ICANN85 Mumbai | CF – ccNSO: Tech Day (1 of 4)
Monday, March 9, 2026 – 10:45 to 12:00 IST

CLAUDIA RUIZ

Hello and welcome to the ccNSO Tech Day. My name is Claudia Ruiz and I along with my colleague Susie Johnson are the participation managers for this session. Please note 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 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 include English, Spanish, and French. If you would like to speak during this session, please raise your hand in Zoom. When called upon, virtual participants will unmute their microphone. On-site participants will use a physical microphone to speak and should leave their microphones disconnected.

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. As Claudia said, my name is Eberhard Lisse. I am the chairman of the .na ccTLD and the chair of the Technical Working Group. And to my dismay or however I call it, I recently got appointed by the ccNSO to the PDP of the GNSO on how to curb domain abuse, which is very interesting from our side of things. This time we have again an interesting agenda. Can you go please to the next slide? Okay, it even works that thing.

As usual, I make some opening remarks for which I have got five minutes and then Ondrej Filip is going to chair the first session. We have for each session decided on having a chair who is sort of prepared to start engaging if there is not enough questions from the audience.

Warren Kumari is going to start an IATF-based presentation about making the local root the default. In other words, how to have a root zone on your local server. There is a lot of implications on this and that's a very interesting topic. So, when they asked, they got.

Then Harsha from .lk Sri Lanka is going to talk about security management practices. We always like to have regional and local contributors and he asked, he got. Then we have again the host presentation, Dr. Tyagi, the CEO of NIXI, has been asked and is willing to give us a presentation on what they do and how they do it.

One of the larger ccTLDs is quite interesting. They had at least two transitions of their backend in the last 10 years from what I know, and so they have got surely some interesting stuff to tell. Then we have given the location, which a country with many scripts and many languages, we thought it might be good to have a round table about universal acceptance, which is handled by the universal acceptance group of the ccNSO.

Then we have a session where Steve Crocker is going to talk about a proposed RDAP framework, how to redact and to access RDAP data.

Jeff Osborn, who I see here, has given a very nice and not too technical introduction on root servers on one of the last ICANN meeting. And I liked it so much that I thought maybe that's a good thing to do it on Tech Day again that everybody can listen why, how they do their things, and what it does and so on.

Then Jaromir Talíř from CZ.NIC is going to talk about DNS hosting with their registry product, which we all have known about because they have spoken often about, FRED. And then we will hear from a new project.

It's the next, Block 5 is all remote participants and they're all in Europe, so time, so on, we put them all together and, in the end, so that is not so difficult for them. And Namingo Registry is a registry written in PHP from scratch, which is very interesting. So, when we heard that they are willing to present, we put them on.

Then Timo Võhmar is from Estonia is going to talk about domain contact quality in the terms of DNS abuse. This is an interesting topic. And then Alexander Mayrhofer from .at is also going to talk about a topic within the DNS abuse framework. In other words, how to assess the quality of a domain name. That said, without further ado, Ondrej, you have got the chair.

ONDREJ FILIP

Yeah. Thank you very much. I think everything was said to the presentation, so I will be chairing the first session. We have three very different presentations. We will spend roughly 70 minutes split into 30 minutes for making the local route the default and the security management practices and then the host presentation.

They are very different because as Eberhard said we had to kind of reshuffle the agenda to accommodate also the remote speakers to make their presentation in a reasonable time for the time zone. So without further ado, I will ask Warren Kumari who has been working his almost whole professional life to make DNS faster and better. So, I think he has another improvement for us. Thank you.

WARREN KUMARI

Microphone. There we go. So, hi all. As Ondrej said, this is going to be a short update on the local work. And the clicker appears to have stopped. Is it outside the window? I can just share from mine. Yep, nope, no love. Should I just share mine? Or, Ondrej can try push it.

ONDREJ FILIP You think I will be more successful?

WARREN KUMARI Nope, I can just share.

ONDREJ FILIP Okay.

WARREN KUMARI There we go. So, this is some joint work with myself, Wes Hardaker, Geoff Huston, Jim Reid, but also a bunch of other people in the IETF, and it builds on some earlier work from Paul Hoffman. As with all IETF drafts, it's subject to change. So as I think many of you know, Dr. Eberhard Lisse, who runs this session, used to be a gynecologist and an obstetrician. So, I thought it was only fitting to start off with something from his area of expertise.

This is a Hungarian doctor named Ignaz Semmelweis. And in 1846, so about 180 years ago, he was hired at the general hospital in Vienna in the maternity clinic. This hospital actually had two maternity clinics, one of them run by physicians and one of them run by midwives.

When he started, he started collecting some data. And what he discovered was the mortality rate in the physician-run clinic was around five, sorry, three times higher than that in the midwife clinic. It was around 18 to 20 percent. That means basically one in

five mothers coming in to give birth would end up dying. This obviously seemed concerning to him, so he did some research to try and figure out what the difference could be. He tried a bunch of different things, copying the procedures that the midwife clinic was using.

There was also a theory that when somebody died, the priests would come along and ring a loud bell as part of the funeral rites, or sort of last rites, and maybe that was causing additional stress to the mothers. He tried all sorts of things and couldn't figure out the difference.

After a year or so, a friend of his, who was a pathologist, was doing an autopsy and nicked his finger with a scalpel. And within a few days, he died of a disease which looked almost exactly the same as what the mothers were dying from. So he came up with this sort of theory or hypothesis that there was something he called cadaverous poisons that were being transferred from the doctors who did autopsies up to the mothers.

