ICANN71 | Virtual Policy Forum – DNSSEC and Security Workshop (2 of 2)
Monday, June 14, 2021 – 10:30 to 12:00 CEST

KATHY SCHNITT: Hello and welcome to the DNSSEC and security workshop part two. My name is Kathy and I'm joined with my colleauges, Kimberly and Andrew, and we are the remote participation managers for this session.

Please note that this session is being recorded and follows the ICANN expected standards of behavior. During the session, questions or comments will only be read aloud if submitted within the Q&A pod. We will read them aloud during the time set by the chair or moderator of the session.

If you would like to ask your question or make your comment verbally, please raise your hand. When called upon, you'll be given permission to unmute your microphone. Kindly unmute your microphone and speak at that time.

All participants in the session may make comments in the chat. Please use the dropdown menu in the chat pod and select “Respond to all panelists and attendees.” This will allow everyone to view your comment. Please note that private chats are only possible among panelists in the Zoom webinar format. Any message sent by a panelist or a standard attendee to another standard attendee will also be seen by the session host, cohost and other panelist. And with that, I'm happy to hand the floor over to Mr. Steve Crocker. Sir, go ahead.

STEVE CROCKER: Thank you very much, Kathy. Welcome to part two of the DNSSEC workshop. In keeping with the fact that this ICANN meeting is the shortest of the three that are scheduled each year, this DNSSEC workshop is also shorter than usual, and correspondingly, this panel is shorter than usual. This panel is a semipermanent part of the DNSSEC workshops. The focus is on DNSSEC updates and multi-signer coordination. These are two of the sort of rough edges or loose ends, if you will, of the DNSSEC protocol area. This is episode five and I intend to, along with my cohost, Shumon Huque, continue the process for as long as it takes until we get the problem solved.

So, what is the problem that we're trying to solve? There are two gaps, as I said, in the DNSSEC protocol specs. One is that the automation of DS updates when there are periodic [key] changes is not a process that was anticipated cleanly in the original design. The primary issue is that when the key changes within a child zone, within a customer zone, if you will, a DS record has to be updated in the parent, usually the registry.

There's no well-defined path for that. Registrars have access to the registry, and if the registrar is also providing the DNS service, then they can update the DS record. But if it's not that registrar that is running the DNS service, we have a problem, because third-party DNS providers do not have generally regular access to the registry. You'll hear lots more about that.

The other is that what started as an issue of how do you make a smooth, glitch free transition from one DNS provider to another in the event that

you're operating in a signed zone is actually a subcase of a more general operation of how do you have multiple independent DNS operators serving the same zone? And there has to be some coordination of the keys between them. And recently, through that work that Shumon and others have done, we have a protocol defined for that, and so part of the focus of this panel and this ongoing series is to focus on the implementation and deployment of this protocol, including all the software that's necessary.

The agenda for this panel consists of this talk which you're listening to, two talks on the DS update issue, and then three talks on the multi-signer protocol implementation. So we have Ondřej Caletka from RIPE NCC, Wilco Van Beijnum from University of Twente, and then Ulrich Wisser from The Swedish Internet Foundation, Matthijs Mekking from Internet Systems Consortium and Peter van Dijk representing PowerDNS.

On DS updates, there are a couple—here I've shown four different ways that it could be done. The arrows on the left which are dotted and blue are ones that are possible pathways but not actually seen in the wild. That is, there are no cases in which this is the way it's being implemented. That is, the DNS provider could have an interface that they could call and push the update either up to the registrar or up to the registry.

In the reverse direction, we have the possibility that the registry or the registrar could reach down by pulling, look at the zone being operated

by the DNS provider, and observe that there is a new key and pull that information back up.

The solid line from the registry to the DNS provider indicates that this is actually in use and that seems to be the solution that has been adopted in at least part of the world. The dotted line from the registrar to the DNS provider is a potential but there is some news to convey there.

So, as this shows, the state of affairs about the first solution is now being shown on the maps that Dan York is providing out of the Internet Society, and we're still in the initial stages of that and we expect that those numbers will grow over time.

With respect to the other possibility, this is a diagram that expands on the notion that the registrar pulls the child zone and then pulls that information and pushes it up through the EPP protocol to the registry. And here is the interesting development that is underway. GoDaddy, which is the largest of the registrars, has announced that they are going to implement this solution.

I check regularly and I have been authorized to present this slide which says development is still in progress. No date is announced yet for operation. However, we have some hopes that by the time we have the next DNSSEC workshop in October, that we will have something different to say about this, and better, presumably.

That's the short story on DNS updates, and as I said, you'll hear two more talks on the subject. With respect to multi-provider coordination and glitch-free provider change, one of the pushbacks that we hear

from time to time is, well, why is this so hard, why don't we just go insecure briefly? Seems easy. Who would notice?

And the answer is, well, if you actually are dependent upon DNSSEC, if you have applications that are depending upon DNSSEC working, then when you go insecure briefly, the effects may be much worse than you first think about. Websites go down, e-mail goes down, other DANE-based applications go down. We could go on at great length about this. We won't at the moment, in part because, as I said, this is a shortened session. But this is a point that we want to make.

