ICANN86 Seville | PF – ccNSO: Tech Day (1 of 2)
Monday, June 08, 2026 – 10:00 to 11:15 CEST

CLAUDIA RUIZ

Hello and welcome to the ccNSO Tech Day session. My name is Claudia Ruiz and I am the remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Community Participant Code of Conduct, the ICANN Expected Standards of Behavior, and the ICANN Community Anti-Harassment Policy. Please observe the following guidelines to participate in this session. They will also be posted in the chat for your reference.

During this session, questions or comments submitted in chat will be read aloud if put in the proper form as noted in the chat. Interpretation for this session will be in English and in Spanish. If you would like to speak during this session, please raise your hand in Zoom. When called upon virtual participants will be given permission to unmute. On-site participants will use a physical microphone to speak.

Please state your name for the record and the language you will speak if speaking a language other than English, and please speak at a reasonable pace to allow for accurate interpretation. Thank you. And with that, I will now hand the floor over to Eberhard Lisse, chair of the tech working group. Thank you.

EBERHARD LISSE

Good morning, everybody. Can you move to the next slide, please. Welcome, everybody. My name is Eberhard Lisse. I'm the ccTLD manager of .na and I am the chair of the technical working group. Today, the Tech Day has got some tech issues, but what can we do?

As usual, we quickly go through the presentations. We have only got two sessions because it's the shorter policy meeting, and the first block is devoted to DNS abuse. I liked in particular to have the .id presentation that I've seen last time to be repeated here on the technical day because we are all faced with DNS abuse and the more tools we find that can be used the better.

Then some evidence-based checks and we have heard from Naveed Bin Rais a few meetings ago about the DNS explore platform so they will do a short follow-up. Next slide, please, Claudia. Claudia, can I have the next slide, please? Thank you.

And then we will have a short presentation or follow-up of Steve Crocker. He has presented about his RDAP project before. Then because many of us use virtual machines, we've come across what's called Proxmox, a virtual machine open-source system to spin up virtual machines. It's used by smaller TLDs and by smaller entities. I see the Austrians shaking their heads at the moment, they use bigger things, but for smaller things which is what we're doing here often is a good way of starting.

Then CENTR has had some work done on Registry Lock, Kristian Ørmen from Sweden will present. Then we have a very interesting topic, how to do the DNSSEC at scale. We all do DNSSEC, we all

suffer from it. Even the big TLDs like Germany, they had an issue the other day. So, it's always very interesting to hear how to do this on a larger scale. And then at the end, Kumari will talk a little bit about BitSquatting. We have session chairs and the first session chair is Régis Massé sitting there. Take it away.

RÉGIS MASSÉ Thank you, Eberhard. Hello, everyone, welcome to this first part of the Tech Day in this beautiful city of Seville. So, thank you for the ICANN staff for making things work for this first session and I will immediately pass the mic to Mr. Hendraswara from PANDI who will talk about DNS abuse mitigation tools.

ERY PUNTA HENDRASWARA Good morning. Thank you moderators and also the chairs. So, excuse me, have you hear my voice clearly? Okay, thank you. Okay, thank you, Chair and also moderators. So, I'm Ery Punta Hendraswara. I am the Vice President for the Administrative and also Technical Operations for PANDI.

PANDI is the ccTLD of Indonesia. The name is PANDI so we are the .id registry. So, right now we have one point four millions somethings for the domains under management per last December of 2025. We are a non-profit, so we are a legal entity since 2006 and we manage the ccTLD since 2007, and also re-delegated by IANA since 2013.

So, we are a multi-stakeholder governance. We are representative from the multi-parties and we have funding and independent members, and also government, internet industry, and also academia. We also have registrar ecosystem. It's on the government's military net ISP, internet service provider, and also the second level domain IDN, which is new, based on the new characters. We have many registrars under us right now, it's local Indonesia, and also work with the global. Okay, next.

Well, you know, ID adoption now is becoming more and more since, not just representing Indonesia, but actually ID is now also associatively with identity. So, in a global use, global world will associate ID as an identity. So now we have a growth in the number of domain under management and also the other problem that we will have beside the growth of the usage, there's also a growth on the abuse.

As we can see that during the 2025, we have like 26,650 reports identified and most of them actually more on the online gambling, which is against our Indonesian law. That's the dominant category on that. And also, the other thing is actually on the phishing is more than 40%. So, we can see from the pie chart that beside the online gambling is also phishing and also malware. And malware, you know, the damage is also quite big because of the abuse on this. Okay, next.