So, he instituted a new policy where after doing an autopsy, Doctors would wash their hands, and before doing deliveries, they would wash their hands again. And fairly much immediately, the mortality rate dropped from around 18 or 20 percent, down to less than 2 percent. So this clearly seems like an obvious win and everybody should start following this, like, let's wash your hands thing before doing surgeries, deliveries, et cetera. Unfortunately, he was widely ridiculed for this idea.

At the time, everybody believed or most people believed that disease was transferred by what was called miasma or basically bad error. So, he kept pushing this idea of like hygiene is important, washing your hands before doing surgery is important, you should do this. Most people ignored him. He ended up actually getting fired from his job because his boss person didn't believe in this whole like, you know, there could be dangerous things that you transfer idea. And he ended up dying before the germ theory of disease became popular.

So why am I telling you this in a DNS presentation? Well, the reason that he was largely ignored is he was working against the sort of prevailing wisdom of the day. We've always done things this way so it must be the best way of doing things. What I'm wondering is, does this mean that we're doing similar things in the DNS? Is there possibly a better way we should be doing things?

So, a quick refresher on how DNS resolution works. The user down here wants to reach a random domain.net and they're asking their resolver, which is newly started up. So, they ask their resolver, "Where is www.a-random-domain.net?" The ISP resolver says, "I don't know, you should go and ask the root." The root says, "You should go and ask .net.1.1.1.1." The resolver asks .net, net says, "I don't know that. You should go and ask ns1.foo.net." The resolver asks ns1.foo.net, and eventually the user gets an answer.

However, users are fickle and they decide they don't actually like a-random-domain.net. So, then they decide that they want to reach

example.com. They go, if they ask their resolver, as I say, the resolver started off with an empty cache. The resolver says, "I don't know how to resolve this." It asks the root, "Where is example.com?" The root says, "I don't know that. You should go and ask.com," and the rest of the exchange happens as you would expect.

Again, users are fickle. Now they want example.org. They ask the resolver. The resolver asks the root again. The root says, "I don't know. You should go and ask.org." And the rest of the exchange happens as you would expect. User wants example.in. Ask the root. The root says, "I don't know, you should go and ask .in." You can guess what happens after that.

Wouldn't it be nice if we didn't have to keep going along and asking the root questions? If we could skip this whole, everything that starts off, you go and ask the root? You know, obviously assuming it's not in the cache, but with an empty cache, wouldn't it be cool if we could skip all of that? Well, luckily, there is already some way to do that.

There's RFC8806. This was originally published in 2020 and it is an update to RFC7706 from 2015. And this is what people generally call local root. Sometimes in the ICANN world, it's been called hyperlocal root, but this basically describes the local root concept.

At a very high level, this is basically recursive resolvers should go off and fetch a copy of the root zone and they should cache it in some way. They should have it locally. This has a bunch of

benefits. One of these is increased performance. If you don't need to send a packet to the root, your round-trip time to the root is zero. If you don't send a packet, there is no latency for sending the packet and getting an answer.

It increases reliability. If you don't need to reach the root to ask a question, you don't need to be able to reach the root actually to ask it a question. There's some privacy improvements. For example, if your user is trying to reach alcoholics-anonymous.orf and they mistype it, or alcoholics-anonymous.org, and they mistype it as orf, you're going to be going off and asking the root, what is .orf, and probably leaking the whole query.

There's also some DoS mitigation parts to this. A relatively common attack or sometimes just a mistake, as a client, we'll start asking for a bunch of TLDs that don't exist, or names within TLDs that don't exist. And the recursive resolver has to try and resolve those, sends the queries to the root. It's additional traffic, additional work.

So, what I said is, for this, resolver should go off and get a copy of the root zone, but how do you do that? It sounds like a really easy thing. To get a copy of the root, you need to be able to know what the root zone looks like. You need to be able to fetch it somewhere. Surely you can't just go off and get a copy of this root zone? Well, it turns out you can.

This is the IANA web page, or one of the IANA web pages, and down here in the bottom there is a link to the root zone, and it's

www.internic.net/domain/root.zone. So, if you would like to see what the root zone looks like, you can just go to this URL and you get a full copy of the root zone.

Here is an example of doing this using Wget, or you could also use Curl. Basically, you open a command line and you, Wget www.internic.net/domain/root.zone, and you download the root zone, and currently it's around 2.1 megabytes or so. If you don't want to use Wget or Curl, you can also fetch it using the DNS protocol itself, dig axfr.@b.root-servers.net will get you a full copy of the root zone. You can also fetch it from b.root-servers.net, c.root-servers.net, d.root-servers.net, f.root-servers.net, which is run by Jeff over there, g.root-servers.net, et cetera.

And some people are doing this. Here are some stats from b.root-servers.net and this graph is unique source addresses fetching the root zone per day. It's around 2,500. We can't definitely tell if these are resolvers doing the local root stuff or if they're just users fetching it manually. At 2,500 or so from B, if you multiply that by all of the root servers that are making it available, you can get an idea of how many are doing it via AXFR.

Obviously, though, having some guy go off and fetch the root zone manually every day and stick it somewhere is not very scalable and not very reliable. So, the major DNS resolvers have all built this in.

This is the example config from the BIND documentation. And so, this is saying a resolver should go off and fetch the root zone, and it should fetch it using the DNS protocol, and it should store it

locally. So, this is basically, you stick this in your BIND configuration and ISC BIND will go off and get a copy of the root zone and maintain it for you.

If you run Knot Resolver, you configure this. Basically, same thing. You should go off and fetch the root zone. Knot uses HTTP, which I think is a more scalable way of doing this and better way. So, it says go and fetch it from the IANA internet site and optionally, it can use a TLS certificate bundle to validate that it connected to the internet website. Basically, that protects the actual transfer.