The big picture on implementing multi-signer protocol is that there's a lot of moving parts. The checkmarks indicate the parts of this picture that are done. The square boxes are the parts that are in progress and the open circles, mainly on the right side and down lower, are the parts that are to be done but in the future, and the regular dots are no information because it's a mixed situation.

So this slide, which is not extremely graphic but is intended to provide a quick look at what parts of the whole picture are done, which ones are in progress, etc., and either this or some improved version of it will be a standard part of what we present in the ongoing meetings.

So the parts are a multi-signer controller, changes to the standard DNS packages, BIND, PowerDNS and the others to provide the interfaces that are necessary, and then operational deployments of this in various operators, and then over on the right, both analysis—the basic analysis has been done. I would love to see formal proof techniques because

there's a lot of moving parts in this and trying to get that right, and then what is also very compelling would be observation.

Eric Osterweil has been doing work looking at recorded records and doing an analysis of what would have been visible and what the validation and resolution status would be at various transitions. And then eventually we hope to have real-time observation mechanisms so that one could observe as these transitions take place whether or not they're working correctly or incorrectly and then actually run experiments demonstrating both correct transitions and incorrect transitions and have the incorrect transitions show up with some visible sort of alarm.

The multi-signer software is a project that's involving several different groups. I'm pleased to report that we have a cooperative sort of ad-hoc consortium, nothing formal, but nonetheless very active. We typically have a call once a week and there is software development underway, and I'm looking forward very much to being able to report positive results from all of this as we're going along.

The basic picture is that there needs to be a way to communicate between independent DNS operators, the keys that each one is using. The public keys, of course, the public parts of the public keys, of course. And what's missing is a piece of software that we call a multi-signer coordinator. That's the part that's under development.

And the other thing that is implied by this is that there has to be a way for this multi-signer coordinator to actually inject the keying

information from one to the other. And so that's the area that requires improvements in interfaces.

We intend, as I said, to go all the way to the point of getting this stuff to work and being able to demonstrate it and being able to demonstrate it not once but on a repeatable basis. And eventually, as I suggested, we'd like to be able to do it in a way that shows both positive and negative results so that, for example, if we violate the timing constraints, we can see that bad things happen, that there will be a loss of resolution or a loss of validation.

Part of this project involves development of testbeds, and there's no need to go into a great deal of depth here but just to give some idea of the complexity that we want these testbeds to have different operators and different points of observation and to see how all of this works.

The components of the multi-singer controller is that it's got to have interfaces to the DNS servers. It's got to have sort of a control mechanism that sequences the steps, and then user interfaces to each of these things and then modules to check on the success and report and acquire statistics and so forth.

With that, I'm going to turn things over to the others. I'll just close by saying that these are the pointers to the previous episodes, as I'm now calling them, as if we're running an ongoing TV series. And in each of these, you can see all of the presentations that have been made before. And then here's the pointer at the bottom to the Internet Society DNSSEC maps which you've heard about previously today.

So, thank you very much for that. Let me now ask Ondřej Caletka from RIPE NCC. Take it away, Ondřej.

ONDŘEJ CALETKA: Hello. Thank you very much. My name is Ondřej Caletka. I work for RIPE NCC, and I'm going to have a brief talk about scanning for CDS records at RIPE NCC. Next slide, please.

So, we started this project because there was a community demanding support for this kind of automation of DS records in the database. For RIPE NCC, we take care of mostly reverse DNS, but also things like enum. We hold the authoritative data in the RIPE database, basically, WHOIS database.

So, what we do now is we have a scanner that scans the updates regularly every day. The scanner has been recently open source, so feel free to look at it and maybe reuse some ideas from it for other projects. We only scan the secure delegation, so those that are already secured by DNSSEC, which makes things much easier because we can trust the data that are in DNS.

And we only consider CDS records because the database holds [directly] DS records so it also makes things easier for us just by scanning for CDS.

We fulfill all the special validation requirements that are mandated like this CDS record should be signed by KSK, not just ZSK, and so on. So we have all those requirements fulfilled. And once we do the scan, if we find that there's some change requested, we push the change to the RIPE

database, which will eventually push it into the DNS provisioning and make it appear in the DNS.

Currently, we have safety limit of 100 domain updates in one go. We haven't reached this limit yet, but if we reach it, it just says that some operator should look at it from our side if there is nothing wrong. So far, nothing wrong happened.

And also, there was a question about algorithms supported. We have no artificial restriction. The scanner uses python and dnspython which supports even EdDSA in current versions. So we have some issues with the resolver that is run for the scans which doesn't support EdDSA, but we're going to fix it. So basically, all current algorithms should work for that. Next slide, please.

Then just for illustration, this is the number of updates in some time, and as you can see, there are actually people using it, so this is the number of records updated. If you see the one that is there, it's just my testing [inaudible] rolling KSK every other day or so, so it's not worth it. But the others are actually our customers using it, so I'm very happy that this thing was actually ... There are already people using it.

