ICANN73 | Prep Week – Knowledge-sharing and Instantiating Norms for DNS and Naming Security (KINDNS) Update
Wednesday, February 23, 2022 – 11:00 to 12:00 AST

ADIEL AKPLOGAN: We will be giving you a brief update on KNDNS, which is an initiative to promote DNS best practice in general. We have translation today, so we will request you to use the translation. A speaker will try as much as possible to speak slowly and articulate so that the translator can do their job.

Ask your questions in the Q&A pod, so that it can be relayed to us. I guess if you want to speak up, you can also raise your hand and they will allow you to speak. Next slide, please, Steven.

So, we will be having this presentation in two parts. I will be taking you through the first introduction and then I will introduce you to [Philip Regnauld] who is our expert on this project who will give you a little bit of a view of where we are on the substance of this project.

So, we have had a session during ICANN72 where introduced KINDNS and explained a little bit what we want to do, so I'm going to start straightforward with what KINDNS is about. It is a knowledge-sharing association known for DNS and [naming] security. In short, this is about sharing best practices and being able to work with the community and the DNS operator in general to voluntarily add their comment to those best practices but also help all of us at ICANN to promote the best practices globally and make sure that we work together as usual to keep the DNS as safe as possible. Next slide, please.

So, some of you are already familiar with MANRS which is an initiative around securing the routing online. KINDNS play a little bit it into that so that we can say that if we have good MANRS and KINDNS, we will do a little bit more securing and keeping the Internet safe for everyone.

What we want to achieve here again is to produce a very streamlined best practices that operates or can easily follow to ensure that the way they operate the DNS is secure and we want to put the emphasis on the DNS operations here because this is going to be purely technical best practices on the DNS operation itself. Usually, around the DNS, there are all the services, there are the applications running and the initiative will try to focus more specifically on the DNS operation, the technical aspect of the operation.

If you are familiar with the DNS, you probably hear about the DNS [inaudible] which has identified all the best practices [inaudible] that are related to the DNS which is a lot, more than 2000 pages talking about the DNS. What we are trying to do is to kind of bring out of all of that what is the most important thing any operator is running any component of the DNS should be doing to keep the DNS safe. Next slide, please.

So, the initiative was [translated] into a few components. The first one is to identify and document the most critical security norms for DNS operations. This is being done with the help of Phil Regnauld I mentioned earlier and the community at some extent because we have a mailing list where we regularly share what is coming up from that work of notifying the norm and getting feedback from the community.

Then, based on those best practices, we will launch a dedicated portal to publish the best practices of course, to provide guidelines on how to implement them, but also provide the community a place where usual or useful information can be assessed where operators who are supporting the initiative, who are committed to implement those best practices can also join and help us promote the best practices, as I mentioned.

So, everything is going [inaudible] published on a dedicated website, on the KINDNS.org domain name. Once we do that, we will start working again with everyone to identify some indicators which can help us see the impact of the initiative on the security of the DNS, identify some of the key indicators that we can measure over time and see how they are evolving in the right direction or the wrong direction and adjust accordingly. That would be the next phase after we run this.

Usually, when we present this in the ICANN context, the impression we usually have is: is this going to touch on the function of the provisioning function of the DNS? That means the registry, the registrar, and the registrant best practices as well.

The short answer is no at the beginning, because as I mentioned, this is focusing and targeting the core function of the DNS. But probably in another phase of this, we will look at how we can map some of those best practices to the provisioning system, meaning the registry and the registrar best practices and add that to the initiative. But it's already [inaudible] but the first phase will be focusing primarily on the operation. Next slide, please.

So, we have focused on certain categories. Again, if you are running the DNS, you are running it in an environment and there are other components that contribute to the robustness and secure operation of the DNS in general.

And as I mentioned originally, we are not going to correct everything, but we are going to focus practically on the components of the DNS, meaning the authoritative server run in different environments, meaning TLDs or people who are running critical zone can be a second level as well. And second level domain name manager in general. So we have the TLD in critical zone, such as [.seal.uk] for instance or any NIC.TLD are critical for the operation of those TLDs, so we are going to add them to that first category. And then the second category are everyone having any registrant managing a second level domain name.