If you use Unbound, same thing. Go fetch the root zone, fetch it using HTTPS, and store it locally. And now you have a copy of the root zone locally, and the name server will automatically maintain it for you.

One of the common questions that we get when discussing this is, but how much data is the root zone? Isn't it going to be a lot more traffic to go off and fetch this and store it locally instead of just getting their answers individually?

To figure that out, I set up a test where I configured Unbound, and I logged all of the traffic coming from the root servers. I then did a query for a name in every TLD, so my machine would go off or my resolver would go off and resolve every TLD and that ended up being about, where's the number? 1.1 megabytes of traffic.

I then compared that to doing AXFR, so basically fetching the whole zone at once from b.root-servers.net and that's around 1.6

megabytes, so a little bit bigger, but not much. Then I decided to see what it's like doing HTTPS. If you get it using Wget from the internet website, the total is around 2.3 megabytes. And that's larger than the root zone itself because of the TLS overhead. There's some certificate exchange. There's some web overhead. But seeing as it's being fetched via web protocols, we can take advantage of this thing called Gzip compression.

Fairly much all web servers and all web clients at this point will do Gzip compression on the fly. So, I configured a web server using NGINX, and in the NGINX config I said when somebody fetches a file, automatically compress it, and the client will decompress it. And then it's around 990k.

So, if you fetch it via HTTP or HTTPS, you can take advantage of these things called CDNs. These are content delivery networks and they're specifically designed to distribute web objects. Basically, every web page that you get on your laptop and for all intents and purposes all videos that you get like YouTube or any other video site is likely served from a CDN. These are designed for high scalability and geographic distribution and these are basically designed to share these types of things.

So, if all resolvers ended up doing this, how would this scale? Would this actually scale okay? Would it be all right? Here is some sort of quick back-of-the-envelope math.

If we assume that there are one million recursive resolvers, which is actually very much on the high end, but it's actually really

difficult to figure out how many there are in the real world because the definition of recursive reserve is a little vague.

But if we assume there are a million or so, and this own file is around 2.5 meg, which is what it is at the moment. And we were to assume that each resolver fetches the 2.5 meg file three times a day, which is, again, on the high end, the total data transferred would be around 7 terabytes a day, or around 200 terabytes per month. If this were distributed using Gzip compression, it's around three terabytes a day or around 80 terabytes per month.

In the grand scheme of things, this is basically nothing, right, this is lost in the noise for a CDN like Cloudflare or AWS or Google Cloud or any of those sorts of things. The transfer rate is also really small. It's around 35 or 36 transfers per second.

But if you were to go off and fetch the root zone from a CDN or wherever you fetch it from, why would you actually trust this root zone? If you were to fetch it from Cloudflare, how do you know that Cloudflare hasn't monkeyed with the information in the root zone?

Well, there's this thing called Message Digest for DNS zones. This was an RFC published probably three or four years ago now, I guess. I can't remember exactly when we did it. And it contains a record called ZONEMD. And what that is, is basically you take the entire zone file, and you create a checksum of it, and then you sign that using DNSSEC.

Here's an example of a ZONEMD record and that is a signature over the entire zone that it's contained in, in this case, the root. So, before any of the implementations would actually load the root zone, they do the ZONEMD validation, and then you can tell the entire zone file is exactly the same as it was when it was signed by IANA, technically by the RCM. Five minutes? Five years. Okay. I thought you were saying five minutes. So basically, you can tell that this is the real legitimate root zone.

So, the documents that we're discussing, it's actually a couple of documents at this time. We broke it up into four of them. And this updates RFC8806, which I mentioned earlier, and contains a few changes.

The first of these is having everybody fetch it from www.internic.net/blah, blah, blah does not seem like a great idea. Having any single point of failure doesn't seem great, or fetching from any, you know, having one CDN, one distribution point does not seem scalable.

So, the document provides a list of places that you could fetch the root zone from and which algorithm you should use. This is extensible, but currently you can do things using like AXFR, HTTP, et cetera, but we could extend that in the future.

There is also procedures for updating this list. So, if somebody else would like to help by distributing the root zone or making it available, this could be added to this list. The document also relaxes the implementation requirements in 8806. The way 8806

was originally written, it required that you have a recursive resolver, and right next to it, you have an authoritative server. And the recursive resolver had to ask the authoritative server each query.

That's not actually how anybody implemented it. The way implementations actually built it is they basically fetch the root zone and they stick it in the cache themselves. You don't have two instances running. So, the update basically says, implementers know how to do this sort of thing. As long as it looks like you're fetching it, we don't really care how you do it.

There are a couple of other small changes. Basically, if you're unable to fetch a new copy of the zone, if you can't grab it and put it in your cache, you should fall back to using regular DNS. There's some other things like other resilience changes, clarifications, et cetera. Here is basically what I said before, but directly copied from some of the documents we were writing.

Again, doing this sort of thing has some benefits. It increases performance. If you don't need to send a query to the root, you don't need to send a query for the root. Improves some privacy, DoS mitigation, and as a bonus one, it reduces the criticality of the root server system. If anything does ever happen to it, it's not an immediate concern, not an immediate issue.

So hopefully, we have lots of questions because we have 10 minutes left, and I figured we'd have some questions.

ONDREJ FILIP Thank you very much, Warren. I think we should have all the presentation with kittens from now. So now it's time for questions, please.

GIHAN DIAS Thank you, Warren. I think that's quite interesting.