Well, as I said before, that not just ID is only for the Indonesians, but PANDI's mandate as a ccTLD, we have to satisfy and also follow a

global standard and also the national legal compliance, which is the things like from the technical perspective from the DNS abuse. We have a term like phishing malware, botnet, pharming, and spam as a factor. And then from the national legal compliance, we have some content unlawful under Indonesian regulation and government instructions. It's basically on the online gambling content, fraud and other unlawful content, and also material restricted by competent authorities.

So, we have our policy on the Section 6.4 on PANDI compliance and policy that we have a coordination and follow-up with the Ministry of Communications and Digital Affairs. Its name is Komdigi and other competent authorities like for the IPR allow. And also, how to suspend and also unsuspend as a documented process. So, basically ID is following two big layers of obligations like from the global and also from the national perspective. Okay, next.

And then how actually are we trying to fight against the abuse? So, there is some innovation that we try to do, which is we call it the IDADX, or Indonesia Anti-Phishing Data eXchange. It's a funding system for handling .id domain abuse. So, we pulled some reports from the registrant and also registrar. And also, we handled the case and also published in the idadx.id if you want to try to have some look into the sites.

So, how is the process? First, we collect report and data on the .id abuse so it's pulled from many sources, including the internal crawling. And then the second one is trying to correlate, which it is

matched with the evidence or maybe the indicator is matched that actually is against the registry record. And then third, we will actually send automatic notify to the registrants, which is, if their domain is actually matched with the criteria for the abuse.

And then number four is enforce. If the content removed in time, the domain is not suspended, but if they don't follow, we will suspend the domain. And then number five is we publish the handled domain transparently on the IDADX. So, the automatic notification from the IDADX, first one is warning plus take down instructions. Then second one is actually a suspension notice if they don't follow the compliance.

So, how to recover? So, registrant can request actually the unsuspend or domain-wide listing via the PANDI form or the help desk at pandi.id, but it should be following the compliance based on the global regulation and also the local and in Indonesia law. Okay, next.

And how actually we doing that? So, we are not doing it independent, is not doing it alone, we are doing it with the collaboration with the multiple ecosystem here. First with the government. Of course, we get the support from the Komdigi as the national authority and then the second one is from the national security coordination is like the National Cyber and Crypto Agency. The BSSN is Indonesians regulatory and also Indonesians body. And then the external sources is also coming from the global like Netcraft, CleanDNS, and TNN and also the internals PANDI

innovations, we call it BIMA. It's PANDI internal crawling engine. So, BIMA is for the breach identification and monitoring assistance. So, it's spelled BIMA, it's not BIMA.

BIMA is the one that detecting the pilot thing and also the one that doing the crawling. And also give the finding directly into the IDADX concept with the external and also government input for the manual maybe or also maybe some report that we get from the respected parties. Okay, next.

A bit technicals, but this is handling workflow. So, report receive and until then verified correction. So, this is the category that we have. Auto-suspend so this is for the abuse criteria on the my.id, biz.id, web.id. So, if the content is not following the compliance there will be auto-suspend. There are also some other suspend with the negative content. Then manual suspend, so this one is actually extension other than the my.id and also biz.id then web.id that having a negative content actually it's a manual suspend.

Okay, basically that's all that we have from PANDI, how we fight the DNS abuse so hopefully this sharing can help and also bring more to the community of the ICANN and also to the others. Thank you from PANDI.

RÉGIS MASSÉ Thank you for the presentation. We've got time for one question if there is one in the room. In the chat there is one, okay. I will read it in the chat. How does IDADX handle false positive and data

validation where aggregating abuse intelligence from multiple sources? And what safeguards are in place to avoid disrupting legitimate domain operations?

ERY PUNTA HENDRASWARA: Well, actually, there is some verifications on that to avoid maybe the false positive. So, that's why on how we handle, some of them is automatically if match the indicators. So maybe some false positive might happen there, but that's why there is some verifications on that especially for the non -- We actually only three second level domain that actually having auto-suspend like my.id, biz.id, and web.id, but for the others actually it's not an auto-suspend. There is some verifications states on that. That maybe answered that question. Thank you.

RÉGIS MASSÉ: Thank you so much for this presentation. So, now the second presentation is evidence based on DNS abuse mitigation. I'll let the floor to Kroopa, please.

