ICANN85 Mumbai | CF – How it Works: IROS and IANA
Saturday, March 7, 2026 – 10:30 to 12:00 IST

FERNANDA IUNES

Hello and welcome to How It Works, IROS and IANA. My name is Fernanda Iunes and I am a Participation Manager for this session. Please note that this session is being recorded and is governed by the ICANN Community Participant Code of Conduct, ICANN Expected Standards of Behavior and the ICANN Community Anti-Harassment Policy. Please observe the following guidelines to participate in this session. I have also posted them in the chat for your reference.

During this session, questions will only be read aloud if submitted within the Q&A pod. Interpretation for this session will include English, Hindi and Spanish. If you'd like to speak during this session, please raise your hand in Zoom. When called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use a physical microphone to speak.

Only questions posted in the Q&A pod will be read aloud during the session as time permits and when directed by the chair of the session. Please state your name for the record, the language you will speak, if speaking a language other than English and speak clearly at a moderate pace. And with that, I'll hand it back over to Siranush.

SIRANUSH VARDANYAN

Thank you, Fernanda. And again, welcome to the session to learn about what is IROS and what is IANA. So, for everyone to understand, IROS is Identifier Research, Operations, and Security function of ICANN. And it's my pleasure to introduce ICANN Senior Vice President and Chief Technology Officer, John Crane, who will be telling us about IROS and staff member, the colleague, my colleague, Adiel is here as well.

And Kim Davis, who will be telling us what is IANA, which is the Internet Assigned Numbers Authority. And Kim is vice president of IANA services and president of public technical identifier here. So, without further ado, for us to start learning, we have the clicker for the presentation. So, you can have it. I would like to give the floor to John Crane to start the session.

JOHN CRANE

Well, thank you. Welcome, everybody, to the ICANN meeting. I assume for some of you, it's your first ICANN meeting, and some of you have been to multiple. I am John Crane. I am part of the executive team and the officer of the organization known as the Internet Corporation for Assigned Names and Numbers. If you don't know what that is, you are probably in the wrong room because that's why you're here. I've been at the organization since inception, which was a lot of years ago and have had various roles.

Today, I have the privilege of being ICANN's Chief Technology Officer. If you talked to me six months ago, that was really about what we're going to talk today about, which is the identifier research operations

and security group. As of recently, I also took over all other aspects of technology inside the organization. So, my role is much bigger than this now, but we're going to focus on this until we find time to change our slides and talk about broader things in future sessions.

So, I don't know exactly what the next slide is, so I'm going to have a quick look at it. So, at its core, the IROS function is separated into two elements. One is what we call the office of the chief technology officer. And you can think about that as a research and engagement function. We do a lot of research around technology and identifiers and things that affect them. And then we engage with the community about that.

That can take the form of a sitting around a table like this today, or it can be in a more formal training environment. And we'll talk about those later. And on the other side is what ICANN was actually formed to do, which is the Internet Assigned Numbers Authority, which is about registering identifiers for the internet. And I'm going to let Kim talk about that later. I'm not going to get into these. I'm just going to do a quick intro.

So, welcome all of you to Mumbai. How many people here is this your first ICANN meeting? Quite a lot. I'm also aware there's people sitting behind me, so I will be turning on occasion. So, I'll start by giving you the advice I give everybody. Chill out a little bit, enjoy. You are not going to be able to know everything that's happening at an ICANN meeting.

There are hundreds of meetings, many, many topics. Try and figure out what is the area that you are interested in. Most of us come here

with a specific interest that we have in our lives. I'm a technologist. How many people here are engineers or technologists? Okay, you're the really good people. You're my favorite people, because I'm a technologist.

How many people here are really more interested in the policy discussion? You people are also cool, but you're just outnumbered in this room. ICANN is where policy and technology meets. That is what's fascinating about this place. So, if you are a technologist, go look at some of the policy stuff. It's actually really interesting. It affects what you do in your day job. You may not realize it now, but policy is very important.

If you're a policy person in this realm, go follow some of the technology tracks. Go and figure out about the technology, because if you're going to be setting policy around technology, you're going to even need to understand the technology, or you're going to need to be able to collaborate with technologists and the folks that do. So, we work together on this.

As I said, welcome to ICANN meeting. Enjoy it. Don't try and learn everything. Staff, we all have a blue thing underneath here that says ICANN Org. Stop people, ask them questions. If you can't find the room, like I had about 30 minutes ago, just ask somebody. We are mostly friendly. If you can't find a staff member, just stop anybody. Everybody here is very open, very welcoming. We rely on getting new voices and new interest into this ecosystem. So, we're really happy to have you here, and what you do is very important. So, thank you.

So we're going to go into talking a little bit about the functions, and I'm going to see which order we put this in. We're going to talk about the office of the CTO first. So, we split this, as I said, mainly into research and engagement. I'm not going to go into these in depth here because we're going to have some slides in a minute.

One area that we also have is we have an operations function. We're a reasonable size group, about 30 people, but we have a lot of things happening. So, there's a lot of operations need to happen. Think about signing contracts, planning projects, all the things that any business unit really needs to do, we have just like any other business. So, we have our own little group that does that. They do that across all of the areas where I work. So, they also do it for engineering and information technology, but it's a pretty important job.

As you go through your careers, you'll find that if you don't do the planning and the operations properly, none of the rest follows. So, we have an operations group. We have engagement folks that Adiel leads here, and he'll talk to you in a minute. And then we have two areas of research. The way I look at this in my mind is that one of the groups of researchers, there's about four researchers in each group, one of the groups are engineers, very deep into protocol work, going to the IETF, the Internet Engineering Task Force.

There's people here know what the IETF is, but those that don't, those are the folks that make the protocols that make the internet work. The internet works because we have open standards that everybody shares. The IETF is the place that happens, not ICANN, but our staff

participate, not just my staff, but many staff across the ICANN business organization.