Then we have the recursive server resolver operator in general. And in that category as well we have three sub-categories that we will be addressing: private (closed) resolver, shared private resolver, and public resolver.

So the best practices that we are going to be promoting are going to resolve around those five categories. But around that, we will be providing guidance as well on how [inaudible] operational environment in terms of service, system, and network and touch on some privacy considerations as well that impacts some aspect of the security in general.

At this stage, if any operators want to join KINDNS and work with this initiative, it will be assessed again—or self-assessed again—against only the two main categories I mentioned which will be defined either in the authoritative server operation or resolver operation in general.

We will publish of course implementation guideline, how to, checklist, and configuration process, examples, etc. There is a question there about which software are we going to provide guidelines on. We are going to us the most popular one, but of course if we present it in a way that people understand the concept, they can easily apply those concepts to any other software in operation out there. Next slide, please.

From here on, I will give the floor to Philip to take us through the different elements we have been looking at in those different categories. Philip?

PHILIP REGNAULD: Yeah. Thank you, AdieI, for the introduction. This is a very interesting project and I hope we can get some feedback from the community about the various best practices that we're putting forward and have identified.

The way we split this up, the way we categorize the types of DNS operators, we decided to go after … Initially, we thought maybe something like top-level domains and second-level domains. Turned out it was maybe a little bit not simplistic, but it was a little bit more complicated than that.

So we decided to first identify whether or not one was going to be an operator of authoritative domain names or an operator of DNS servers offering recursive resolver service.

So, to start with the authoritative server, what we did is we looked at the type of zones out on the Internet by hierarchal nature, the DNS of course is more critical the higher you go with obviously the root being the most critical and maybe the most—not vulnerable, sorry, but the most sought after target if you're going to compromise it if you're going to have a security incident, that's going to be the choice more so.

So, TLDs have an important place. But we also looked at other domains that could have an important operation. For instance, if we look at the name servers for a lot of countries, it's not uncommon to see, for instance, in where I live right now in Denmark (DK), the name servers for the DK domain are placed in a sub-domain called nic.dk. Well, arguably, these second-level domains will be probably just as important as the TLDs that their name servers are living in.

So, therefore, we decided to split these categories as such. Critical zones, we call them. And obviously the top-level domains but all the auxiliary or so-called support zones that are used for providing service of some kind, either hosting name servers or similar dependencies.

And another one that we decided to include in the critical zones are not directly tied to the operation of the DNS itself, but if we look at it from a ccTLD perspective, some zone files or DNS names that are associated with very critical, either health care, e-governance, citizen services identification systems, there are such that are maybe more important than others, but in this case, I took an example. In Denmark, we have a national identity system and I realize, well, if myid.dk is down, then a lot of people are not going to be able to log on. And you might consider how does that impact the DNS, well this is going to be driven by a self-assessment model and what we really want to do here is not necessarily dictate which zones are critical or not. It's more a framework for people to identify: am I running critical zones? Or for an organization to identify: are we providing critical services? Therefore, what are the guidelines? What best practices should we be observing to protect these services?

Of course, we've also included things like finance and banking sites which can be considered critical for the functioning of an economy and a country.

So, this is again, kind of, I wouldn't call it arbitrary but it's our decision to structure it like this. It also, in terms of vulnerability and impact to a constituency, an economy, a country, what happens if these domain names go down obviously with the ccTLD itself being the most critical? Next slide, please.

Then we have other domain names. That's of course all the other domain names below the top-level domains. They're also critical. They'll be providing all the various services and websites and e-governance and e-commerce and all the things we know to be on the Internet. Those have to managed responsibly as well.

So, it's not because they're second-rate citizens or anything, but maybe there will be a few less constraints on the kind of best practices we want to observe and the reason is we want to have this operation is maybe to encourage people to come to get started with KINDNS, to start implementing these best practices and not be put off by the complexity, and therefore, well if you're running a second-level domain, you're probably not going to impact other people so much since you don't run delegations but you still have … You could be subjected to cyberattack. You could be subjected to system failures. And are you doing the due diligence? Are you running the best practices that would allow you to either mitigate or at least recover from these incidents?

And it doesn't really matter how important the domain is in that respect. Everybody should be doing this because a misconfigured or a domain that's been hijacked is going to be disruptive in one way or another. Next slide, talk about … Yeah.