Okay, Next slide, please. And this is the last thing of my brief talk. This is results from one day, and as you can see, mostly nothing changes. There are some updates pending—there are 40 updates—but I also want to say that we constantly see 17 zones that publish CDS records, but those records are not signed properly by KSK which is mandated by the RFC, so we don't do any updates for them even though they're appropriate DNSSEC signed, but not with the special validation

requirements. So we just keep them there until those that are publishing it will fix their publishing software.

Okay, thank you very much. That's everything from me. And if you have any questions. I don't see any. Thank you.

STEVE CROCKER: Thank you. Questions. "Steve, you mentioned a proposed solution to update information from private DNS providers to registry and registrar. My question is, is there any coverage required in terms of ICANN relevant policies or otherwise?"

The answer is not yet. That is—I've really tried to say two things. No, but there should be. It's not clear whether or not the best path is to try to get that as a requirement or simply to have market forces take over. If GoDaddy implements their solution, then presumably, that will first of all cover a large part of the market by itself because they're very huge, and the other is that it may stimulate others to do it and make the questions of whether there's a requirement less important. Thank you.

All right. Let's move on to Wilco Van Beijnum. And Wilco, help me understand how to pronounce the name of your university. I used an English pronunciation but I have a feeling that's not how it comes out in Dutch.

WILCO VAN BEIJNUM: It's the University of Twente. So, let me start then. Thank you for allowing me to present this. I'm currently researching the state of

DNSSEC automated provisioning for my bachelor project, and I want to tell you some of the results that I've gathered so far. Next slide, please.

So I looked at three things. The first of them is software support. Here, I've looked at both parent side support which is in most cases currently the registry, and child side support which in most cases is DNS providers.

For the parent side, I've looked at authoritative DNS software first of all. There's currently not much support, only BIND 9 features a tool DNSSEC CDS which can automatically update DS records based on a list of CDS records.

On registry software, however, there is a bit better support, FRED, Which is created by CZNIC, currently features full automated provisioning support. It is also used in practice, and it can automatically scan domains for CDS/CDNSKEY records and update the DS records of the child zones accordingly.

And thirdly, I've looked at some scanner software. The CDNSKEY scanner program is actually part of FRED, but it can also be used as a standalone program, so you can give it a list of domains and it will scan them and update you on the status of the CDNSKEY records. And as you heard on the previous talk, RIPE also has a scanner now which scans for CDS records and that is open source.

Then on the child side, there's a bit more elaborate support in DNS software, Knot DNS fully supports automatic provisioning. It can enable DNSSEC and also do automatic key rollovers without human

intervention. And BIND 9 and PowerDNS both feature ways to publish CDS or CDNSKEY records whenever you're signing your zone. Next slide, please.

The second thing I looked at is support top-level domains. In the table on the right, you can see some of the top-level domains that currently support automated provisioning. Some of them use the CDS record and others use CDNSKEY record. All of them also allow you to enable DNSSEC from an insecure side listed in RFC1878 and you have to wait between three and seven days and not change the records in that time for DNSSEC to be enabled.

I've also sent out a questionnaire to a lot of top-level domains, about 20 top-level domain registries responded to that, and eight top-level domain registries indicated that they are currently working on adding support for automated provisioning or are planning to do so in the future. Some of the top-level domains are listed on the slide. And a further eight indicated that they might add it in the future.

Some o the limitations that were indicated for adding support were mainly a lack of demand from DNS providers, lack of priority and lack of resources. Next slide, please.

So the third thing I looked at was top-level domain support for automated provisioning. For this, I've utilized the OpenINTEL dataset which collects daily snapshots of DNS zones of [inaudible] domains. Firstly, I looked at the dataset of .ch domains. There, I found that 7.2% is using DNSSEC, but over a period of a month, only 0.75% of domains ever published CDS records, which is not that much in my opinion.

CDS is actively used to enable and disable DNSSEC. It's not used very often. I've also looked at nameservers of DNS providers that publish CDS records and looked at which ones are at about 100% CDS record publication, and already, some Swiss DNS providers do this, so no other big DNS providers seem to be publishing CDS or CDNSKEY records at the moment automatically.

Finally, something interesting for registries that want to implement automated provisioning. I've also found quite a lot of malformed CDS records. For example, records that had the zone signing key instead of the key signing key or that did not have a corresponding DNSKEY. So that is something to consider as a registry, how you want to handle those records.

Finally, I've also looked at the Alexa one million domains dataset for domains that have automated provisioning records published even though top-level domains don't support it yet. Here, I found 2% of domains use DNSSEC and also 2% of domains actually already publish a CDS record, and 1.7% CDNSKEY records, so it seems to be used more than at the top-level domain which already supports automated provisioning, which is definitely interesting.

.com has the largest number of domains with a CDNSKEY or CDS record, which is also logical because they're the most prominent in the dataset, and .dev has the largest percentage of domain names with CDS or CDNSKEY record publication. Next slide, please.

So in conclusion, DNSSEC automatic provisioning is already used in practice, but not that much yet. That seems to be mainly because of a

lack of demand, priority or resources. I suspect that as more top-level domains or different registrars add support for automated provisioning, it will also be used more in practice.