We operate in that area, and we're interested in the development of the protocols. And then the IANA function, as Kim will talk about, registers those protocol identifiers. So, those are all engineers. And then the other side are much more academically trained researchers come out of universities, are very interested in the data statistics. And both groups operate in a slightly different way.

The academics are really good at academic process, and when they publish things, they may go through academic peer review before they get published by ICANN. So, it can take a year or more sometimes for us to publish something because we go through the academic review process. And the engineers, they tend to publish more open documents, less, I wouldn't say less rigor, but less of the academic process.

And the two of them together, just like policy people and technologists, they actually make for better discussions and better products. So, that's why we have the two different areas of research, that they have different methodologies and they really combine well to produce better product when you put those two skills sets together.

And then I'm going to, is the clicker working? I was going to say, where's the computer? Where are we at? And we have a broken clicker. Let's find somebody who knows how to do technology. I was

just too far away, I guess. Okay, so I'm going to pass over to Adiel who's going to talk about some of our technical engagement. Adiel?

ADIEL AKPLOGAN So, yeah, the technical engagement team is mainly coming support to the rest of the OCTO team in terms of building and developing collaboration relationship within the technical ecosystem. As John mentioned, we have a research team that does a lot of work looking at behavior around DNS operation and ICANN mission in general. My team will take some of the outcome of those researches and engage within the community in general.

In addition to that, we work a lot with operators on the ground to promote best practices related to DNS operation in general. And also collect feedback that are used by the research team to know where they should focus. So, overall, it is the team that is the public facing part of what OCTO does.

We also work to build tool that can help us promote those best practices. Tool like KINDNS, you have probably heard about it. KINDNS is a framework similar to manners. That is used by operators to assess how they are doing in operating secure, scalable and resilient DNS in general.

We promote the best practice using that framework these days and partner with many organizations globally to make sure that operators being authoritative operators or resolver operators follow some of those basic best practices.

The technical engagement team has a specific particularity which is it is a regionally distributed team. Although within OCTO, we have most of the team work remotely and are located all across the world. But the technical engagement team specifically have a regional focus in the way they operate. We structure that way to make sure that we take into consideration the diversity in different region because different region had different level of maturity in operation, different level of needs.

So, the team that are based and positioned in different region work with operators in the community in those regions to identify the needs and to also provide the support that they need specifically. So, it is important for you to know who in your region is responsible for working with the rest of the ICANN team in the region, but focus on the technical engagement.

So, we have a representative in Africa and Middle East, it is Yazid, you may have hear about him already. In Europe, we have Ulrich Weisse. He's the technical engagement manager for Europe. In this particular region, Asia Pacific region is Champika. He must be, no, I don't know if he's here yet, but he will be around, he's coming. Latin America.

SIRANUSH VARDANYAN We will see him tomorrow.

ADIEL AKPLOGAN Exactly, he's presenting tomorrow, exactly. In Latin America, we have Nicholas Antonio. He's based in Uruguay and support the activity

down there. So, as you can see, we really want to make sure that we build the proper connection between different region, between operators and between actor on the ground with what the rest of the OCTO team does.

So, I think that's cover the technical engagement aspect of this. We work as well a lot with other sister organization within the I-star, what we call the I-star ecosystem, ISOC, the NOG, the RIR in general. Most of our outreach and capacity building activities are developed and done with them. Happy to answer any question and interact around this further if needed. Next slide.

ALFREDO CALDERON I do have a question. This is Alfredo Calderon. I'm a mentor for the NextGen program. My question is, you mentioned MANRS, you mentioned this other initiative. What is the difference between both and what does MANRS mean?

ADIEL AKPLOGAN Good, yeah. MANRS is the Mutually Agreed Norms for Routing Security in general. So, MANRS is a framework that operators use to ensure that the way packets are flown on the internet in terms of routing are securely done, taking into consideration several best practices. KINDNS, however, is the same thing, but for the DNS, right? So, the routing is one part of the internet. It's how packets move. And the DNS is at the top gluing the two identifiers, the name and the IP address together. ICANN core mission is around the DNS.

So, we have designed KINDNS, taking into consideration the experience of MANRS to give operators a simple framework that can allow them to know how to deploy, operate the DNS based on the best practices, streamlining the best practices to the most important and critical one so that they can implement them.

You can go to KINDNS.org to hear more or see more about KINDNS. It is an initiative that ICANN is supporting, but basically it is a community supported initiative. Our goal is to continue to support it, but to bring more operators, more people from the community along that line. So, we want KINDNS to also become a reference when it comes to best practices operation in general as MANRS is today.

JOHN CRANE

Okay. So, I'm going to talk a little bit about the first research team. This is led by a gentleman called Matt Larson and is once again, geographically dispersed, not by design though. So, all of OCTO is geographically dispersed. We have people around the world, but as Adiel was saying, the engagement team, that's actually a very deliberate choice of where we put people.

For these kinds of roles, we've hired them where they are. So, we have a lot of very experienced, I would say well-known in the ecosystem engineers who are some of the best in their field. So, we hired them where they are. So, we have people in the US, we have people in Europe and we have people in other locations as well.

Actually, not many in the US, they tend to, a lot of the skillset around here tends actually to be in the European realm for some reason. I think the universities there and the way the network built out of the universities there just produced a lot of these people. There's not many of them, there's only three or four in each group.

The idea really, and one of the big goals of the Office of the CTO in general is to provide trusted information to the community so that when we're building policy, that it's built on a framework of factual data and a factual understanding of the underlying technologies. And this is what I was saying earlier, the policy and the technology must go hand in hand.