So, the recursive DNS operators. That's the other half of the ecosystem, right? We've got the authoritative operators and then we have the recursive DNS operators and here we had to consider what kind of recursive resolvers do we have out there.

They'll either, in a nutshell going to be public or they're going to be private. When you look closer, there's going to be private resolvers and those are going to be fully corporate networks, fully closed, not accessible from the outside, VPN type access and usually on private address space and those are going to be companies and organizations, such as healthcare and banks and most businesses are going to be structured like that. But to some extent, also, home networks and maybe residential networks.

And a little bit more open, we have shared private resolvers and it's maybe a fun name but we try to find something that … I didn't want to specify necessary ISPs or any kind of providers. We say shared private because they're private to a set of customers or a set of maybe institutions. I imagine a university network. They're not accessible from the outside but they're still being shared by possibly legally distinct entities. For example, multiple customers sharing an ISP's resolver or it could be in the context of a federated set of organizations under one common technical administration.

And these are going to be one other category and, after that, we're going to have public resolvers. Public resolvers, we know them immediately. We think of maybe 8888 from Google or Quad9 and these kinds of services, but in between that, and actually closed DNS service, we have commercial DNS filtering which may or may not be open in one form but … Sorry, I'll try and talk a little bit slower.

For the public resolvers, we have, as we mentioned, Google and similar, and then we have semi-open or—how do you call it? Open resolvers with a commercial component, where with the right agreement or contract, you'll get additional service from these particular resolver operators in the form of scrubbing or passive DNS service where your DNS traffic is being analyzed to see if maybe you have hosts that are compromised or infected by malware.

And these types of operators are public usually but they will have some kind of access control mechanism and maybe some fee to be paid or a contract to be entered, so that you may use their service and benefit from the added—how do you call it? Value-added services they will be providing.

So, that's the different categories, and in the section above, we identify how are these resolvers—how is access restricted to these resolvers? It'll be a mix of IP address access or it could be certificates or it could be VPN. It doesn't really matter. What really matters here is that we try and place each of these operators in a category where they can go and look at KINDNS and say, "Well, okay, this applies to us. What are the best practices?"

And also, to some extent, maybe for end users and organizations. For instance, a business that will want to look up, "Hey, I'm connecting to my ISP. I'm using their resolver to forward my DNS queries to the Internet." What best practices should they be observing and are they sticking to it? Are they actually following the program?

Then you can refer your provider or your ISP or your IT department. Are we following KINDNS for these resolver best practices? Are we protecting our user's privacy in the way that it is described here? Next slide.

So, recommendations for private resolvers. We haven't really covered a lot of the recommendations for the authoritative resolvers but that's not too important right now. Let's focus on the private ones.

Private resolvers, as we mentioned, are on private networks. They're in, some cases, part of a trusted competing domain like active directory, things like that. So you'll find them in a lot of Windows environment.

And what we did when we focused on those, we focused primarily on the security of the network and we identified also the need for transparency which means that certain recommendations that might make sense other places like DoH or DoT, a lot of these networks are still uncertain about enabling these services but they can still benefit from using DoH or DoT, for instance, and forwarding their queries to an upstream resolver that has security—sorry, encryption enabled, allowing you to at least secure eavesdropping of the queries on the way out of the network. So that's a slightly different scenario for each of the operators.

Then, availability and resiliency of services. Well, as Adiel mentioned, some of these best practices will cover good old systems administration best practices, and therefore we don't go into them right now, but they'll be detailed on the website and in the program so that at least there will be something to refer to in terms of system hardening and proper system administration.

And very prominently here, we do see DNSSEC validation. I forgot to mention this. For the authoritative servers, we will of course expect that one of the leading best practices that we'll be promoting for authoritative zone operators will be DNSSEC signing, and here in the resolver side of things we'll be encouraging—or rather make it a requirement—for resolvers to do DNSSEC validation.

This is fortunately already the case for a lot of the software out there but we'll make this, underline it, and say you must be doing DNSSEC validation at this point in time. Again, it's probably already the case for a lot of software. Next slide.

So, shared private resolvers we've discussed are ISP types. They're ISP or similar. They'll have similar requirements to a lot of the private resolvers but they'll have a different spectrum of customers because they'll be dealing with maybe a mix of mobile, cable, fiber, residential. There's going to be some access control.