If you want some more information on the results I gathered, there are a couple of reference slides at the end of this presentation. You can find them on the ICANN website, and I will also publish a paper in the future on this topic. So if you're interested in that, you can contact me via e-mail. Thank you.

STEVE CROCKER: Thank you very much. And I'll just flip through here the slides that are included in this deck that are available online include these references. There's three slides there. Four. Five. Six. Seven. Thanks. And Wilco, you're obviously in an excellent position to provide updates to the maps that Dan York maintains, so that'll be extremely helpful.

WILCO VAN BEIJNUM: Yes. I have already sent him an e-mail about it.

STEVE CROCKER: Good. So that covers the talks on the DS update part of the two problems that we're focused on. Let me move to the multi signer protocol and implementation. Ulrich WIsser from the Swedish Internet Foundation. Ulrich, the floor is yours.

ULRICH WISSER: Thank you, Steve. Hello. This is Ulrich from the Swedish Internet Foundation and, yeah, I'm here to give you a small update on the multi-signer project. Next slide, please.

So one thing we have been working on is making this list of capabilities that we need to implement multi-signer, and then we looked at this, how you could do this, and we came up with basically three ways to do it. You can do all the updates you need to do through the command line, you can do it through dynamic DNS, or there could be some kind of REST API or something that you could do it through.

And just to start with it, we think that the command line is a good first indicator to see if a software has capability, but it's not really usable in everyday operations because no DNS operator would allow me to get in, log into their nameservers and start using the command line to make changes to their setup.

So that leaves us in real life with dynamic updates on REST API, and we think that actually, the dynamic update [is] the way to go, but we included everything we looked at.

So what you can see here is that software usually has problems with ... They have automatic key management and so when we then want to inject additional ZSK that they don't have the private part of and stuff, then things start to get complicated and that is usually very poorly supported.

And the same goes for CDS and CDNSKEY, because this is handled automatically and we start working against the automation. So we

need to find a way to make this work, and we're actually in contact with these vendors of the software working on this and we'll hear about this later on in the program.

So what we found is that in PowerDNS, actually, for CSYNC it should say PowerDNS 4.5 because the upcoming version of PowerDNS will support all of this. And here you will see that PowerDNS comes with a REST API, so we included it in this test. And this is actually not saying that we didn't test the REST API on BIND, and not because they don't have one so they can't do REST APIs. Next slide, please.

But the world of DNS is not only software, the world of DNS is a lot of service providers, and so we looked at service providers too and you will see that—so we worked closely with deSEC, and so they support all of this. Obviously, they don't support dynamic DNS updates, but they support a REST API and a web user interface where you can make these changes. And as soon as the new PowerDNS will be released, they'll hopefully support CSYNC records.

And Shumon has been in contact with NS1 and Neustar and we hope to see implementation of this later this year, but we hope to have an update in the next ICANN meeting for this. And then we hope that we could get in contact with more service providers, so if you're a service provider in this session here, please contact me afterwards, we would be happy to include you in this and work with you to make this work. Next slide, please.

So, what have we been doing since the last meeting? We have tested the algorithm, we have tested to do it in PowerDNS, BIND and Knot, we

have seen that the algorithm holds, we actually did a few workarounds, sometimes when the software didn't really like what we did, we imported even the KSK just to make it publish the keys we wanted it to publish or the records, and so we got this working, which was a big step because it‘s always good to see that something you wrote up is really working.

What we have not yet tested is rollovers, and that includes key rollovers and algorithm rollovers, and that is future work we need to do. But the focus now is to get the first step working and then take the more complex operations [inaudible]. Next slide, please.

Yeah, so we have a draft written on this. It’s obviously a work in progress. We would like for you to read it and comment on it or send text if you have something you would like us to include or if something is unclear, please say so because ... this is actually hard to really express in words, the specifics, and it would be really good to have more eyes on it and get some feedback if we nailed it or not. Next slide, please.

Yes, and then we have the multi-signer controller, and that is the software we’re working on that should implement the algorithm and then control different softwares either through REST API or through dynamic updates, and it has a plugin architecture, we have a plugin for dynamic updates with PowerDNS and the REST API for deSEC, and at the GitHub where you find the draft, you'll even find versions of the GitHub for the multi-signer controller. So [if you're happy to look at it] and if you want to contribute or want to include a plugin for your specific provider or software, we would be happy to take code or work

with you to implement code, and we hope that we'll be having alpha release this month. Next slide, please.

Done. Thank you.

STEVE CROCKER: Thank you very much, Ulrich. This slide here is the first slide for Matthijs, so without any further transition, the floor is yours, Matthijs.

MATTHIJS MEKKING: Thank you, Steve. I'd like to start with a small clarification on Wilco's presentation because it said BIND uses a tool to publish CDS and CDNSKEY records. Actually, BIND 9.16 has a DNSSEC policy configuration and that will publish those records automatically when it is time.

Ulrich presented the capabilities that are needed in the software for a multi-signer project or the things that the software needs to support in order to make this work, and it can be done with [three] things: BIND can do it with the tools or the dynamic update. The tools is maybe not that interesting, but I would like to mention it anyway, because it allows people to actually write their own API against the command line tool.