KROOPA SHAH: Thank you. I think you can all hear me. So good morning, everyone. I hope you've all had an uneventful trip to Seville. My name is Kroopa and I run the registry services business at Identity Digital. I'm here to give you a small update on our DNS abuse mitigation activities, specifically those that are evidence-based. Next slide, please.

Before we get started, a little bit of information on Identity Digital. We are a registry operator for over 275 TLDs. These are TLDs that we own and operate and we provide the infrastructure for over 200 others. So, these are ccTLDs and gTLDs including .org, .au, .ai and others. This includes the registry, EPP, DNS, DNSSEC, registrar support and for many, DNS abuse mitigation activities. We support over 35 million domains under management across all the TLDs on our infrastructure, sponsored by over 1,800 registrars and that's, again, across all the TLDs on our infrastructure. Next slide, please.

So, let's take a look at what the process looks like at Identity Digital to act on DNS abuse. At a high level, a reporter reports an instance of DNS abuse through a web form on the identity.digital website or through an API. And DNS abuse is, as many of you are aware, phishing, botnets, malware, spam, when it is used as a vector for other forms of abuse. And so, reporters can report that through the web form on our website or through an API.

The API is JSON-based, makes the reporting a little bit more consistent, makes it a little easier. When that report is received by our DNS abuse team, it is investigated. And if we deem that we have sufficient evidence, then we escalate that report to the sponsoring registrar for review as well as escalation or action.

The registrars maintain the relationships with their registrants and so the report is referred to the registrar if we deem that we have sufficient evidence to proceed further. If the registrar does not act in a timely manner, that's generally 48 hours of the escalation, then

the registry acts on it and that action essentially means applying a server hold to the domain name and pulling it out of the zone to prevent further resolution.

Now the registrar can appeal the decision. So, let's say that they were unable to respond in a timely manner, but they have more information to demonstrate why that name should not be suspended, then they can submit it to us. We will review it, investigate, and then again depending on the information provided, we may unsuspend or you know remove the server hold from the domain.

Now there are a few caveats to this process that I'll call out. For example, if you get a report from a trusted reporter on an instance of CSAM or something similar, violent imagery, then we will immediately suspend the domain, right? We will let the registrar know, but we will immediately take action. Next slide, please.

So, I talked about evidence that gets reported into us when reporters report or when reporters provide complaints of DNS abuse. And here are some examples of acceptable evidence that we have used to act on domains in our portfolio. Screenshots of some of the pages here. On the left is the infamous toll road message. If you're in the United States, at least, you've been the lucky recipient of a message like this, you know, over SMS asking you to pay your toll bill or else your account will be suspended. And then the two screenshots on the right as well are attempts to have the user enter information.

So, these are some of the types of evidence that get reported to us that we are able to use to investigate further and then act. And again, that action means either reporting it to the registrar or if you've already done it and there's no action, then the registry will go in and remove the domain from the zone by putting a sort of a hold in the registry. So, what has that approach led to? Next slide, please.

In the last quarter of 2025, we've reviewed over 85,000 domains that were reported for a variety of different reasons. In total, in 2025, we reviewed over 590,000 domains for reports of abuse. And you'll see that the largest category here is phishing. Most of the domains reported are for phishing. We received over 66,000 reports of phishing for over 21,000 domain names.

You'll also see, if you look at the last category there, it says other, and that represents, you know, over 66,000 reports. And this other category is sort of a, you know, bucket for things like third-party reports, maybe it came in through email, it wasn't reported through the website, it was on a block list that got reported. There's a number of other categories that are accounted for in other that's listed here and this is what we've been able to do.

Now, this information is publicly available through Identity Digital's anti-abuse report. We publish a quarterly anti-abuse report where we talk about what we've done for the past quarter, as well as interesting trends that might have come up. So next slide, please.

In our last anti-abuse report, I'll call out that there's more information, but there's also some tidbits about bitsquatting, if that's something that you're interested in. I know that there is a presentation here or in this forum about bitsquatting a little bit later today. So, I'll just walk very quickly through some learnings that we've had as an organization regarding DNS abuse. Actually, before we go into that, let me just walk through the action timeline and what some of these images that you see on the screen mean.

What we've seen is that within the first 24 hours, there are some registrars that take action before they receive an escalation from the registry. So, what this means is that the reporter who reported the domain to us, they may also have looked up the sponsoring registrar. They may have gone to the sponsoring registrar's website. They may have reported the domain to them as well. So, by the time, you know, that first 24 hours passes, we've seen that some of the registrars, they take action proactively on their own. And in the last quarter of 2025, that represented 541 domains, 559 cases reported to us.