And because these resolvers are shared between many different customers, there's also a privacy issue. There are such things as cache snooping and other things we're not going to talk about right now. But you want to be sure that when you're offering DNS service you do so in a privacy friendly fashion, and therefore one of the things we do recommend here is to enable DoH or DoT or both on your resolver service so that those customers of yours that feel more comfortable forwarding queries to you, they can do so. And also not be—not tempted but feel that they can use your services and not have to find a resolution elsewhere for the sake of latency. They can use your network by offering DoH and DoT. They can use your systems and they know they already have a business relationship with you as an ISP. Then you should be offering DoH and DoT. That just makes sense in terms of privacy. And this is of course alongside existing traditional and encrypted DNS which is going to be around for a while.

Well, availability and resiliency of the DNS service. There's a number of recommendations there that we make as well. Those are on the Wiki and on the website soon, and those are again good system practice, hygiene, hardening. And once again, DNSSEC validation. It would be expected at this point that anybody running an ISP resolver would be doing DNSSEC validation. Next slide.

For the public resolver operators, we are talking about mostly the big open ones, like Google and all the others. Of course, they have their own way of doing things, but I'm sure that a lot of them already implement a lot of these best practices and one of the things that we do put [inaudible] now is DNSSEC validation, and luckily the big names are doing DNSSEC validation and privacy consideration, DoT, DoH. All the big names offer either DoT, DoH, or both so it's possible to forward your queries to, say for example Quad9 or Cloudflare is 1111 and use their DoT or DoH service.

Another thing we mention is Qname minimization and forgot to mention this for the others, but Qname minimization to avoid leaking fully qualified domain names unnecessarily towards the root. By turning this on, you make sure you only reveal part of the domain name, because whether we want to or not, domain names are revealing. They disclose information about our users and their habits and what they're viewing. That's not always for public consumption.

And the same goes for the closed public resolvers offering service for either payment or based on some agreement or access control, they will also need to offer DoT/DoH which is probably already the case for a lot of them.

And Qname minimization is almost always enabled nowadays as a standard practice but now we're making this, I would call it, a requirement so to speak. And DNSSEC validation, obviously. Next slide.

I don't know if that's for me to cover.

ADIEL AKPLOGAN: I can take it from here. Thank you. So, that part from Phil gives you an overview of how the different best practices are going to be structured. Again, keep in mind that our goal here is to streamline this and focus on the most important one, wanting to only have a very certain number of best practices to implement, not more than our target is 7:10 practice in total for people to identify themselves with, so not to confuse them further.

Next is to give you a little bit of review of how all of this is going to be presented and structured when we launch this. As I mentioned, it will be hosted on the KINDNS.org which will be a dedicated website for it, supported and sponsored by ICANN primarily, where people who are interested to join the initiative in different ways can have information they need.

So there will be a session that touches on the different categories that Phil mentioned. Support and engage area will be where operators who want to join and support the initiative can enroll either a sponsor, member, domain operators, or ambassador of the initiative.

Then we will have a tool session where operator can self-assess themselves. We are still discussing and looking at how those self-assessment tool is going to be. What we want, we want it to be a very easy, simple, straightforward because we are not making this mandatory. This is a voluntary engagement, so people must engage at their end of committing to implement these best practices in general.

So, the self-assessment tool as well is going to be based mostly on people's self-commitment here into implementing some of those practices and answering those questions up front, for instance.

We will have a dashboard. We will develop a dashboard that will give us some information. And as I mentioned earlier, being able as well to track some of the identifiers that we will have chosen to allow us to see how this is impacting the landscape in general. And then we will also have the guidelines on how to implement the best practices or giving more detail and more information about any of the best practices or consideration.

For instance, if you look well at the slide that Philip mentioned everything related to hardening the core or system security and all those things, which are not going to be part of the core best practices are going to be highlighted in the guidelines, for instance, where people can go and look at them. But they are not part of the core of what we want to do. And then we will have the blog, events, and all the other things related to that.

So, as we are working toward this, summarizing, identifying the most critical best practices, we have also started developing some of the guidelines. We have released recently a guidebook for DNSSEC signing in one of our [OCTO] documents. Those kind of documents are going to be referenced is well in KINDNS and many more guidelines are going to be released soon. Next slide, please.