So these engineers, these researchers, they are actively involved in the DNS communities, but ICANN's mission is not just about the domain name system, it is about the unique identifiers of the internet. And I think Kim will touch on roughly how many of those there is, it's more than 10, there's quite a few of them but everybody focuses on the DNS because that's the one that people interact with, right?

People know names because they're the human recognizable part of it. So, obviously a lot of the policy also focuses on that and ICANN specifically has a role in setting domain name space policy for the DNS system but not for many of the other identifiers. That's why ICANN is so DNS centric. So, we spend a lot of time working on those technologies, understanding those technologies.

If you go to the internet engineering task force, you will see that often members of this group are actually the authors of some of the new

drafts and some of the protocols going back many years. So, really deep, deep in DNS technology and protocols. We, and I think this can be on the next slide, we study a lot of DNS traffic and we have a lot of projects just like the engagement folks do. We have measurement projects and boards where you can go and see information around the DNS and the internet's health.

ITHI is indicators of what is happening on the network. They're all sorts of different indicators, not just DNS. There're indicators about what's happening in routing, et cetera, and all downloadable, obviously not for us, but for the community.

We've measured the concentration of DNS resolution. So, a few years ago, there was a lot of discussion about the large resolvers and how they were dominating the resolution ecosystem. I'm thinking about the big public resolvers that run on things like 1.1. 1.1, 2.2.2.2. They have this pattern where they just get the single and the double digits, like 8.8.8.8 is a very well-known one. There's many of those.

So, we actually did a study on this and published the document that demonstrated that, in fact, although everybody was saying that, when you actually looked at how DNS traffic flowed and where the resolution was going, wasn't true. But there was a lot of discussion in the realm about what are we going to do about this problem of everything going to these few servers. So our job is to actually bring facts into the discussion. So, we did the research and wrote a paper and things went, oh, we all assumed that this was happening in the

technology, but when you measure it, it turns out not to be. So, that's the kind of thing we do.

DNSSEC error reporting or DNS error reporting in general, that wasn't something that the protocol did. So, some of our staff wrote a draft into the IETF and now there's an RFC. The RFCs are the documents or protocols that's introduced error reporting into the DNS so we can get better error reporting. And we track a lot of things at the DNS route.

ICANN is fortunate and privileged enough to run one of the 13 route servers for the internet. So, we have access to that for our researchers so we can look at a lot of things of what's happening at the top of the DNS. So, that was that protocol side of things.

The other side of it is of course, these then are more the academic side is the security, stability and resiliency research. They focus a lot, next slide please. They focus a lot on issues like how is the DNS being abused and what kind of security concerns are we seeing?

The other side also looks at security concerns but more from a protocol side. So, a lot of this is about understanding what state of the art is in researching abuse. We publish a lot of documents around this. We have our own research and we have our own tools. So, we produce tools for the community.

Often when we are producing a tool for the community, the tool of itself is of course important but what is really important from our side is actually to demonstrate the science and to show how things can be done. So, many years ago, there was this concept that you couldn't

actually measure how the DNS was being abused. So, we built a tool to demonstrate how you could actually measure this. But the important part was not the tool, it was that we published all of the methodology and all of the research behind it so that other people can go and build similar tools.

And if you look at the way we measure things now in the ecosystem, there are companies out there doing it, commercial and non-commercial. And a lot of the methodology goes back to that research. And a lot of the early work that we did and they've just improved on it. And that's our role is to kick off that science and get people to actually understand that we can measure these things. And then generally educating the community so that they have more informed dialogue when they're trying to set policy in this area.

Now, this week, there are new policy initiatives kicking off around DNS Abuse. So, you may see some of my staff getting to the microphone and talking about some of the facts behind it. So, if the conversation veers off into a little bit of mythology or concepts that people have in their mind that don't necessarily match what's happening like from a science or from a fact standpoint, then we might get up to the microphone and say, hey, we've done this piece of research or there are people over here that have done the research.

You might want to go read this and we'll be what we call SMEs or subject matter experts. And we also have funded research outside of our group. So, obviously we are only three or four of these researchers and we can't research everything that happens on the internet. So, we

funded research projects with universities, et cetera, so that we can leverage their skillsets and also have them publish on things that are a little bit harder for us to remain neutral on.

So, we can go out to a third-party university and say, here's a question, would you do this? Now, sometimes they come to us, sometimes we go to them, but so we do also sponsor some of the research projects and also, we occasionally sponsor like conferences and things like that where we can bring researchers together. Next slide, please.

So, these are some of the tools we've done. Domain Metrica is a measurement platform for currently focused on the reputation of domain names. So, not as a name being abused, but does a name have a reputation of being abused? Very different things, the terminology matters, right?

There is a lot of work out there in the industry reporting that names are abused or measuring that names are abused, but they don't always provide the evidentiary data to demonstrate that it's been abused. So, to us that has something that has a bad reputation. It's not necessarily abusive. It just has a reputation of being abusive, and that's what I mean about having a little bit of rigor in your, both in your studies, but also how you talk about these things.

I'm trying to be very clear about what we're doing. INFERMAL was one of those projects that we had done by somebody else. So, it's not an ICANN piece of study, although we were involved because we helped them with some of the methodology, but it was done by the University of Grenoble and they published that. It talks about a lot of things

about how people choose names to abuse them and why they do it. That wasn't our area of expertise. So, we went to these folks and had them do it.

We're looking at things like we have a current paper that is about parked pages. Does everybody know what a parked page is? I think a few nodding, but maybe not. So, a lot of times people will register a name and they don't actually put active content on it. You know, it's not somebody's website. It just has like an advert there. Those names can be used in different ways. They can be completely benign, right?

It can be just a name that is registered with a registrar and the registrar has an advertising mechanism where they put a standard page there, completely benign. But it can also be maligned. There are attack vectors where people put code on those pages, et cetera. So, we've got some studies on what that looks like and how you can identify the different kinds of things.