Within that first 24 hours, the registry took action on over 1,600 cases, or I guess over 1,800 domains. This is where we have evidence of abuse. We feel pretty confident that this domain is used for abuse. We put a protective hold on it. It's been reported to us by a trusted source, a trusted notifier through our network and so we've put that hold on it proactively and that represents over 1,800 domains in the last quarter of 2025.

Between 24 and 72 hours we saw that the registrant took some action to remediate the case of abuse or the complaint. And there were 158 domains in that category where the registrant did something or some corrective action, which meant that we didn't have to take further action on that domain name. And then we also noticed that there were over 1,500 domains where the registrar reached back out to the registry and provided us with more information on the actual complaint itself, which means that we didn't take any action.

So, when the registrar responds and says, “Well, here is more information to demonstrate why further action doesn't need to be taken,” we review that, we investigate that. If we deem that it looks good, then we’ll get back to the registrar and let them know, but we will also not take further action on the domain name. And then separately, we had over 1,300 domains, again, this is the last quarter of 2025, where the registrar took an action once the registry reached out to them. So, once we reached out to the registrar and let them know, sent the evidence that we had received, then they took action.

Now, generally, we give the registrars 48 hours to take action, and then after 48 hours is when the registry comes in, takes a look, and if the registrar hasn't done anything, then the registry will go in and again suspend the domain, put a server hold on it, remove it from the zone file. So, in the last quarter of 2025, the registry took action

over 2,200 domain names, and we also had a number of compromised domains reported.

Now, for compromised domains, you know, where there is legitimate, you know, users, legitimate websites, or the domains are being used for legitimate purposes, they're not being used for DNS abuse. But for some reason, you know, they've been compromised, and so they're being used as a vector for DNS abuse. In that case, again, we communicate back with the registrar. We don't go ahead and suspend the domain. That would be impacting a legitimate registrant quite negatively.

Our full year 2025 stats, we saw that there was less than a 1% abuse rate in all of the TLDs that we managed. Again, we support over 35 million domains under management. That's across all of the TLDs on our infrastructure. For our own TLDs, these stats represent the actions that we took on TLDs for which we are the registry operators, TLDs that we own and we operate.

I want to talk a little bit about some of the key learnings, at least from our perspective, on this approach that we have taken with regards to tackling DNS abuse. The first is enabling reporters to provide as much information as they can up front really helps us tackle DNS abuse more effectively.

So, when someone sends in an email saying, “This domain is bad,” that doesn't really help us, but we have a web form on our website that ask for very specific information, the descriptions, a full URL, evidence, files, all of the things that we require to investigate and

do things. And the fact that, you know, we have that in place has made it that much easier for reporters to report DNS abuse and that much more efficient for our abuse team to review and take action in a timely manner.

The other thing that we've noticed is, again, because we are managing over 470 TLDs, we've been able to observe abuse activity in one or patterns of abuse in one TLD or a set of TLDs. And we've been able to take action there, see that abuse move to, say, another set of TLDs, but act on it proactively, act on it quickly before it does more harm in a larger way. And so that's also helped us tackle DNS abuse a little more effectively, a little more proactively.

I'm not going to go ahead and say, oh, we can all prevent DNS abuse. I don't think that we can prevent DNS abuse as such. What we can do, though, is control the effectiveness and the speed of our response to it. And that's what we've seen, is that the persistent crippling of abuse, the persistent actions has resulted in malicious actors eventually giving up and moving from say our set of TLDs to another set of TLDs eventually away from our portfolio to other portfolios. And with that I will take some questions if there are any in the room for me. Yes, sir.

EBERHARD LISSE Okay, Eberhard Lisse, .na for the record. I am a ccNSO appointed member on the Associated Domain Check PDP Working Group of the GNSO. So, I'm interested, do you do any associated domain

checks yet? For example, if you see a batch of domain names registered same day or same credit card details and whatnot.

KROOPA SHAH That's a great question. So, we don't necessarily capture credit card information from registrants because, again, our interaction is with the registrar so we don't necessarily do those on the credit card. What we do see is that if there's a pattern of registration, sometimes it's just what looks like random strings, they don't seem to have any real meaning, then that's something that we will review and investigate more proactively.