We now have a contractor on board to help us in designing and developing the new website, probably going to look something like this, but we are at the very early stage of [helping] them to find out the final layout and the final design of this. But really that is going to be the very next outcome for this project in general. Next slide.

So, we have had to adjust our timeline a little bit since the last time we made this presentation for various reasons, but our target now is to be able to launch this by the end of the first quarter of 2022 that goes to end of March approximately.

Maybe we may not have all the functions of the website but we will launch the most critical one which can allow us to start running with this, start unrolling, and watching operators on this.

We will update the community and the mailing list as we go. If you are interested, I recommend that you join the mailing list. It is an open mailing list. Anyone can join and provide us feedback and contribute to the discussion of that.

In the meantime as well, we have set up a Wiki page where we are publishing most of the things that we are working on now as a temporary repository. Those are going to move into the formal website when we [add them].

I think that is the last slide. Yes. Thank you, everyone, for your attention. We would like to hear from you any questions, comments, suggestions. This is what this session is for. Thanks.

KINGA KOWALCZYK: We have a question in the Q&A pod. Adiel, would you like me to read it out loud?

ADIEL AKPLOGAN: Yeah, please.

KINGA KOWALCYK: So, we have a question from Sivasubramanian. "Would KINDNS also develop a shared common commitment, not to bind users to one resolver but maintain multiple redundant resolvers that anyone could use? A [rough], inaccurate, somewhat farfetched scenario. Is that [inaudible] university blocking students from using external resolvers or the Mayor insisting that everyone in town do not use another city resolver? This and other relevant shared commitments apart from sharing of best practices."

ADIEL AKPLOGAN: Thank you. Very interesting scenario. The questions are pertinent and relevant to the registrant part of this meeting, the user part. And most policy related because these are policy decisions when somebody asks their customer, their user, their community, their resident to use a certain operator or a certain service. And policy aspect are not really covered here in a certain way because we have very little control on those things.

From the registrant perspective and running operations, of course you can set up several resolvers in your set up and that is a general best practice for network management, general best practice for service providing or resiliency and redundancy in the DNS service in general. But we haven't focused this on that aspect specifically because it's more on the registrant and ISP best practices side, not the core of the DNS operation, per se.

KINGA KOWALCZYK: We don't have any other questions. Please, everyone, if you would like to ask a question, raise your hand and we will unmute your microphone and give you the floor or type in your question into the Q&A pod and I will read it out loud.

ADIEL AKPLOGAN: I would like to hear from participants about one aspect of this. Phil mentioned it. I think that was on the slide talking about the public resolver where we highlight the privacy consideration which is an aspect that we have discussed a lot internally.

As you all know, DoH and DoT at the beginning were very controversial, has raised some concerns at the beginning, but more and more resolver operators are implementing that and it's a good practice for privacy in general.

The question is how much of those privacy considerations should be part of KINDNS? When you look at usual best practices documentation, until last year, most of them … The privacy consideration is not highlighted enough. But over the past few years, privacy has become a very important consideration for users online.

So we felt like privacy should be considered in KINDNS in general. Qname minimization is across the board, both from resolver perspective of them has to do this, it's straightforward, it's most of the software so I don't think there is an controversy on that. I don't see any.

However, on DoT or DoH, we have here sometimes people raising their eyebrow. "Oh, are you going to touch on that? Are you going to mention that as best practices or not?" On the mailing list, there hasn't been hard proof [inaudible] until now on those two, particularly on DoT, but I would like to hear what people may think about privacy, from privacy consideration point of view and DoT and DoH. There is a hand raised.

Ulrich has his hand raised. Can somebody give him, open his mic so that he can speak?

ULRICH WISSER: Hello. This is Ulrich. Yes. Thank you, Adiel, for allowing me to speak here. I wanted to say that I think … DoT solves some of the privacy considerations, because obviously it doesn't allow people [inaudible] to listen to your request, but you still have to put a lot of trust in your operator. That is obviously only can be solved with Oblivious DNS but I don't think that Oblivious DNS is now ready for being anywhere near best practice.

ADIEL AKPLOGAN: I agree.