We don't have plans on adding our internal REST API for this, although I should say that it's not on short-term. We have discussions on how we can improve [our NEC] and making use of the REST API, but that is further ahead.

When we look at the capabilities, what BIND Is able to do, I think mostly BIND supports this already. It uses the dynamic update protocol with no real specific things, so it should work just like your regular operations. Ulrich mentioned that DNSKEY without access to the private key could not be added to dynamic update and cannot be removed, but that is not true in my opinion, so if there are issues with that, Ulrich, that's fine, because it should work normally with dynamic update.

The only issue that we have in BIND is adding the CDS record with dynamic update, and that is because we have a small check that will look if there is a corresponding DNSKEY in the zone that matches that CDS record. And if that DNSKEY is in the zone from the other provider, obviously, that is all true. And so adding a CDS for a different DNSKEY is not supported currently, although we have scheduled to work on a fix next month. Next slide, please.

So, how does this work with dynamic update? We have [confirmation] so you can actually authorize who can update your zone. There's two ways. There's a simple allow update configuration which is IP based and which you can authorize with TSIG, and this can be set on any level, like [options, view level] or zone level, but it is really if you care about security and recommend it to be set on a per-zone basis.

The update policy allows you to have a more fine-grain control. You can define your own rules who can authorize it, who you authorize. This actually must be set per zone, and I've provided two examples here, for example, the first line here provides provider A access to a specific zone

and for these specific records. So it can only update the DNSSEC-related records if you provide this [rule.]

The second rule just provides access to provider B to all zones that the other provider has. And this might be useful if you have multiple zones that you are transferring or if you have multiple zones that you have a multiple-provider situation for. And that's all I have.

STEVE CROCKER: Thank you very much. And move to Peter van Dijk, PowerDNS.

PETER VAN DIJK: Hello. Thank you. I think this slide speaks for itself, so I'll just provide some background. Back in 2012, SIDN, the .nl registry, did some examples with secure transfers between providers using the EPP key relays, so no CDS or whatever, but this did require registrar cooperation to help them with those experiments. We added our direct DNSKEY setting which allows the mixing of local and foreign keys, and because we did those experiments back then, it turns out we were quite well prepared for all the things that Steve and Ulrich and Peter have been asking us the last year. And indeed, as Ulrich says, [TSIG] will be 4.5 which we should release in a few weeks. All other things that are needed have been in there for quite some time. And like Matthijs, I would like to provide an update to Wilco's talk. Yes, we have a tool, but the tool just turns CDSKEY publication on or off. If Wilco got the impression from the documentation that the tool does a publication right now instead of

automation, then maybe we should fix the documentation. That is all I have.

STEVE CROCKER: Thank you very much. So there's been a couple of comments about whether or not the data presented in one talk is accurate with respect to one of the other speakers. This is a workshop as opposed to a formal publication, and so I consider it a benefit and a positive outcome that as part of presenting the status of things, we also get a little bit of interaction between the speakers and that facilitates forward progress along all these fronts. So I expect that we'll have clarity and improvements over time as this goes forward.

And let me add along that line, one of the things that I should have included that has come up in discussions is that there are some limitations listed in the RFCs with regards to algorithms having to match, for example, which we think are not actual requirements from—they should not be required because there's no necessity for them, and in fact, those would be stumbling blocks if you're trying to move from one provider to another and they don't share the same algorithms, but it should still be possible to do that transition.

So, over time, all of these things will get straightened up. So I guess we're out of time. Kathy, you should say whether or not we can take any questions. I don't see any open questions in the question pod.

PETER VAN DIJK: There was one question but I typed the answer.

STEVE CROCKER: Good.

KATHY SCHNITT: And if anyone wants to ask a question verbally, just go ahead and raise your hand, we'll be happy to unmute your line.

STEVE CROCKER: All right, this completes the panel on automation of the provisioning aspects for DNSSEC, and I turn the program back over to you, Kathy.

KATHY SCHNITT: Steve, wonderful, thank you very much. We can now move on to the moderator of our next session, which is Jacques Latour, and we have some presentations from Adiel and Michael. Jacques.

JACQUES LATOUR: Hello. Can you share the agenda?

KATHY SCHNITT: Sure.

JACQUES LATOUR: So I did have a question about the other session. CSYNC. I think next ICANN meeting, next DNSSEC workshop, Steve, we should look at CSYNC also and who's doing what in terms of implementation there, because that's parent-child synchronization of nameservers, and that's

really interesting, and also there's other stuff. So I think it goes along with CDS automation.

STEVE CROCKER: Thank you. I've made a note to that.

JACQUES LATOUR: Okay. So the next session is Adiel. From what I understand, he's from Montréal, 5:17 AM, and he's going to talk about KINDNS project.

ADIEL AKPLOGAN: Thank you, Jacques. Thank you, everyone. I will quickly take you through a new initiative that ICANN is starting, which is KINDNS. That's how we call it. So I'll try to share my screen, and hopefully it'll work.