Sometimes we see that the names are registered, but nothing has really happened. But if we see that there's a pattern of names that's coming from a specific registrar and some of those names are being used for, say, DNS abuse in, say, the next day, the next week, et cetera, then we'll evaluate the whole portfolio. We'll evaluate that whole group that came in, or the whole batch, and then refer it back to the sponsoring registrar.

RÉGIS MASSÉ Thank you. Another question in the room? I don't see it on Zoom that's someone in the room? Okay.

CLAUDIA RUIZ We have a question in the chat.

RÉGIS MASSÉ: It just arrived. Thanks.

KROOPA SHAH: I'm just reading the question, sorry.

EBERHARD LISSE: I can read it. Enoch from Malawi Youth IGF question, "Are abuse rate metrics like 0.4% standardized across registries? Or do difference in detection methodologies affect comparability between TLD operators?" Before your answer, there was this Interisle report which had immense higher figures, 10 times higher figures, but they just looked at block lists. So, the question is interesting.

KROOPA SHAH: I haven't read the Interisle report, so I won't comment on that, but that's a really good question. The 0.47% rate that I mentioned was kind of average across all of the TLDs that we operate naturally, and sort of the average, so to speak. It's not for one specific TLD. It's slightly higher for one TLD, while other TLDs may have no abuse that are in our portfolio. So, it's really the average across all of the TLDs in our portfolio, again, TLDs that we own and operate.

So, I guess the short answer is yes, different detection methodologies, they will impact the comparability, you know, in abuse rates amongst different TLDs across different operators. Thank you.

RÉGIS MASSÉ Kristian.

KRISTIAN ØRMEN Thank you. This is Kristian from the Swedish Internet Foundation. You mentioned that there was 1,500 of these reports that went to registrars and they got back to you and you didn't then suspend the domain. So, I WAS just curious, are these 1,500 potential false negatives, or is it because they cleared the abuse or fixed the abuse and got back to you?

KROOPA SHAH That's a really good question as well. So, I think it's a mix of both where the reporter came to us and said, “This domain is abusive, here is some evidence.” We looked at the evidence. We have specifically defined thresholds and based on those thresholds, if we deem that, yes, this warrants further investigation or an escalation to the registrar, we'll go and we'll report it to the registrar.

Now, if the registrar came back and either, like you said, they fixed it on their own or they provided an explanation where we deemed it meets our criteria, then we would go ahead and not take any action on those domains. Right? So, does that answer your question? Thank you.

RÉGIS MASSÉ Okay. Thank you, Kroopa. Big applause for Kroopa, please. And now the next presenter is normally on remote. Do you have connection with…? Yes, he's on. Perfect.

AMREESH PHOKEER Hello. Good morning.

RÉGIS MASSÉ Okay, so we now talk about evidence based on DNS and Mr. Phokeer, you've got the floor.

AMREESH PHOKEER Thank you very much. I hope you can hear me okay.

EBERHARD LISSE Yes, we can.

AMREESH PHOKEER Thank you. All right, so thank you for having me remotely today. Unfortunately, I couldn't travel, but some of my colleagues from ISOC are on site and probably even in the room. So, any questions I'm sure they can answer or I've put my email at the end so that you can also interact with me, but happy to take also questions at the end.

So, my presentation is about evidence-based DNS resilience. The previous presentation was also evidence based as you can imagine having accurate data and facts about how the internet and the DNS

in this regard is functioning is actually important. So next slide, please.

So, no need to tell you about how important the DNS ecosystem is to internet services and applications nowadays and generally, of course. And on top of that, there have been many regulations such as the NIS2 and the CSF2 that are actually putting more emphasis on those critical infrastructures and, you know, putting regulations and directives to enhance the security of those applications. So it is in this regards that we at the Internet Society, we think DNS is a critical component of the Internet and we thought that measuring how resilient it is, is important.

And in thinking of how to measure the resilience of the DNS, we thought that, oh, why don't we use existing frameworks? So, there is that framework called Knowledge-Sharing and Instantiating Norms for DNS and Naming Security, which I'm sure you're very familiar with, KINDNS, provides a set of measurable and non-measurable practices that operators need to put in place to be able to make the DNS more resilient generally. And it is using the KINDNS backdrop that we have run this. we are actually running this measurement study. Next slide, please. I can also share my screen if that is more convenient.

EBERHARD LISSE No, carry on. It just takes a moment.