ONDREJ FILIP Please, state your name before you --

GIHAN DIAS I'm sorry. I'm Gihan from .lk, Sri Lanka. So, I understand so this would be mostly useful for people with invalid root or top-level domain queries, right? If it's a very top-level domain you would have got cached anyway?

WARREN KUMARI Yes. I mean, largely, except that for each query that you have in a new cache, you know, that isn't in your cache already, you need to go off and fetch it. So, it helps with both the initial query when you need to fill the cache and also if there are unknown requests. There is another thing which also helps --

GIHAN DIAS Right. I mean, 99% of the time, any valid query would be served from the cache.

WARREN KUMARI Yup. Generally, like almost always, root TLDs, you know, the TLDs are already in the cache.

GIHAN DIAS Right.

WARREN KUMARI However, you still need to get them there, so it's primarily reliability.

GIHAN DIAS Right, but that's like once a day or once every --

WARREN KUMARI Once or every two or three thousand queries.

GIHAN DIAS Whatever. Right. Then you didn't really talk about a protocol for figuring out the sequence number and figuring out is this root -- I mean root doesn't get updated every day. So would be just good if we can just figure out how the sequence number change, which of course AXFR does, and whether you could use IFXR rather than AFXR. Could you comment on that?

WARREN KUMARI Yep, there is text in the document. I didn't want to get too deep down into the details, but basically, if you're doing it using, for example, HTTP, you can use the expired --

GIHAN DIAS So that's maybe my question. Why use anything other than DNS? Because it's a resolver anyway which knows how to do DNS. So, why not use just IFXR which would be very easily just get a few kilobytes or whatever of changes every so often? Wouldn't that work?

WARREN KUMARI Currently none of the root servers are doing IXFR as far as I know. There's a little bit more overhead if you want to do IXFR because you need to keep the current one and some changes. Generally, though it would work the same way that a secondary server generally works. You check to see if the SOA has changed and if it has, then you can download it again. If it hasn't changed, you don't.

If you do it using HTTP, you can also do the same thing using the HTTP expire header. So, you don't need to fetch the entire zone each time. You can do a expire query and see if it's changed. So, we have that.

GIHAN DIAS It's a different timeout than the DNS timeout?

WARREN KUMARI Pardon?

GIHAN DIAS That's a different timeout?

WARREN KUMARI Yeah.

GIHAN DIAS So that might -- Okay, there may be a difference between timeouts?

WARREN KUMARI We have a bunch of text in it, but it's basically, use the same timers that the DNS does, right? You have the retry timer and the expire timer. So, it basically says use those and fetch it using whichever protocol. You can check to see if it's changed with the cache headers. But yeah, look at the document, it has all of that. Anyone else? There must be more queries. Oh, Jeff.

JEFF OSBORN Silly me. I was using the Zoom hand up.

WARREN KUMARI Oh, sorry.

JEFF OSBORN And it's always one or the other. I have a thousand questions that I don't want to trouble people here from. Can you and I definitely get a half hour sometime?

WARREN KUMARI Sure.

JEFF OSBORN Okay. I'm going to go with the big question. I would love to see this happen. I spend, I've got a little nonprofit that gives away free software that does the thing you talked about before, that provides with DNS just like Ondrej does. And it's a terrific thing to do except finding a couple of million extra dollars to run a root server system is a real burden. And getting rid of it plus the political implications of why all of this stuff, I would love to get rid of all of this.

So, you kind of opened with, why are people so stuck on the past? I'm just wondering who that is because I would run screaming to do this. So the real question is, I'm concerned about this implementation at a level that we could get to scale really fast because you're going to see Andre next week, I suspect, my Andre, other Andre. And let's just make sure this is ready for prime time before we roll it out.

And then the one actual other question I had is, do you expect this to be fairly incremental? Does this start on a Thursday and the root server system stops the next Monday or?

WARREN KUMARI: No, no, no. The root's going to have to be around basically forever, right? Like we've seen potentially the amount of traffic hitting the root would go down over time.

JEFF OSBORN: But it's tremendously less critical.

WARREN KUMARI: Yeah, it becomes less critical. And so, if it has an issue or an oops, it's okay.

JEFF OSBORN: Great.

WARREN KUMARI: But yeah, so the root system is going to have to be around forever for all intents and purposes. So why isn't it happening more?

I mean, currently it requires people to go along and make some changes, right? You need to enable this in BIND. You need to enable it in Knot. If you're a resolver operator, you need to go and twiddle the knobs and configure it. If implementations started

doing this basically by default, there would be a large stair step as new updates roll out.

Potentially, the current IANA system would be more than capable of handling it, but I believe that a number of large CDNs are more than happy to just provide the transit for it because, as I said, in the grand scheme of things, it's lost in the noise. So, there are a number of CDNs who have already started saying, “Yeah, sure, I won't even be able to measure this load, so happy to provide it.”

Obviously, one then needs to have the list of who's doing it and the theory is IANA will sort of maintain that list we chatted with Kim already.

JEFF OSBORN Okay. It’s just --

WARREN KUMARI Did that actually answer your question or not? I'm not sure if I answered a different question.

JEFF OSBORN There's a million more.

WARREN KUMARI Of course.

JEFF OSBORN: But let's just make sure we coordinate. I don't want to back you and accidentally say the opposite of something and then have to --

WARREN KUMARI: Of course, thanks.

EBERHARD LISSE: Let me -- Eberhard Lisse .na, the former gynecologist and obstetrician. We see this a lot. Evidence-based medicine is the in-thing and from there, evidence-based decision-making has progressed into other areas. And we see that older people, like me, find it difficult to adapt to new ways. It has worked, it ain't broke, that kind of thing.