Okay. So just to put this into context, everyone participating in this session knows a lot about the DNS and we all know that the DNS although originally designed as a very simple and straightforward protocol, over the time with the evolution of the Internet, it has also evolved to address and cater for several other needs, which are important, of course, but those evolutions which are an additional layer to it make it sometimes a bit complex, and you know about the DNS [inaudible] project which has tried to identify all the standard and the recommendation through RFC that talk about the DNS, and it ends up showing that there are even more than 297 RFCs that talk about the DNS with more than 200 pages for that. And there are different kinds of evolution. Of course, not all of them are implemented everywhere.

Those evolutions happen at the level of different components of the DNS, I must say.

So this matched the feedback that we have been receiving in our department, which is technical engagement where we try to engage, as much as possible, with operators, and something that we realize is that people who are running small operation kind of get lost in all of this.

So, that complexity is there. For people running the DNS for small operation, community network, internal IT infrastructure, they want to provide the DNS service that is fast, that is reliable, that just does the fundamental job, but also, they don't want to be the weakest link in the whole ecosystem security. They want to make sure that they're running properly and they're providing a secure service to their user.

Although they want to provide that kind of secure operation, secure service and reliable one, when they go back and try to look at what are all the recommendations, all the standards there, yeah, it kind of gets scary, and they kind of back down and try to focus on the basic one and get it running and forget it somewhere there. But big operations, they may have the means to follow all of this, they have specialists that get to attend all these kinds of events that we have where best practices and evolution are being taught.

So to be a specialist has become difficult for some kind of operator. So, what can we do to help everyone? What can we do to level up the operation of the DNS overall so that we can at least agree on certain key components? And to be honest, the idea is to get small operators, those

who don't run very big DNS service to be able as well to say at least we have implemented what is important.

So we know that recently, the encryption has been added to the DNS between stub and resolver, but there are other evolutions that are coming up. How does small ISP keep up with that evolution? How does medium-sized enterprise assure that they're following all the elements and guidelines, as I mentioned? Community network operators, we're seeing a lot of community networks coming up in various areas of the world, and they run DNS, of course, and sometimes they don't have all the information to be able to ensure that they are running the best level of service, they are running DNSSEC. Many of them think that it will be very complicated for them to maintain, so they go for the easier way.

Corporate IT managers used to run a lot of DNS infrastructure even though recently, many are kind of pushing that out to third-party vendors. But even if they are doing that, they need to understand and know what exactly to expect from their provider.

So that's where this initiative comes from. So KINDNS stands for knowledge sharing and instantiating norms for DNS and naming security. We pronounce it kindness. And as many of you know, this plays a little bit into MANRS. So the idea is that if people are running the DNS with KINDNS and trying to apply the most important best practices, we will contribute to a little bit more secure and safe Internet when it comes to the DNS.

So the whole idea at the need is to provide and publish something simple that everyone can refer to, and being small, big or medium-sized

operator, that will point them to the most critical element for operating a secure DNS service. We are putting the emphasis here on the operation of the DNS because that's where we want to start with, probably extend that to other components of the DNS, of course, such as the registry and registrar interaction, maybe getting software vendors as well in. But to start with, we want to really focus on DNS operation.

So, how are we going to do that? The first step—that's where we are right now—is to identify and document the most critical security norm. Most critical here because again, we want to streamline all the best practices to the most important one. While doing that, we don't as well want to go through the [mass], the lower level, but we want to have the critical mass, say, the critical 20% that can help us achieve 80-90% of a secure DNS operation.

We have discussed this with the community and got some consensus around those important things, they'll be documented and published online for operators to consult, and also to join the initiative. Actually, the success of this initiative will depend on how much and how active operator will join the initiative, becoming kind of a flag holder of those best practices, and also help us relay the information and the knowledge around those best practices.

We want as well to make sure that doing this, we are able to measure the impact, so identifying some key indicator that we can measure over the time to see the impact of KINDNS on the secure operation of the DNS in general.

I mentioned the registry-registrar-registrant aspect that is very core to ICANN policy function, and as we are doing this, we will certainly evolve this into those areas as well, but again, we are starting with the core operation, then we'll move it to the procurement side of the DNS, which addresses this.

So, as I mentioned, we are covering the operations side. We can't talk about the operation without talking about having the operational environment itself, so that would be kind of a category that would touch on all the different elements, but then we will specifically target authoritative server both at the top-level but also at SLD level, because there are many registrants who run domain names and who run their own DNS, and we should be able to give them as well practices that they can implement.

And also, recursive resolver, we want to address both those who are running recursive resolver for their own infrastructure, those who are running them from public consumption, that means they're open and they provide the service either to a set of customers that they have or have it completely open like we are seeing more and more these days.

We are also planning to target resolvers that are run within enterprise or ISP infrastructure, or those who are also run using Anycast. Because we don't want to touch on everything, we plan on having a specific session dedicated to some evolution that are not critical to the core of running a secure DNS but are important to consider. And when operators are going this direction, to know what to expect, to know what that will imply on their infrastructure, and we will address stuff

like DOH, Anycast, other stuff there and we will develop those aspects as we go along with operators and experts in that realm.

In defining those best practices at this point, we are using the framework of what must be done, what should be done or what may be done, knowing that the focus of KINDNs will mainly be around the musts particularly, and sometime on the "shall." So to really streamline the most important.