AMREESH PHOKEER

Okay, thank you. Right, so in terms of resilience metrics, of course, as I just mentioned, KINDNS as a framework has both what we call measurable metrics and non-measurable metrics. Because they are mostly associated with some organizational processes that need to be put in place by either authoritative DNS operators or recursive DNS operators to enhance the security of the overall DNS ecosystem. So, for us to be able to provide some statistics about the overall health, we therefore focus on mostly the measurable practices of both the authoritative DNS side and recursive side. Next, please.

Yes, so KINDNS is itself broken down into two parts, the autorotative DNS resilience part and the recursive resilience part. I will now talk about the autorotative DNS resilience measurements. Excellent. So, the measurable practices for KINDNS, there are more than six practices, but the one that I've put here are the measurable practices.

So, practice one talks about DNSSEC and key management, which means that operators of top-level domains and secondary-level domains need to put in place the necessary framework for assigning their zone and also do proper key management. For practice number two, they need to limit a zone transfer, which means transferring the information from one zone to another. And practice number four is the authoritative and recursive servers need to be on different hosts, meaning that if you are operating an

authoritative server and also a recursive server, they should technically be on distinct and different networks.

For redundancy, practice number five says that you need to have two distinct name servers and practice number six provides some recommendation on software diversity. You might better have different software coming from different vendors. Network diversity, meaning that you need to place your name servers in different networks and not put everything in one network and geographic diversity as much as possible, put your name servers far from each other so in different geographic regions. Next please.

So, what are we exactly measuring? We looked at the number of name servers. This is to protect against, for example, node failures. We look at the number of IP addresses that are serving those authority name servers. We look at the number of ASes that are hosting those name servers and also, we look at whether those name servers are using anycast addresses. As you know, many name servers are already using anycast addresses and this can provide protection against, for example, site failures or DDoS protection as well.

On the right, as you can see, we are using a variety of datasets and most of them are open datasets, so datasets that are provided by either universities for their research programs, or it could be also from companies such as IPInfo for geolocation. We also look, therefore, at the geolocation aspect. So, IPInfo again provides us quite useful information about geolocation. So, put that all

	together, it gives us some data points that are allowing us to actually measure, to some extent, the authoritative resilience of the DNS. Next slide, please.
RÉGIS MASSÉ	Amreesh, did they reach you? We have a power failure here.
AMREESH PHOKEER	Oh no. Oh, I'm sorry to hear.
RÉGIS MASSÉ	Okay, we are on Zoom.
EBERHARD LISSE	What's the consensus, Régis? Shall we carry on, on Zoom?
RÉGIS MASSÉ	Okay, just carry on. All participants are supposed to have Zoom running so they can listen.
AMREESH PHOKEER	Okay.
RÉGIS MASSÉ	Yeah, carry on.

AMREESH PHOKEER All right, I just need to move to the next slide.

RÉGIS MASSÉ There you go.

AMREESH PHOKEER Okay, thank you. So, in terms of data collection, we wanted to have a global view of all TLDs, as many as we could. So, we used a series of different data sets put together to build that corpus of domains. So, we used CT logs. So, as you know, Certificate Transparency logs from different providers gives the domains that have certificates and when the certificates are expired or not, they are added to the list.

We have zone files from the ICANN Centralized Zone Data Service, so CZDS. We have some opendata. So, for example, there's .se or you know, reverse DNS zones from RIR will also provide some information. And also, top-lists. As you can see, there are a number of top-lists like Tranco, Majestic, Radar, that provides a set of domain names that can be useful for this study. Next, please.

I'm sure you must be familiar with the University of 20 OpenINTEL dataset. As the dataset that we have used, they are also collecting from using multiple data sources. And right now, they have around 308 million domains that they are currently getting information from and we are actually working with them to actually use their dataset because they refresh it on a daily basis. We haven't yet

implemented this dataset, but we are in process of doing that. Next, please.

So, these are some of the early results that we have. So, we, for example, took the Tranco top one million domains, and we were looking at the different metrics that we have highlighted before. For example, the number of name servers, the number of IP addresses hosting those name servers, the number of ASes, and the number of anycast addresses. So, our initial analysis was on the top, Tranco Top 1 million domains. I will give some more details later. Can we move to the next slide?

So, we do have a dashboard. So, if you can scan this QR code, you will go on a website. So right now, the website doesn't have a domain. It's connecting directly to the IP, but you can have a chance to go and see the proof of concept. So, we have a dashboard that shows, for example, the ccTLD zones and gives you, again, for each of those zones, those different metrics that we have highlighted before. We have a section for gTLD zones and you can compare one gTLD or one ccTLD to one another.