ULRICH WISSER: But I think that is something that might should be mentioned, that DoT solves some of the privacy problems but it's far from solving everything.

ADIEL AKPLOGAN: Thanks. That's good feedback from an operator perspective. You mentioned something which is important and interesting as well. The trust relationship that already exists somehow between user of resolver and provider. For instance, in the context of ISP, it is obvious there is a service agreement. There is some constraint there, which may make the privacy consideration move in some sense or if you are [inaudible] where you have a policy that makes you use their resolver of your company, the privacy aspect is kind of toned down a little bit because anyway you are using your corporate network, so you have to abide to certain policies.

But when you are in the wild, in the open, maybe that's where the privacy becomes more critical and that's why we kind of put the privacy consideration—we're trying to put the privacy consideration more highlighted in the category that covers the public and open resolver side where anyone can decide to use any resolver and that's where they may be more concerned about the privacy and what they are signing to the resolver.

Anyone else have something to share on the resolver aspect or on the [inaudible] of course as well on KINDNS and the best practices we are exploring?

Okay. If there is no other question or comment, that means that we are in … I read it that we are in the right direction. I would like to thank everyone who joined the mailing list and provided input so far to the initiative. I think there is one question. Oh, there is one question in the pod. I missed that. Could you read it?

KINGA KOWALCZYK: Yes. So, is measurement a part of the KINDNS design, a form of measurement where the community has ways of assessing the resolver data maintained by various resolvers and the method of comparison? Is the measurement a part of the KINDNS design?

ADIEL AKPLOGAN: Well, as I mentioned, we will identify an indicator that allows us to see how this is impacting some security aspect of the DNS operation in general. What those indicators are going to be are not yet defined clearly. We may be also sharing some of these identifiers from other ICANN initiatives, like [ITHI], for instance or other measurements, initiatives that already exist.

But in the question here, there is a mention to assess any resolver data. I don't know which kind of data is being mead here and the relevance in this specific case.

Usually, from our perspective, when we are measuring this kind of thing, we try to measure what we can measure from our perspective, not trying to get access to private data that we don't have control of. And we will probably measure those things according and only related to the best practices that we have in our [inaudible] may be able to say, well the number of resolvers that have a privacy consideration [inaudible] all this way over the past month, year or so by measuring those of the resolvers who have DoT or DoH or Qname minimization active. That is a way we can measure that directly.

But I don't know what you mean by assessing the resolver data, because when you start talking about assessing the resolver data, that gives me a completely different view on that. We will measure what we can see publicly, but only related to the best practices that we have previously [adhering to].

Also, we want to engage and get as much resolver as possible to join the initiative and commit to those best practices. If, in the course of that, we manage to have a way of cooperating and working with some of the resolvers to do some advanced measurement and advanced study on what they are seeing, of course that can be added to the initiative.

But again, we will make this as straightforward as possible, so that what we are presenting doesn't pose or present any controversy to both user but also people who are going to [inaudible] the outcome of the initiative.

Okay. The mailing list is open. Feel free to join and we can continue some of these discussions there. You can bring up other considerations or questions you may have there or you can reach out to us directly at OCTO. You can write directly to OCTO@ICANN.org as well if you have a direct question.

I see one question from Desiree. Thank you, Desiree. Yeah. As I mentioned, there is a communication and outreach component of this that is pretty much highlighted as well. Of course we work with the community in general. But that's why we have an ambassador program hidden somewhere there where we will work with the community to promote this. But ICANN as well will put some resources in communicating and promoting the initiative as we announce it is launched. Of course the success will depend heavily on the promotion, perhaps, and how we make this easily accessible and consumable by the community.

Thank you, everyone, for joining. We only have nine minutes left in the time allocated to this. I'm not seeing any further questions, but hopefully happy meet you all on the mailing list offline to discuss this. Anything else that, Kinga, that we need to talk about? Otherwise, back to you.

KINGA KOWALCYZK: No, we answered all the questions. Thank you, everyone. The presentation will be published also on the website on the ICANN73 website in probably a couple of days. Thank you.

ADIEL AKPLOGAN: Thank you, Kinga. Thank you, Steven. And thank you, Phil, as well. Bye-bye, all.

PHIL REGNAULD: Thank you.

[END OF TRANSCRIPTION]