We plan then to help those operators in order to implement those best practices to publish guidelines, checklists, configuration process, example, etc. where they can, after understanding those practices, also have all the support and the guidelines they need to get this done.

So yeah, we are still at the very early stage. When I say very, there's already some output coming out. We hope to launch this by the end of this year, 2021. At this stage, we are streamlining the best practices, trying to engage the community. We have a mailing list that exists. Once we have stable best practices, we will start launching the dedicated website. Right now, we have a Wiki page where we started publishing some of the information around the project. That'll be a temporary repository until we have the dedicated website.

We are really interested in having feedback from the community on all of this, particularly while identifying the best practices we want to focus on. So the mailing list KINDNS-Discuss is open. If you want to contribute, you're welcome to join the discussion there. Over the past few days, there has been an interesting discussion there on the fact that

we want to make validation a must for instance, what that means and how the community are seeing that was interesting to watch.

So you are welcome to join and then provide input there as expert in operating DNS. That's what we want to hear. But while doing that, we need to keep in mind that we want this to be straightforward so that anyone, no matter the size of the infrastructure, are able to implement them.

We have a subject matter expert consultant who's joined us to work on this, Tim Wicinski who you may know. So that is it on KINDNS, and I will be happy to answer any questions or respond to any suggestion that you may have. Thank you.

JACQUES LATOUR: Thank you, Adiel. We'll do the Q&A after Michael's presentation. So, Michael is going to present DNS resilience program and why [inaudible] RFP. There's a question, "Will your presentation be available on the ICANN workshop page?"

ADIEL AKPLOGAN: Yes, it will.

JACQUES LATOUR: Thank you.

KATHY SCHNITT: It's actually now there, so you can go and find it now.

JACQUES LATOUR: Michael.

MICHAEL HAUSDING: Okay. Thank you, Jacques. Welcome, everyone. My name is Michael Hausding. I work for SWITCH, the registry for .ch and .li. I'm going to talk about our DNS resilience program. And with that, we issued an RFP for DNSSEC measurements that raised some attention, and I'm going to talk about the program and the RFP.

So good thing is I don't need to talk about DNSSEC and explain what DNSSEC is. I'm just going to tell you where is .ch with DNSSEC. If you look at this chart, it looks quite good. We have 150,000 signed domain names. But the bad news is that if we compare this with others that are leading the DNSSEC for their TLDs, you can see that .ch with about 7% is quite behind other ccTLDs that promote DNSSEC for a long time.

And this is why the office of communication, which is the regulator for the .ch domain name, mandated Switch with giving financial incentives for DNSSEC signed domain names. And we have a contract with the office of communication that basically puts a steering charge on the domain name at the time of registration and renewal that we collect, and then later redistribute it to domain names that are signed with DNSSEC.

I'm just going through this quickly. What we do is from the 1st January of 2022, if you renew a domain name and it's signed with DNSSEC, you pay the price that is valid at that time, and if you renew a domain name

that was not signed with DNSSEC at the time of the renewal or if you register a new domain name, you will have to pay one additional franc for that domain name.

And all these additional francs go into a piggybank, so we collect them until the end of the year, we have some fixed compensation for the measurements and also for the work that Switch is doing to run this program, and at the end of this year, we're going to redistribute all the money that we collected to the registrar by their share of DNSSEC-signed domain names.

So there are two incentives here. The first incentive is that registrars have to pay one additional franc if the domain is not DNSSEC signed, and the second incentive is the more DNSSEC-signed domain names they have, the more they get back at the end of the year.

The question however was, how are we going to measure which domain is validly DNSSEC signed and which domain will get some of the piggybank?

We talked to other registries and they talked about their experience and they said, yes, just looking at a DS record in the parent zone is a bad idea because then registrars will sign the domains and will forget it and you will have a lot of broken DNSSEC domain names. And that's why we thought we need to measure the signing of these domain names and only if the measurements indicate that the domain name was validly signed, then we will give the lower price and also give a point for the cashback.

We have something called—we checked nameservers off that run a lot of domains for .ch, but we thought that because we're going to collect and redistribute a large amount of money, we really want to have an independent expert who is doing the measurements for us. Also because of operational issues, we're not used to really large-scale measurements and we think there are people out there that can do these measurements and offer this to Switch as a service.

And this is why we issued a request for proposal on May 7th. We sent it to seven potential bidders where we thought, "Okay, they're already doing DNSSEC measurements," and we also sent it to the DNSSEC-COORD mailing list and asked for proposals for these measurements.

We only left two weeks of time to answer the RFP. We know that was short notice, but as we have to start on the 1st of January 2022, there was unfortunately not more time.

We received five proposals and we evaluated them. We gave some criteria on the evaluation. We said the measurement itself is the most important thing and that will give 50% of all the points you get for your proposal. We have optional services like showing a dashboard or publishing statistics about DNSSEC, can give 10% extra. We looked at the qualification of the bidder, his experience and his knowledge on DNSSEC and DNS and if he had publications about measurements and DNSSEC.