And then the other thing is why many companies, like I was just saying, why should they be bothered? Big content delivery networks won't notice the difference because the difference for them is so little. That said, I think from what I hear, it's the right thing to do in principle, but you must convince commercial interests, I think.

JEFF OSBORN: I'm going to come from a very opposite direction at 3:00 something this afternoon in the third session. So, it's going to be interesting to compare and contrast the two presentations.

ONDREJ FILIP Okay. I think we will thank Warren for the nice presentation with kittens. Thank you so much.

WARREN KUMARI Thank you.

ONDREJ FILIP Yeah, please, and the next presentation comes from Harsha Saparamadu and it's about best security management practices. Please, the floor is yours.

HARSHA SAPARAMADU So, thank you, Filip. Perfect. Again, it's not working. Right, so good morning, everyone. So, my name is Harsha Saparamadu and I work as an information security engineer at the LK Domain --

ONDREJ FILIP Can you speak a little bit louder, please? Or closer to the microphone?

HARSHA SAPARAMADU Yeah, yeah, sure. So, and I am an ICANN85 fellow, so it's a very pleasure to be here today and share some ideas. And share some of our experience with you regarding the security management practices that we experienced during the implementations of information security management systems, ISMS at the LK Domain Registry.

So, before going into that, I would like to briefly introduce my organization first. So, we are the .LK Domain Registry. In simple terms, we are the ccTLD country court top level domain for Sri Lanka, which is .lk.

At present, we have more than 60,000 domain names registered on the .lk domain. So, in addition to that, .lk, we also manage two IDNs, Internationalized Domain Names, .lanka and .ilangay, which represents the Sinhala language and the Tamil language. So you can see some logos here, which are the initiatives by LK Domain Registry as a not-for-profit organization to Sri Lanka. And if you would like to know more about this works, you can simply visit our website www.domains.lk.

So, moving to the topic, why is security essential for a ccTLD? So, of course, security is important for every organization no matter industry. Every organization needs to protect their information, systems, data and the services. So, when we're comparing ccTLD with others, ccTLD has a special responsibility, that is, ccTLD represents a country.

So, for an example, if someone finds a phishing website under the .in domain, people quickly see that it's an Indian domain. So even before checking the details, the country name comes first. So, because of that, as a ccTLD, we have a strong responsibility to manage our namespace in a secure and trustworthy way to protect our public trust.

In Sri Lanka, we have 60% market share on the domain name. So, with that space, there are many important domain names including government websites, banks and the payment gateway. So, imagine what would happen if the DNS servers become unavailable even for a short time, it can create a serious disruption and even can affect the country's economy in a bad way.

So, the responsibility of a registry, I believe, is not only to manage a domain name, it's also to maintain stability, security, and trust in the namespace. So that's why both ccTLDs and gTLDs must pay close attention to security and operational stability in order to provide a reliable service to the community and to the wider internet ecosystem.

So, I would like to raise three simple questions you can maybe ask from yourself or maybe you can ask from your organization. So, one, have you identified all the security risks that are relevant to your organization?

Number two, are there necessary controls in place to mitigate those risks? Number three, are your policies aligned with international standards or the best practices? So, these are some simple three questions, but they are very important. They help us to understand where we stand and what improvements we may need to take. So, moving to the next slide.

What are the risks that ccTLD may face? So, I have grouped them into five main categories, but there may be a lot of this than that, but I try to cover the most common and important risks here. So,

yes, of course, the cybersecurity and the cyberattacks become huge risks for the ccTLD. Maybe DDoS, or maybe a DNS amplification, or maybe a DNS hijacking inside the threats.

These are the common to each and every organization, but when it comes to the second one, the service outage, this can happen due to maybe a hardware failure or maybe a soft app bug or maybe a system misconfiguration. So even a small configuration error can make a DNS service unavailable.

So, the third one is data breaches. This may not directly affect to the DNS resolution service, but registry data is always sensitive. For example, registry database may contain passwords, use information and Person Identifiable Information, PII. So, if this information is leaked, it may damage the trust and it may create a legal and reputational damage to a registry.

So, the fourth one is about the phishing and fraud. This is also related to cyber-attacks, but often appears at the early stage of a large attack. So finally, the fifth one, the third-party risk because ccTLDs does not work alone, we work with many other parties like registries, registrars, hosting providers and resellers. So, any of the third party, any of the risks of the third parties may indirectly impact your ccTLD system.

So, when we think about the risk in a ccTLD, we need to look at the full picture. So sometimes something that looks like a smaller risk

can turn into a big problem later. So that's why we say each and every possible risk needs to be identified carefully.

So, to address this risk, we need to have a proper and structured approach. So about three years ago, before we started this implementation, we had a similar situation. So, at that time, we believed that we already had the right controls in place to handle most of the risk, but in reality, that was not true.

So, I remember one incident where the few of our DNS servers were hit by the DDoS attacks, but that was not a large one, but it was enough to get our attention. So, after that, we worked with the local ISPs and implemented the layer 3 DDoS mitigation solution for protect those servers.

And in other occasion, I remember, we noticed that our DNS server management system had not been reviewed for the access quite sometimes, but after that, because we started use access review then and we even plan to review access in every quarter, but over time that process slowly stopped happening.

So, this made us realize something important. Most of time when an issue appeared, we try to fix only that specific problem. We applied a specific solution just to close the issues, but we did not always look at how to prevent that same issue from happening again.

So, we understood that fixing problems one by one in an ad hoc way is not a real solution. So, we need a proactive approach, not

just a reactive one. So, we started thinking of a better way to handle security across the whole organization. And while doing that we realized that following a recognized international framework could really make things easier.