Did some work on looking at blockchain-based names and where they appeared in the route. So, there's all kinds of different things that we're doing here. Basically, trying to figure out what's happening in the ecosystem. How's the ecosystem being abused and misused and whether or not we can measure that in a scientific and reproducible way.

Next slide, please. So, where can you find some of our work? So, the way we publish things is through what we call the Office of the CTO document series or OCTO documents. Pretty standard way many

organizations do. You just have an initialization like OCTO or acronym and then a number.

These are some of the recent examples we've done. We've been doing these for years. So, we don't publish like one a week. Research takes time, at least rigorous research takes time. So, we might publish some years we may only publish one and some years we may publish four. Depends on what the active topics are. So, we're at 41. We've got three more in lineup for hopefully this year.

Some of these are very technical and some of these are a little abstract. So, I talked about, I'm trying to say, no, that's not even in this one, the one that was on part pages. So, evaluation methodology or an evaluation methodology for block lists, reputation block lists. Folks heard of those? These are lists of names that have this reputation for being bad. We use them for our research.

Now there are dozens of providers of lists like this and they get used in technology day to day. People put these on their networks and they use them to protect their clients, et cetera. So, if you're on your network and you get a warning that this name is dangerous for some reason, it might be on one of these block lists. But because we use them for our research, we have to evaluate how much faith, if you like, or trust we can put into each one of these lists.

And everybody that uses this data should really be trying to do that. But we realized there wasn't a methodology available for people to do that and do that consistently. So, we wrote one and published it. Now, all of these go out for comment, et cetera, so we got inputs, but

we actually wrote a methodology for evaluating these lists because we realized it was missing and people are using them, but without really having a good methodology to evaluate them and we thought, well, that's probably not a good idea. So, that's why we wrote that one.

Classification of unsolicited email threads. So, how many of you get email that is maybe not good, malicious in some aspect, right? We all get it. Well, it turns out it's not straightforward, right? People will say something is a phish and you can look at it and say, well, is it really a phish or is it a malware drop or is it something else?

And we realized that folks were classifying things without really doing the analysis. So, we built a tool based on the LLM, so using machine learning to look at, I think it was millions of names that had been flagged or emails that had been flagged as malicious is a feed that you can get from one of these providers and it has a lot of emails in there and they're all flagged as phish, but they're not. But we didn't really have a good methodology to actually decide which ones were and were not.

So, once again, we thought that's probably not good, we should probably go and develop that. Now, because we're a nonprofit organization with a mission that's driven around these things, we were able to go and spend the money to do that research and then just publish everything, like our algorithms, our software, et cetera, so that other people can actually take this and we can improve the ecosystem. And that's the kind of thing we do.

We've written some documents on blockchain technologies, trying to understand them, trying to help the community understand them. A while ago, that was a very big issue. It's still an issue, but it doesn't seem to be so prominent in the discussion anymore. So, we try to demystify some of that area of what do these technologies actually look like.

We treat them as though they're one technology. Are they really? Turns out, no, they're not. There's lots of different ways of doing this, but nobody had really documented that from a neutral perspective. Lots of documents from the people promoting this technology, some documents for people that were very much anti the document, but nothing really that we could find that really had a neutral aspect to it. So, we wrote a couple of documents on that.

OCTO 41 is about how do you randomly select candidates for an election process? Very different, but it's a question we got from another part of the community about when we're selecting chairs and when we're selecting candidates, and not just the ICANN community, but also in the IETF, et cetera. How do you do that using technology to have, I guess, not predictable, because that's not the way, but consistent mechanisms for actually selecting candidates? It's not something you'd actually expect from ICANN, but we get a community asked to solve a problem.

We had the expertise, so we went and did this. We've got documents on quantum computing. We've got documents on all kinds of things that you might not directly associate with what is happening in the

ICANN world, but they do actually have some tangential effect on what we're doing. Post-quantum cryptography is important, for example, for DNSSEC, so we have a document on quantum computing.

So, this is how we mostly publish our viewpoints and our data on technology. Next slide, please. We have a lot of other ways of doing things. Obviously, we have blogs and we have speaking engagements. That's for OCTO in general. This is mostly done by Adiel's group, but not exclusively. So, obviously, if we're going to go to an academic conference and talk about an academic paper we've done, we would send the academic. We won't send one of our regional guys. We may send them both so they can learn from each other.

Next, please. I think that was it. So, what I want to do now, I think, is ask for some good opportunity for some questions now, and then I'm going to hand over to Kim. So, I see, oh, one, two. Let's start here and go back, and then I'll look behind me. The gentleman in the white shirt.

UNKNOWN SPEAKER Hi, I was looking through the OCTO publications and my question is about, like, if someone is not working for ICANN, but like us, like the NextGen, the ICANN fellows, is there any possibility of, like, students contributing to the publications? Thank you.

JOHN CRANE Helps if I turn it on and not off. You can always contribute by giving comments, right? So, we publish these, we ask for comments on

them. So, you can always contribute by giving your input. We don't really have a mechanism for non-staff members to write into this series, but there are lots of other opportunities in the community to publish things, the various blogs that are not staff-related.

Now, if you have a specific topic that you're interested in and you would like us to write and you want to talk to us about it and give us input, we can always have that conversation. But there are lots of opportunities in the ICANN realm to, like, put out opinions, et cetera.

If you're really looking to write a scientific paper, this is not necessarily your best mechanism, especially if you're an academic, you might want to focus on some of the conferences, et cetera, out there that are better for that, but like the one conference or something like that. But we can talk to you about that. We're always open to, like, advise and to mentor. So, if you want to talk to one of our researchers about what you're working on, just reach out.