There were some other requirements, like do the bidders support DNSSEC themselves? And in the end, we also took into account that we

want to extend the contract for other security standards like DMARC and DANE in the future.

We had seven interested bidders. In the end, we received five proposals. Two potential bidders chose not to offer because they either had no resources because of the short notice of this RFP, or they didn't want to offer a service that has kind of an SLA where they need to have operational stuff. So that was one of the points where research-oriented organizations didn't feel quite comfortable.

But in the end, we received five proposals. We evaluated them with five people, so we had three people giving points and we had someone checking all of the technical aspects and someone checking the legal aspects. And in the end, we came up with a candidate that was our favorite, and all the others were—yeah, so all of the five proposals fulfilled all our needs. So they all could measure DNSSEC, they all could offer most of the services we wanted to have and also, most of them offered all the optional service we asked for.

So there were some weak points of the other proposals. First of all, not all of the bidders used DNSSEC themselves in a way where we thought that someone who is measuring DNSSEC would be suited for as a service operator. Some of them were lacking experience in the operation and the operation of large data and large-scale measurements, and the third point was the capability to build the service to a schedule was in question by some—they could meet our schedule, but they were somehow limited in the features they could add to our schedule.

In the end, we had one clear candidate. Unfortunately, I cannot say who it was because we are still in contract negotiations, and the reason to choose this bidder was that he fulfills all our requirements, like all the other bidders, but one big advantage here was that he's already running a measurement service and he had no issues in scaling and in getting up the measurements running because he already does these measurements, and this is for us the smallest risk for operational issues.

We see an issue here that the way that results are presented to the registrars and nameserver operators are limited, but we received an offer that improves this graphical user interface and we hope that we can also help others here if the graphical user interface of the service is improved.

And third of all, like all the others, it was a very trusted organization and we feel comfortable that this measurement is a success. So, what is next? We need to complete the contract and then we are continuing with the measurements for the next five years under this contract.

We plan to start offering the measurements to DNS operators in the end of October this year, and from the 1st of January 2022, we will collect the additional 1 franc for each domain name that is not DNSSEC signed, and in the end, we redistribute the collected fee to all the registrars based on the number of domains that are DNSSEC-signed.

For 2022 and 2023, we only have DNSSEC as the things we're measuring, and for 2024, 2025 and 2026, we want to add additional DNS-based security standards that we're going to measure and incentivize.

As you can see, we hope that the number of DNSSEC signed .ch domain name is increasing and that means that the amount of money we're going to redistribute to the registrars will drop over the five years. And we hope that in the end of the program, we catch up with the other ccTLDs that are around 60% of DNSSEC domain names.

So, this is our plan for the future, and I hope we can improve the rate of DNSSEC signed .ch and .li domain name with that program. Thank you.

JACQUES LATOUR: Thank you, Michael. So now we have Q&A for Adiel and Michael's presentation. There's one question for you in the chat, Michael. Can you explain dashboard key features that you mention in the presentation about DNSSEC measurement?

MICHAEL HAUSDING: Okay, so basically, the RFP was quite open. We said we wanted a dashboard where registrars and DNS operators can see what goes wrong with signing. So we have some requirements that you need to fulfill if you want to get the lower price. That means that the domain name needs to be signed with a current algorithm, it has to answer correctly for a given label and for NX DOMAIN, and we excluded a certain amount of algorithms like algorithm seven.

That all should be visible in the dashboard. And basically, what we want to do is to give a registrant overview of broken DNSSEC configurations. So our goal is that registrar sign with DNSSEC or DNS operators, and if they do mistakes with DNSSEC, they will see it in the dashboard. So

we'll send them a mail every day and in this mail, we will have a link saying you have a domain name that has DS record in the .ch zone but the signing is not valid or has some minor issues, and then he can click on this link and this link will take him to the dashboard where he sees what's wrong with the domain name, and most likely, he will also get hints on how to fix it.

There are other things in the dashboard, like for a given nameserver, what is the most common configuration error? So if you have a registrar, we can tell them, "See, you have this error that is responsible for 50% of all your wrong configured DNSSEC domain names," and then he can look at this issue especially and improve it.

JACQUES LATOUR: Thank you. I had one question for Adiel on KINDNS. From what I understand, MANRS is a simple framework to comply with. Are you looking at KINDNS to be as simple?

ADIEL AKPLOGAN: Thank you, Jacques. Yeah, that is really the goal. The goal is to keep it really simple, and that's why, as I was mentioning during the presentation, we really want to streamline those multiple best practices to the most important one. Of course, this is not only about DNSSEC, it's about securing DNS operation in general, and DNSSEC is going to, of course, be one of them. Maybe the most important one, but we will try to add some other elements as well to have a secure ecosystem. But sure, our goal is to make it simple, to make it

straightforward, very limited number of best practices that can be impactful.

JACQUES LATOUR: Okay. Thank you. I don't see any other questions, any panelist questions, or raise your hand, or any questions. If not, then this panel is done. Kathy.

KATHY SCHNITT: Thank you. Well, thank you all for joining us for virtual ICANN 70. Thank you, everyone. Have a great day.

[END OF TRANSCRIPTION]