And then you can also use a feature that we have implemented where you put your own domain and your own subdomain, and then you can test and it will generate the statistics on the number of name servers, where they are located, whether they are located in the same geographic location, whether they're using anycast, et cetera. Right now, we have only tested that on a sample, but not on the big corpus of data that we have of several hundred million

domains, but rather just on less than a thousand domains. Next please.

Right. Now let me talk about the other side of DNS, which is about recursive resolver. And I will also look at different metrics that we are measuring to capture how resilient the recursive resolver ecosystem is around the world. Next, please.

So, these were features that were technically extracted from the KINDNS guidelines on how to operate recursive resolver servers. So, the first one is about being BCP38 compliance, so meaning that a network should not allow spoof traffic. This is an important practice that needs to be adopted by recursive resolver operators so that they prevent their own network to be used as a vector for DDoS, for example.

The next one is about QMIN minimization. As you know, it is important to protect the privacy of users, so QMIN minimization is an important standard that should be implemented in that regard. The next one is about MANRS compliance. So, MANRS is the Mutually Agreed Norms for Routing Security. Again, it's a set of best practices that a network operator needs to put in place to protect its own network, but also other networks. And in MANRS, there are a few things such as whether your network is doing what we call RPKI so routing public infrastructure to protect the internet routing. And also, whether it is doing global coordination in terms of putting the right information in the databases for coordination for routing.

We also talk about software diversity. Are you using different software for your resolver? Is your host only allowing DNS traffic and also importantly is your resolver doing DNSSEC validation to protect users from DNS cache poisoning? Next slide, please.

Again, we find very similar practices with the authoritative DNS part where we talk about geographical diversity. So, if you're operating a set of resolvers, do you have geographical diversity? So here we are mostly talking about open and public resolvers that usually are using anycast. Topological diversity, again, this means having the resolver in different networks. What are the caching best practices? Anycast, I just talked about that and one, I would say, bonus information that we can get from all the measurements that we are looking at is also the consolidation effect of maybe resolvers. So, are we seeing most of the queries going to a subset of resolvers rather than having it very well distributed? So, this I would say a bonus that we can get from this measurement study. Next please.

So, for our study, the data sources that we have selected are four, so mainly two that are looking at the spoofing side of things. So, we have a project we are working in collaboration with, TU Dresden in Germany, where they are hosting a project called OpenDNS API. And using a technique with transparent forwarders, they are able to understand in which network there is spoofing happening.

Another project which used a different technique called CAIDA Spoofer basically needs the one running the measurement to actually download an application and run this application. And this

will basically send some feeds back to CAIDA, telling whether spoofing is allowed in this network. So, put together, those two projects gives us relatively good information about spoof, whether spoofing is allowed in a network or not.

We are using OpenINTEL data and this OpenINTEL data, as you know, it provides us information about domain names and real-time zone feeds and to some extent is helping us in our measurement. And for DNSSEC validation, we are using the APNIC Labs data. So APNIC Labs have their own measurement campaign they use using Google Ads. Using that, they are able to give a percentage to each network they are able to reach on whether they are able to do DNSSEC validation. Next, please.

RÉGIS MASSÉ Amreesh, excuse me. One or two minutes left to finish to let time for the last presenter, please.

AMREESH PHOKEER Okay, sure. So, in that case, I will just skip that slide, I will go straight to the report. So, in our preliminary analysis, what we have seen for the open resolver reports, 93% of ASes with open resolvers have been observed to host less than 10 open resolvers and we have observed that a single open resolver per AS for 46% of ASes. So, there is a single open resolver for 46% of all ASes. And we have observed 28,000 open resolvers in almost 5,500 ASes in 183

countries. So, this is to give you a global overview of how the open resolver distribution is around the world. Next one, please.

In terms of closed resolvers, so as you know, there are open resolvers and closed resolvers. Closed resolvers are usually hosted behind ISP networks and usually only available to ISP clients. We have observed in 94% of those ASes with closed resolvers, they host less than 10 instances of an open resolver. And we observed only a single closed resolver per AS for 48% of those resolvers. And in total, we have observed 72,000 closed resolvers in 9,000 ASes in 212 countries. Next please.

So, as I was mentioning that the OpenDNS API was using transparent forwarders for the spoofing networks, so the number of ASes that are allowing spoofing that you are able to track are more than 2,000. And CAIDA is providing almost 700 ASes that are providing spoofing. So, this is the data that we are working with for spoofing. Obviously, it is not global because we only have a subset of this information, but it is allowing us to have some amount of information. Next, please.