SIRANUSH VARDANYAN Can take one more question for now and then we'll move on. Please, go ahead. Lady first.

JOYSA KAUSHIK Hi, my name is Joysa Kaushik, for the record, and thank you for the session. So, my question is that you mentioned about a project, a tool called Domain Metrica. So, at present, how many scripts does it include? Is it just, like, Latin-centric or is it inclusive of other scripts as well?

JOHN CRANE

It includes all the TLDs namespace. So, it's all the generic top-level domains, including the ones that technically everything is in ASCII because it's DNS data, but the ones that are representations of other scripts, the IDNs are also in there, and some of the ccTLDs. So, yes, it's universal acceptance and the ability to use other scripts is a big thing inside ICANN.

So we try and make sure that all of our tools incorporate. I think we're at 96%, might be 95%. There are some of the tools that we use are big corporations that we just don't have the influence on. But, yes, it does include other scripts, yeah.

SIRANUSH VARDANYAN

Thank you, and I would like to give the floor now to Kim to talk about IANA. Kim, please, and then we'll have some time at the end for questions as well. Thank you.

KIM DAVIES

Thanks, Siranush. Good morning, everyone. I'm very happy to talk to you about IANA, which is my bread and butter. Excuse me. So, you heard what IANA stands for, Internet Assigned Numbers Authority, earlier in this session. But really, what is IANA all about?

Essentially, IANA is a part of ICANN that serves as the official record keeper for unique names and numbers that are used for the internet. So, what does that mean? It means that for the internet to work, we

need everyone to communicate on a technical level in the same language.

We do this through internet standards, and those internet standards rely on a number of different ways of communicating that need to be coordinated through using a common vocabulary of internet identifiers. What I just said probably doesn't make a lot of sense, so let me expand a little bit upon it.

We're all familiar with domain names, I think, in this room. Imagine a world where you type in a domain name, and depending on what part of the network you're on, you go to a completely different company. That wouldn't work so well. So, it is really critical that we coordinate the domain name system in a way that every domain name has the same meaning, no matter where you are in the world.

And that is fundamental to why ICANN exists, and also why IANA exists. Now, it's not just domain names that have this property. It is important on the underlying communication protocols that the internet uses that those protocols work the same way no matter where you are in the world as well. Anytime you use the internet as a user, you're often using, for example, a domain name, maybe you're using an app, but you don't actually get exposed to a lot of the ways that the communication you're making transacts over the internet.

Every time you request a piece of information or send something, there is a whole range of things that are happening behind the scenes that enable that communication to be successful. For example, your ISP, your Internet Service Provider, is communicating with other

Internet Service Providers and making quick decisions about where to send the internet traffic, like over which wires is that connection made.

When you're requesting a web page, there's often a negotiation happening. For example, we just had a question about supporting other scripts and languages. Whether you're presented with, for example, a web page in English or whether it's presented in a different language, there is a negotiation happening. Your web browser is telling the remote web server, here are my language preferences. And that web server is making decisions based on your language preferences and sending back the content in the language that best suits your preferences.

All that is happening in a way that's not really user visible, but does depend on a common language on a technical level being used to express the communication. So, whether it's deciding which route to take over the internet, whether it's your language preferences, or there's literally thousands of other technical decisions that are made in the course of just performing a simple online transaction.

And to make those decisions, to make those discussions happen, requires a common technical vocabulary. And that is at the heart of what the IANA functions are all about.

Now, the IANA functions predate ICANN. Really, they were born at the same time the internet itself was born. My predecessor, essentially, was the person with the record keeping book that any time there was a new device added to the internet or a new technology added to the

internet, he wrote it down. And he essentially served individually as the person you went to if you wanted to add something to the internet.

Now, that evolved. It stopped being a person in a notebook. It started being databases. That person, John Postel, eventually became a team of people working in academia. But it was in the 1990s. It was a recognition that the IANA functions were critical and really needed a home that was beyond one individual or one team. And that organization needed to have proper global oversight given how important it was.

So, creating a proper long-term home for IANA was really why ICANN was created in the first place. Excuse me. It's only Saturday and I'm already losing my voice. Is it Saturday? Is it Saturday? It is Saturday. It is Saturday. All right. So, yes, creating a home for IANA was the primary reason why ICANN was established in 1998.

Now, when we talk about those identifiers, I'll get to breaking it down in a little bit. Generally, we talk about those identifiers, we use a specific term. We call them protocol parameters because these are the parameters that go into the technical protocols that the internet runs on. But because they're so specialized and there's a lot of special considerations that go into them, we typically call out number resources or internet numbers as one specific category and also domain names that I think you all said you're familiar with.

But at its heart, IANA maintains the official databases, the official records of what protocol parameters have been assigned and for what purpose or to whom they've been assigned. And we do those

allocations within a team that I'll describe in a few slides. And we do it according to policies that have been established by the technical community.

So, I mentioned why these exist in terms of interoperability. If we're not speaking the same language on a technical level, then nothing's really going to work on the internet the way we expect. That interoperability is key to why we exist as a global organization. Can you imagine if there was an IANA function in every country and it just did its own thing?

The internet as we know it today wouldn't be an internet; it would be 250 different national networks that didn't talk to one another. That's why it's so critical that we have bodies like ICANN that are truly global in nature that allow for global participation because the internet really relies on that connectivity, that interoperability that works no matter where you are around the world.

We often think of the internet as free from sort of global control. And that is true, but no one, there's no entity that controls the internet at the global level that says, here is what you can and cannot do on the internet, here is how you must use the internet. But we do require that common technical vocabulary to everyone following the same sort of rule book in terms of the technologies for it to function as a global system.