So there's a lot of standards like ISO, GDPR, [inaudible - 00:43:43], so there's a lot of, status are there. So, the good thing about using a standard is that, that's already designed. So, the process are clearly defined and we do not have to build everything from the beginning. So, in the simple term, the wheel is already there, so we just need to learn how to use that properly.

So, our next step was to identify the most suitable standard for our organization and start building our security practices around it. So, after looking at different options, we decided to use Information Security Management System, ISMS, as our main framework. So, this comes under international standard called ISO 27001.

So, there were several reasons to why we select this framework. It provides risk management capabilities. It has predefined policies, procedures. Also, it has good governance structure especially the controls that have to identify and mitigate security risk, the main reason to why we select that ISO 27001 as our base framework.

So, if you look at the ISO 27001 Information Security Management System, ISMS, so this shows a very high level what we did at our organization. So, ISO 27001 already gives us a structured framework. It includes more than 25 standard documents to guide

to an organization on the policies and the procedures needed to protect their system and information.

In addition to that, it also defined 93 security controls, which has helped organization protect from three key areas. So, confidentiality and integrity of the information and the availability of the information. So that's the basic idea behind the ISO 27001.

And what we did was quite simple. We use ISO 27001 as our base framework, then we look at the policies, procedures, and the controls defined in the ISO 27001 and selected the practices that were important for our environment.

So, after that, we used them to standardize and improve our process we already had. So, in simple terms, we did not start from the beginning. We followed an already well design framework and adapted it to our organization. That's the whole idea of this presentation.

So, if your organization wants to standardize your current practice, processes, or maybe security management approach, or even if you think you already have good practices and controls in place, so it is always good to evaluate them against the recognized well-defined framework like ISO.

So, the exact framework you choose will depend on your own requirements. So, in our case, we mainly focus on two things. That is the protection of the information and the availability of our entire system. Because as a registry, as I mentioned earlier, even a short

disruption to our system can affect many services. So, the security and the availability were our top priorities. So that is why we selected ISO 27001 as our base model.

So, in the same way, any organization can choose a suitable framework based on your requirement so then can compare your current practices with that framework and understand where you are. So this is our high-level summary and this is about what we implemented during our implementation time. So, based on the guidance from the ISO 27001, the first step for us was to identify the most important policies, procedures we needed in order to improve our security standards.

So, after that we grouped them into five main areas. Strategy and governance mainly look at how we manage at the security or at the organization level, and how decisions are made during the operation. One key area is information asset management. Before we can predict anything, first we need to identify what are the assets we have.

So, if you don't have clearly identified what are the assets, systems data become very difficult to protect them. Because most of the time, organizations focus only on their critical systems and of course the UAT system and testing system need to be considered.

Sometimes attackers don't straight for the critical system because those usually have strong security controls instead of -- so they may

try to get in through the less protection system because like test environment or smaller internal system.

So, another important part is the risk assessment, regular risk assessment and proper risk treatment always help us stay prepared. As I mentioned earlier, the continuous identifying and addressing risk is one of the best ways to keep our system protected.

So, second area is about the technical operation and the defense. So here we focus on protecting our digital infrastructure. So, we implemented several procedures in this area.

System access control to review who can access our systems and the virus and malware protection to keep our system safe from the malicious software. And the vulnerable assessment and management, then the log management to track system activities and detect unusual behaviors. Then the backup and recovery which is very important, to make sure we can restore the system if something goes wrong.

So, the third data is about the change management. So, this is about building and updating system safely. So, under this we have change and release management procedure, that ensure the update or modification of every system is controlled and documented properly. So, nothing breaks and nothing is done by mistake.

So, the fourth data is the physical and human resource. Here, the media handling and the disposal procedure is very important. So, they make sure that both physical and digital information is handled securely and no data leak happens accidentally or intentionally.

So finally, we have resilience and response. So, this area focuses on what to do when something goes wrong, especially during an incident or a disaster. At this procedures we built around core principle that is confidentiality, integrity, and availability. We call it as CIA in the ISO. So, this was our main focus to make sure the information is kept safe, accurate and available when needed.

So, throughout every part of this process, continuous audit is very important so it helped to stay on track and make sure we don't miss anything. So, at LK Domain Registry we do internal audits ourselves twice a year, carried out by our own engineers and in addition we also do an external audit once a year by a completely independent party.

So, this is important because sometimes internal audit may miss certain issues. So, we might not even see something as a risk because we are too close to the system. So external auditors can look at the things from the fresh perspective and identify risk we might overlook.

So apart from audits, there are some other important activities that help keep system effectively. So, these include the review controls, continuous monitoring, and learning from the past incidents,

lesson learned, and the management review. So, all this together makes sure that the process we put in place keeps working properly over time.

So, what are the challenges we have during the implementation? So, the main challenge was cultural resistance. Some people in the organization saw ISMS as extra work. So, to overcome this, we conducted multiple awareness to explain why ISMS, Information Security Management System, is important especially for critical organization like ours.

So, another challenge was the large number of documents. So, we implemented 18 procedure documents and there are even more policies procedures beyond that so sometime it was confusing to manage all them. So, to resolve this, to solve this, we prioritized the most critical document first and simplified them as much as possible to avoid conflicts between the policies.

And deploying ISMS without disrupting our existing workflow was also a huge challenge. So, sometimes some compliance requirements mean we had to change our existing workflow, but instead of creating a separate process, we worked to integrate new controls into workflow that we already had. So, that's one challenge we had.

