to speak, rather than in transit. So this was something for ccTLD operators to use internally on their systems.

And we also saw various updates of Zonemaster presented by Mats, including support for CDS/CDNSKEY testing as well as additional languages that the tool now supports. And he also emphasized that Zonemaster is basically—at least that's what I learned—it's basically a flexible component that you can use for various purposes.

And I kind of took drew the parallel with internet.nl, which is a site that we have here in the Netherlands where you can check the validity or the security of your domain name or of your Internet connection, which also is based on a generic component. And I think that Zonemaster is actually quite similar, but then specifically aimed at name servers for DNS infrastructure.

So then the other group of talks that we saw were on DNS measurements. So the first one was by [inaudible] about the DNSSEC bug that SIDN labs found in Google Public DNS. It was kind of a difficult one to carry out, I think, but at least the folks at Google managed to repair it within a month and a half, and the bug bounty was donated to an open source project.

And then we saw the work by Roy Arends who did scans for DNS resolves to get the root [inaudible] from them to check which root servers these resolves were using. And he also emphasized that there was a lot of future work to be done as well over there.

Then the other group was basically what I call DNS infrastructure/engineering. That was basically Brett on his own,

showing what it takes to transfer a gTLD to Nominet or perhaps also even outbound, so to speak. He presented a detailed cookbook or manual of five phases to make that happen, which I'm sure looks easy on slides but is kind of difficult to carry out.

And then lastly, we had a bigger chunk of talks on anti-abuse projects. So we had two presentations about detecting abuse based on machine learning models, for example. And I think that this is something that's key to further secure the DNS ecosystem.

But what lacks there is basically what Jordi just spoke about—the data sharing components—because, as we all know, the Internet is a collaborative system, and so is the DNS. So if you want to secure the entire infrastructure, you will need to collaborate. Right?

And the one of the ways of doing that is sharing information or, alternatively, sharing the models that you developed for these machine learning systems. So that's something that we're actively engaged in ourselves at SIDN labs, together with our colleagues at DNS.be.

And a recent development for the data sharing might also be the Amsterdam Data Exchange. I'm not sure if you guys are familiar with it. They also tried to do something similar, but on a generic basis. So not specifically for cybersecurity, but for any type of information that you want to share. So that might be something you may want to look into as well.

And also, maybe we can chat out later on, on the DDoS Clearing House because I think we already did quite a bit of the work there, although the application is different. And I totally agree that the complexity is in the legal bits and not so much in the technology.

So these last three, this last group of presentations about anti-abuse work is actually going to be an important one, especially with the NIS 2 Directive coming up here in Europe.

So I think, at least in my mind, the conclusion is that we looked at four great topics today was excellent presentations which, again, I think gave an excellent overview of recent developments in our community. And I think, on behalf of Eberhard, I can thank all of the speakers and Eberhard himself for this great afternoon.

Eberhard, back over to you.

EBERHARD LISSE: Thank you very much. I appreciate this. And we'll see each other in Kuala Lumpur. I will try my best to make sure that I am not going to miss it this time. Have a nice evening. I am already missing the Indonesian food evening that I was going to have tonight. Good-bye.

[END OF TRANSCRIPTION]