So, the records that my team maintains are essentially voluntary. So, no one's saying you have to follow what IANA does, but everyone does because the internet wouldn't work if everyone didn't follow it, but it

is a voluntary system. But it's a voluntary system that's followed by software vendors, ISPs, businesses, application developers, really anyone that wants to connect to the internet follows the set of assignments and allocation databases that IANA maintains because they want to work with the internet as it is. The internet as it is, relies on these databases that we maintain.

So, the core deliverables we have, I sort of mentioned protocol parameters. I will say on protocol parameters generally, mostly it's just software developers that care about protocol parameters because as I mentioned before, it's sort of the technical plumbing of the internet. How a web browser negotiates language preferences for a webpage.

That's probably just of interest to people that write software browser software, that write web servers, but as an end user, as long as it works, I don't really care. And that's true for a lot of things that we do. It's the engineers, the software programmers that write the software that really have a stake in how it operates.

And when the technical standards are developed, IANA has a role in operationalizing those technical standards. And so, our community of interest when it comes to protocol parameters is really working with software developers, standards authors and the like to make sure that their applications are properly supported.

For names and numbers, it's different. So, for names, domain names, domain names are assigned on a very hierarchical system of allocation. Like there's a bunch of registries around the world,

registrars and the like. There's no one official record that contains every single domain name in existence.

Instead, it's a distributed database where different parts of the database are stored in different places. What my team's responsibility is when it comes to the domain name system is we manage the most sort of central database there is, which is called the root zone. The root zone is the official record of what is a top-level domain. Top-level domains you're probably familiar with. There's .com, .net, .org, sort of very old top-level domains that started with the DNS itself. They operate at a global level. They're available to anyone around the world.

There's been a lot of those added over time, .ninja, .diamonds, .info, .museum. There's roughly 1,000 different options for global top-level domains, often referred to as generic top-level domains. But there's another class of top-level domains called country code top-level domains.

Every country is issued with its own two-letter code. Here in India, it's .in. There's also, additionally, representations of country names in other scripts and languages. So, here in India, I know there's a lot of .bharat in different writing systems. So, for countries that predominantly do not use either Latin writing system, well, essentially don't use a Latin writing system, there's options available in other writing systems as well for representing the country name.

So, that list of what top-level domains have been assigned is the root zone, and that is operated by my team. Alongside the root zone is a

critical security element called the KSK, the key signing key. This is essentially the cryptographic key that proves that the root zone hasn't been tampered with. And so that key is operated alongside the root zone itself and provides integrity protection for the DNS. Now, let me talk about internet numbers for a moment. So, what is an internet number?

Well, every device that's connected to the internet, whether it's your laptops, whether it's a mobile phone, a smartwatch, but many other devices that you can think of, like practically everything is connected to the internet these days. How does the internet know when traffic is going to a specific device?

The way it works is that each device is given a unique number called an IP address. IP addresses are the way that when internet traffic, whether it's over a wire, over Wi-Fi, or what have you, works out exactly what device to send the transmission to. IP addresses need to be assigned to make sure that they're not conflicting so that each device has its own unique number.

And IANA is, again, responsible for that at the global level, which doesn't mean that we have a record of every single device on the internet. Instead, we maintain high-level global allocations that say, for example, here is a block of IP addresses that are assigned for the Asia-Pacific region. And then there is a system of regional and local assignments that are made.

So, for example, in the most typical case, there will be a regional internet registry for a region. IANA will give them a big block of IP

addresses. That regional registry will then, in turn, provide ISPs with a block of IP addresses. Now, an ISP could be the company that you pay for your internet access. It might be a large corporation.

In fact, this ICANN conference itself has its own block of IP addresses just for the conference. So, it can be a conference, for example, that gets a block of IP addresses. Then that network operator will individually assign IP addresses to customers. So, when you connected to this network this week, the first time you walked into the conference using another one of those secret protocols you don't normally see called DHCP, the conference network assigned an IP address directly to your device.

And so that will happen seamlessly without you noticing, but it was from an allocation of blocks that ultimately originated from IANA. So, that's number allocation in a nutshell, and that's another function that we perform.

Next slide, please. So, who's involved in the IANA functions? So, the IANA functions are, I mentioned at the beginning, sort of a part of what ICANN does. ICANN is ultimately responsible for the IANA functions, but does not operate them directly. There is another entity at play here, which is PTI. I'm not going to get into too much detail about what PTI is for this level of presentation, but PTI is a separate company from ICANN.

It is another nonprofit organization, and ICANN contracts PTI to perform the IANA functions. ICANN then oversees how PTI does the work. Now, why is it structured this way? Essentially, it's about having

an accountability mechanism that's separate from ICANN to make sure that the IANA functions are performed correctly. Hasn't always been this way. In fact, this structure that we see here began in 2016, but it was an evolution required by the global internet community that it be structured this way to ensure that the IANA functions are properly overseen.

Before 2016, IANA was just a department within ICANN, but now we are a legally separate entity. But the truth is today, we still operate like a department of ICANN. We operate in the same offices. We have the same legal team. We have the same IT team, the same HR team. So, a lot of it, the way we operate within PTI is the same as if we were a part of ICANN directly.

ICANN funds PTI. So, PTI, 100% of its costs are paid for by ICANN. But crucially, ICANN maintains all the accountability mechanisms to make sure the IANA functions are done correctly. There's a lot of oversight bodies that govern IANA, given our importance to the internet, given the criticality of the services that we provide. Of course, it's important that there is accountability too.

Imagine if we did our jobs just the way I felt was appropriate without accountability, then I'd probably make some bad decisions. But I am accountable to global bodies that oversee our performance to make sure we follow policy and we do our jobs well. And those bodies are operated by ICANN.

Next slide, please. So, I mentioned all this, I think PTI, Public Technical Identifiers, created in 2016. But PTI is the body that actually runs IANA.