And during the initial stage, conducting a risk assessment effectively was very difficult. So, we could not always identify all the possible risk within the organization. At that point, we take

help from external auditors and train our staff to improve our risk assessment process.

So finally, resource constraints were also a challenge because our team is small and then our daily operations already keep everyone busy. So, to manage this, we divided the project into phase wise and implemented step by step.

So, all these challenges throughout that implementing ISMS is not just about the document or control, so it's about the people, process, and the carefully planning.

So finally, after going through the whole process, there are many benefits. So, first, we gain confidence in our own system. We know our system policies and processes are strong enough to face a risk. So, second, our teams know what to do when something goes wrong. That is a good advantage from this implementation.

So, at the LK Domain Registry we conduct annually cybersecurity drill, disaster recovery drill, and the fire and safety drills so this ensures that everyone is aware of their responsibilities and can act quickly during the incident.

So, another benefit is clear visibility of risk. So, you know what risk exists and how they are being managed so the fewer surprises that means. And problems become easier to handle because they are already being identified and addressed. Yeah, yeah, so this is how we implemented the Information Security Management System at the LK Domain Registry.

So just remind, we are not officially certified in the ISO 27001, but we follow its recommended process and procedures in our own way. So, this help us to ensure the security and the safety of our information and the system. So, thank you very much, guys for your valuable time.

ONDREJ FILIP Thank you very much, Harsha. We have time for very quick question. Is there any? I don't see so again thank you very much. Thank you.

HARSHA SAPARAMADU Thank you, Filip. Thank you.

ONDREJ FILIP And then now it's time for the typical host presentation. So, who is going -- ? Rajiv Kumari will deliver the presentation, okay. Thank you very much. So, we have some small change in the agenda, but anyway we will not miss the presentation, so the floor is yours.

RAJIV KUMAR Thank you, sir. It's okay? Yeah, thank you, sir. My name is Rajiv Kumar. I am giving a host presentation on behalf of NIXI, National Internet Exchange of India. Good morning, good afternoon, good evening to all of you as per your time zone because a lot of people are joining online as well.

So just I give a pre-introduction why we introduced AI-powered detection of malicious .in domain name. NIXI is an operational body of .in domain name. We are only operators. There is a lot of LEA, Law Enforcement Agency, in India, who is monitoring each and every domain name. Some people know the who is detailed, but some don't know the who is detailed at all.

They come to us why the content is abused. At least the content abused is not a responsibility of any TLD as per the ICANN guideline, but our law enforcement agency, at least you monitor who is properly.

Then we start a proactive action to take action to whatever the complaint received from the law enforcement agency. We take the domain down, put the domain on server hold as per the direction received from the LEA.

But this is not sort out our problem at all because some domain is run by a proxy, someone run a domain name on behalf of someone else. And if I provide the detail to our LEA, LEA told that it's fake information, it is not right at all. Then we start evaluating each and every domain name daily basis, whatever the domain created, and this is also not sort out the problem at all. It's reduced the abusive domain name approximately 80%, but not 100%.

So, we appoint a local company, the Sesta IT, to monitor our domain name registration based on the AI technology. And the motive of that, make a domain name trusted and intelligence. We start monitoring the four layer. The first layer we evaluate the

domain name, whose details. Second layer, we take a visible registrant like their phone number and mobile number used any time in the abusive domain name in any TLD. Then we can recognize that domain name registration also.

Third, we can evaluate a domain content, website content. Fourth, we evaluate the AI IP number monitoring also from where the domain is registered. And on behalf of individual capacity, I request ICANN to kindly allow whose data as the geo-tagging so that no need to any whose verification is required. And third, we can evaluate the content also.

Based on this four layers of evaluation, our abuse is reduced approximately till 98%. Only 2 or 3% is not able to detect now, and 90% abusive domain is verified. And whatever the abusive domain detected, we can give our LEA to further verify and instruct us to put the domain name on hold.

CLAUDIA RUIZ Rajiv, can you speak closer to the microphone, please? Thank you.

RAJIV KUMAR Yeah, thank you. Thank you. So, our motive is to make .in secure as well as to make the global TLD secure. We book 3,000 domain name daily basis. We monitor each and every domain name and if the fraudulent domain name, like bank, e-commerce, government portal website is found, we can put the domain name immediately on server hold and ask the registrant to submit a KYC document.

If any sensitive domain registration is found, any time we can put the domain on hold immediately. We can launch, steal, disappear, repeat the scammer daily basis.

Why the traditional method is failed? We already told that it is a human, cannot scale the thousands of domain daily basis. Reporting based take down every day is a long lag of time. During the window, hundreds of users are compromised and fraudulent sites are visually interested and legitimate one is compromised domain.

We can start AI-based power solution. There is four layers. First is domain name analysis. Second is network graphic analysis. Third one is content inspection analysis, and behavior monitoring as well.

Domain detected at 9:00 AM, morning, after one minute, we flagged the domain impersonation and detected. And in the second minute, we can notify the authorized owner and the authorized person to take a further analysis and give a direction for domain blocked. And within 30 minutes, we can block the domain name currently and the user has zero compromise basis.

We can evaluate 2 million plus domain name within hours, less than hours, and found that 24,000 domain name are abusive domain name registration. It's a historical domain name, like first domain name. NIXI managing approximately 4 million domain name currently. And out of 24,000 domain name, 143 domain name fake website of SBI, State Bank of India, with the prominent

website, a commercial organization in India, the Bank of India. Suppose that 140 as a big huge number if the one single bank is duplicate or [inaudible - 01:04:56] of domain found that.