So again, these are just like a global overview of which countries are doing the most, the highest level of spoofing. For example, we have Brazil. We have found more than 1,000 ASes that are doing spoofing. In India it is around 800 and in Bangladesh it is around 150. So, it gives you an idea of which ASes in these countries are allowing spoofing, where we have actually found open resolvers.

So, there are more details in the slides. I will stop here for the sake of time, but happy to take questions now or offline as well.

RÉGIS MASSÉ: Thank you. Thank you for this presentation, very interesting. To let the last presenter to have his time, I will go directly to his presentation and after that, if you have a question for your [inaudible - 01:00:21], we will just take one or two minutes from the break to give the last presenter having his time. So, I think we can jump in the last session now.

EBERHARD LISSE: And in the agenda, the names of the presenters are clickable email links so if you wish to engage with one of them just look at the agenda and click on the corresponding name.

RÉGIS MASSÉ: So now we will talk about DNSXplore and I leave the floor to Shahzaib.

SHAHZAIB JUTT: Yes, I'm audible?

RÉGIS MASSÉ: Yes, go on.

SHAHZAIB JUTT

Thank you for having me. So, I will be presenting DNSXplore. So, DNSXplore is basically a platform for DNS archival history. So basically, what we do is, we do the DNSSEC validation after every specific interval, like two or three weeks, store all the DNS exchange for all the domains and then we present the history. Can I have the next slide, please?

JOKE BRAEKEN

Apologies, a request to the speaker, could you please be closer to the microphone? Thank you.

SHAHZAIB JUTT

Yeah, sure. So basically, this is our team. We have presented at ICANN84. We have received feedback that we can improve the validation methods, how we can do the validation. So, we have acted upon that feedback and we are here again. Can we have the next slide, please? The next one.

So basically, I would give you a recap of what is DNSXplore. As we get all the data, we get the domains from all the zone files from the CZDS data, we do the validation for all the domains in the zone files. We store that data and we provide the DNSSEC stats as well as individual validation for all those domains. Next slide.

Regarding the feature, currently we have around 250 million internet domains and our crawl rate is basically 10 million domains a day. Currently, we haven't reached a milestone of one year history, but currently we have done two to three iterations and we

	have that data in the history and we are providing that data for the community. Next slide.
JOKE BRAEKEN	Thank you. Apologies to request to the speaker again, the audio quality is not so good. Could you perhaps try to take your earpods out and try to speak into the microphone of your device directly?
SHAHZAIB JUTT	Yeah, sure, I'm sorry for that.
JOKE BRAEKEN	Thank you, let's give it a try.
SHAHZAIB JUTT	Yeah, so is it better now?
JOKE BRAEKEN	Thank you. This is much better.
SHAHZAIB JUTT	So, [inaudible - 01:04:25].
JOKE BRAEKEN	Apologies, the audio quality is worse again. Could we go back to…?

SHAHZAIB JUTT Yeah, sure. Anyway, is it good now? Sorry for the [inaudible - 01:05:00].

JOKE BRAEKEN Thank you. We hear you loud and clear now.

SHAHZAIB JUTT So basically, how we do the validation is we get the domain from the Zoom file, we [inaudible - 01:05:13], we do the whole chain validation from [inaudible - 01:05:20].

EBERHARD LISSE Sorry, we can't understand almost a word of what you're saying. Please try to use the microphone from your earplugs again, please.

SHAHZAIB JUTT Hello? This is some quality have improved or is it still the same?

EBERHARD LISSE Noise is in the background.

SHAHZAIB JUTT Okay, just give me a minute. And so, I think I have [inaudible - 01:07:03]. Is it good now? Can I have the next slide please?

RÉGIS MASSÉ Sorry, we ran out of time and we can't hear you.

SHAHZAIB JUTT Oh, I'm sorry for that.

RÉGIS MASSÉ Yeah, we've got your email on the site, I think. So, if there are some information or questions after the session, I think we will join you. Thank you for your part of presentation. Sorry for that, but I think we have to stop there.

So, big applause for the presenters and remote, please. Yeah, thanks. Thank you. Okay, so that's the end of first part of the Tech Day this morning. Good luck, Jaromir for moderating the second part. It was not very easy this morning, but okay. See you at 11:45 for the second part, thanks a lot.

[END OF TRANSCRIPTION]