The IANA team is employed by PTI, but essentially, we operate like a subsidiary of ICANN. PTI has its own Board of directors. Five-member Board. One of our Board members is Indian, Anupam Agrawal.

John Crane is also on the Board and I'm on the Board. So, that's three. This slide's a little out of date. We recently changed our chair. Tobias ended his term. He's been replaced by Lisa Furr, who is from Europe. And rounding out the Board is Xavier Calvez, who is ICANN's Chief Financial Officer. And I'd be remiss not to note the IANA staff. This slide is also a bit out of date, but essentially, we have roughly 20 people working in the IANA team, delivering the services that I just mentioned.

Three of which will be on site this week. Myself, Raquel, who's sitting right behind me, and Selena is currently stuck in an airport somewhere, but should be joining us soon as well.

So, the IANA operations, I think I really talked about this. We process requests, essentially. People want registrations in IANA databases. They reach out to us. We assess them against the policies, and then we implement changes in our databases based on whether they're eligible or not. Slides are jumping all over the place. Can you go back a little bit? All right.

We also act as subject matter experts, both internally to enhancing our systems and tools, but also to ICANN more broadly on all the different things that ICANN is doing. A lot of what we do is making sure we operate in terms of service level agreements. A lot of what we do is

operationally critical. So, it's important that we do it within established timeframes, and we do it in an accountable way.

So, a lot of what we do is measuring and recording what we do to make sure SLAs are adhered to. Generally, we operate in normal business hours, but for some critical operations like root zone operations, we have 24-7 coverage so that if there is a need to do an emergency root zone change, we can.

Some of the other functions sit within the 20 people I mentioned. We develop software to support the functions that we do, product management, that cryptographic key that I mentioned. I could spend an hour just talking about that, but I don't have the time. But we do a series of events called key ceremonies where we administer the key in a very public transactional way.

Yeah, next slide, please. And there's also sort of the meta work that we do to sort of, in a way, oversee ourselves to make sure that we're delivering the best service that we can. And this is a variety of different programs that we've implemented. We have a series of third-party audits where we have a third-party auditor come in, make sure that we're following policies and procedures, and demonstrating that we do to the auditors, and they provide annual reports to the community.

And this provides confidence to the community that we're doing our job in an appropriate way. We manage, obviously, risk within the company, making sure that any challenges to our operations into the

future are remediated, so we're in the best position to continue delivering our service.

Customer satisfaction is critical to us, so we have a variety of different ways we monitor and manage surveys. And in fact, there's a QR code, I think, at the end of this slide deck, where every time we present, we ask people to give their feedback on how we do our work. And that really feeds into our Continuous Improvement mindset, which is that even if our customers say we're doing our job well, we're always looking for what's the next thing we can improve so that next year we're doing the job even better.

Project management, we have a lot of projects going on. Community engagement is part of what we do as well. And so, all of this stuff needs to be coordinated and managed within our team.

Next slide. Here's some sort of like, this is almost the end of the slide deck for me. It's worth noting some of the things we do and don't do because there's a lot of misconceptions. I just told you that IANA runs all these critical databases that the world relies on. And there's an easy to take away that, well, IANA decides who gets a domain name. IANA decides who gets an IP address. IANA decides this, IANA decides that.

Fundamentally, we don't decide. We are checking that applicants meet policies, but it's not that we just, you come to us and say, I want to do this. And it's just our personal decision to do it. What we're

doing is taking community developed policies and coming full circle to this meeting and why we're here this week.

Policies will be developed in the corridors and the meeting rooms throughout the week, this week in various different community groups. Once those policies are ratified, they'll be given to IANA to implement. But IANA will be implementing language that the community has developed and agreed.

The community will say, here are the rules for when these kinds of things happen. And it is my team's job to follow those rules. So, my team doesn't create policy. My team doesn't set the rules of the road. We're here as an operational team that implements the rules that you as a community have decided.

So, what do we do? We create registries based on those policies. We then maintain those registries. We receive requests to add and remove things to those registries. And then we publish those registries. All of our records are essentially public records, but we don't create policy. We don't even interpret policy. We follow rules that have been set by the community.

If the policy is really vague and we can't make a determination in an objective way, we go back to the community and say, your policies are lacking. Please build your policies better so we can make proper objective decisions against those policies. And probably worth a call out, I mentioned Country Code Top-Level Domains. They're a very interesting case. There's a very particular set of policies that govern how ccTLDs are selected and operated. But essentially what it boils

down to is that the decision-making process for who runs a Country Code Top-Level Domain is performed within a country.

IANA's job is to recognize a manager, but not decide who the manager is. And I think there is a subtle distinction there. We recognize managers who have been selected through an appropriate national process, but we don't pick who the manager is.

If you come to us and say, I want to run .in, for example, I will say to you, show me that a process has happened within India to select you as the manager of .in. And there's a lot of details around how that actually happens, but essentially, we're seeking to recognize that an appropriate process has happened within the country under national law that is appropriate for that country. And only once you've satisfied those criteria will IANA recognize you as the operator of a ccTLD. But we're not deciding who gets to run a country code.

It's not as if a bunch of companies come or entities come to us and all compete for our attention and we go, I pick you. No, we tell them to go back inside their country, have a locally appropriate process within your country, come to a consensus decision in your country involving governments, private sector, academia, civil society, and whoever else needs to be involved in that decision. But once that consensus decision is made in the country, come back to us and we'll recognize that decision.

Next slide. Want to know more? Apologies, this isn't my slide deck, so I'm kind of looking at this for the first time as well. There's a lot more to know about IANA that I can't share in 30 minutes, but I'm happy to

go into. And we have other presentations we give in other forums as well. We have presentations we typically give at other meetings. We have resources on our website, iana.org, and I'm happy to provide you with pointers to it.