The roadmap, pre-registration detections, expense of beyond the .in to global TLD borders coverage, cross-border collaborations, continuous model of information and advice tactics solves. And one of my suggestions to give a geo-tagging of the whole detail, so that we can at least detect the domain registered location.

Because in some cases it is found that the domain is registered in the address of India, but the actual person is sitting outside of India, for saving a GST tax in the Indian government and for doing an abusive activity. So, geotagging is one of the most important parts.

Strategic, shift cybersecurity postures from reactive to proactive basis. Scalable, non-linear increase to cost with the domain name volume. Protect the brand trust for businesses and public confidence in digital infrastructures. Align with Indian borders, digital safety, and financial inclusion goals.

Every fake website catching, someone's savings protected, someone's identity kept safe, someone's trust in the internet preserved. The internet should be safe for everyone. AI makes that possible one domain at a time. Thank you.

ONDREJ FILIP Thank you very much. Are there any questions? Yeah, please.

VADIM MIKHAYLOV

Thank you, Rajiv, for your amazing presentation. Vadim Mikhaylov from .eu. We have the same research, but in the research phase. Am I understand correctly that you are in production phase so it's working system that analyze it? Can you tell, have you any assessments of false positive rate of the domain detection? And it's the first question.

And the second one, do you have any plans to expand that model for a huge number of your IDN domains? Thank you.

RAJIV KUMAR

Thank you, Vadim. Yeah, last month, my team detected 286 domain names using the historical registration based on our AI-based tools. We share the data with our LEA for further analysis and give a direction because NIXI don't have the power to block the domain name, historical domain name. And we found that 201 domain from that data is, also LEA confirmed that, yeah, it is abusive domain name, please, put and do it that. One of the benefit I got it that we can give data because LEA don't have the historical one registration base. So, this is the benefit I got.

Yeah, your second question for IDN base, of course, we are working closely with our team for the IDN as well and very shortly we can launch our IDN-based analysis to laws.

ABDALMONEM GALILA: Yes, this is Abdalmonem Galila for the record. Thank you, Mr. Rajiv, for this important presentation. It is just a comment rather than a question. Actually, thank you very much for highlighting the role of one of the emerging technologies like AI in order to foster the registration of domain names and DNS abuse mitigation. And otherwise, from my position, it could be used to accelerate the adoption of universal acceptance.

So, by the way, I will invite you from here for the next session for the take day that will highlight the role of AI in order to accelerate the universal access to adoption. Thank you very much.

RAJIV KUMAR: Thank you, I will be there.

ONDREJ FILIP: Thank you. Any other questions? Please.

KEVIN BROWN: Sorry, just a quick one, I didn't get your name.

RAJIV KUMAR: Rajiv Kumar.

KEVIN BROWN: Sorry?

RAJIV KUMAR Rajiv Kumar.

KEVIN BROWN Rajiv Kumar, thank you.

RISHAB DHANIYA Okay. Hello, everyone. My name is Rishab. I am a NIXI fellow. So, my question is that when you have all these AI checks in place for malicious domain names detection, it is possible that sometimes AI flags some of the correct ones and maybe some of the malicious ones get passed. So, is there a manual check in place which you implement? Thank you.

RAJIV KUMAR We continue with manual checking till the AI 100% did not match and all the things, but yeah, after some time we stop manual checking and shift to AI base, 100%.

ONDREJ FILIP Okay, we have one more question.

ROHAN SACHDEVA Thank you, Rajiv. My name is Rohan. I am from .au and I'm also an NextGen fellow. My question is the AI system that are being used, what algorithms you use and how did you train them and how did you develop it? Thank you.

RAJIV KUMAR It's a typical question for me. I am the only implementation agency and implementation person. The right person is sitting behind me. I introduce you after that. They can teach you how they develop the software.

ROHAN SACHDEVA Thank you.

RAJIV KUMAR Yeah.

ONDREJ FILIP Any other questions? We have two more minutes.

UNKNOWN SPEAKER Yeah. This is [inaudible - 01:11:31] for the record. Sir, I wanted to ask that as AI helps detect and block malicious domains more quickly, how can registries, registrars, and the ICANN community collaborate to scale such solutions globally and make the internet safer for everyone out there in other parts of the world?

RAJIV KUMAR Thank you, [inaudible - 01:11:57]. We are the start. We initiate that this process, maybe other ccTLD also initiate the similar process in their country, but not in this platform. But this is the platform where we can start the initial page and we can approach other

ccTLD to kindly ask and do it that. And also, the ICANN person will go ahead with the same process.

ONDREJ FILIP: Thank you very much. I think that's the end of the session. Thank you. Thank the presenters for interesting presentation, thank you very much. As already said by Abdalmonem, you know we will reconvene at 13:15 the local time, but before that Eberhard has some technical information for us or more administrative information for this session.

EBERHARD LISSE: I forgot to mention when I went through the agenda that Kevin Brown, a member of our technical working group sitting over there will go and do the wrap-up, which is why he asked Rajiv for his name so he knows who presented what.

ONDREJ FILIP: Excellent, thank you very much. Enjoy lunch and see you at 13:15. Thank you.

RAJIV KUMAR: One additional announcement, NIXI is providing lunch for 9, 10th, and 11th. Kindly collect a coupon from NIXI booth or you can visit directly to the lunch area. You can get a coupon from there. Thank you.

EBERHARD LISSE	Sorry. Where [inaudible - 01:13:29]?
ONDREJ FILIP	At the NIXI booth, I think.
EBERHARD LISSE	So, the NIXI booth?
ONDREJ FILIP	Yeah, very kind of you. Thank you so much. Thank you for your hospitality.

[END OF TRANSCRIPTION]