This slide is from the last meeting where we had a follow-up session where we talked about all this stuff. So, it's kind of, apologies, we didn't strike this slide off this time around, but I can't point you to a session later this week where we're going to be talking about all this, but I can point you to a recording of the last meeting that we did where we did talk about a lot of this stuff.

But with that said, I'm happy to answer any questions that you have and get into the details. Otherwise, please do approach me throughout the week, and I'm happy to talk to you one-on-one. But I already see hands shooting up, so let's do some questions and answers.

SIRANUSH VARDANYAN I know we only have 20 minutes, and I know that we will not have time to accept all the questions, but these guys are here during the week. Please use this opportunity to ask them. Please go ahead, Millenium. Come closer to the mic, yes. And then please be ready.

MILLENIUM ANTHONY Thank you, Kim, for the presentation. My name is Millennium Anthony. I'm from Tanzania. I'm an ICANN85 fellow. I wanted to know, you mentioned that IANA implements what the community

says, but it doesn't really make the policies. But I was wondering how does IANA make sure that they have operational independency, that they just don't take anything that the community says and implements, you know?

And then the second question was like, we know that the internet infrastructure is growing. How does IANA ensure that, or how is IANA ready to expand its scope apart from protocol identifiers, the domain names, and numbers? Yeah, thank you.

KIM DAVIES

So, I think in response to the first question, how do we not just do what the community says, but sort of make sure it works, I guess is at the heart of that question. It's multi-part. Firstly, when the community is developing policies, generally they're not doing it in isolation.

One of my team is usually participating as a liaison or a consultant to that group, providing our subject matter expertise. In fact, my next meeting today is to go talk to the ccNSO, the Country Code Name Supporting Organization, because they're looking at developing new policies. And my role there will be exactly that, to provide operational insight into their work so that when they develop their policy, it's informed by facts on how that would actually work in real life.

The idea there is that the community is advised by staff so that their policies are grounded in reality, they are operationalizable. And hopefully the policies that come out of that work, whilst we're not a

decision maker in the policy, having been informed by that input will be a practical one. But if it's not, there are other safeguards.

Policies, when they're developed by the community, don't just go into effect right away. Policies need to be approved by the ICANN Board. ICANN's Board, one of its fundamental roles is to approve policies before they're implemented. And the Board asks those kinds of questions as well. They say, is this implementable? And it's not just, is it implementable in a practical sense, like do the words on the page make sense? But maybe it's too costly, like a policy could just be completely impractical from a cost perspective.

And so, the ICANN Board asks those kinds of questions, then we'll only approve policies once it recognizes that it's not only implementable and appropriate, but it's funded, et cetera. And then it will prioritize it against all the other things that ICANN and IANA are being asked to do.

As for growing into the future, I mean, that's a constant source of discussion for me as the leader of IANA, our Board, like we look at a strategic lens to make sure that IANA is fulfilling the job it needs to do today, but also fulfills the job it needs to do in five years and beyond as well.

So, we're constantly asking ourselves, like what are the next challenges that the internet faces and how do we meet that challenge? Part of it is just organic growth. Our job continues to get bigger; it never gets smaller. There're always new technologies that are coming on the horizon that we need to support, for example. One thing that I

spent a lot of time on last year was that there is this new protocol for drone identification.

So, all the sort of civil aviation sort of entities around the world are working on a global standard for how drones talk to one another using internet standards. And so, there's a lot of work happening on how drone identification will happen and it will involve a registry of drones. And again, IANA plays a role at a global level, but then we make assignments down to national bodies that will do drone registration within a country level allocation scheme.

So, that's just one example of something that if you'd asked me 10 years ago would never have been on my horizon, but it is today. So, we continue to monitor and develop to support future needs.

SIRANUSH VARDANYAN Thank you, and we'll take one question from a virtual participant: I'm from Mongolia. How do emerging technologies like AI or automation help improve the management of internet registry?

JOHN CRANE So, any disruptive technology that comes forward of which today I think the most prevalent that people talk about is LLMs and generative AI. Now, most people, when they say AI, that's what they're talking about. They come with both pros and cons. So, any new technology that is disruptive, obviously there's the disruptive elements, but we do, in my other role, where I'm in charge of all the AI and IT, we are looking at what kind of tools can we bring in.

Now, it doesn't mean we're developing tools. We may just be purchasing tools with some form of machine learning in the background, whether it's an LLM or something else. So, absolutely, you can use these tools to be more effective. You can use them to help write code, but you need to understand that these are emerging technologies.

They are not fully flushed. So, we also need to understand the negative side of bringing in these technologies. Many of you will have heard of hallucination by generic AI, and there's a lot of evidence out there, and scientific research on how you can actually force them to hallucinate, how you can actually manipulate these tools.

So, when you're looking at all the kind of commercial tools for inside business work, we're no different than any other business. So, we are looking at what can we use, how can we use them. ICANN and IANA are not unique in that sense. We're a business like everybody else. We write documents, we keep financial data, we have projects just like everybody else does.

There is an interesting discussion ongoing at the moment, and we're going to bring out either a blog or an OCTO document in the coming weeks or maybe months around this, and that is how does generative AI and machine learning in general make use of the identifier systems? And is it different from other disruptive technologies?

Now, you will hear everybody tell you that it is, but a lot of the time when you dig deep down, you'll find that actually it's not that different, but where it does differ a lot is on scale and speed and just

the rapid change in the way these things are working. And then what are the kind of risks that they're going to bring to it?

And the other question that we have, which is really sort of relevant to that last question as well, will they use different identifiers? When you look at agentic AI, how do the agents communicate? Do they need to communicate? Do they need some kind of unique global identifier registration? We don't know the answers to these, but we are constantly looking from the ENIT side of what are the tools we can bring in, just like any business does, but also from the OCTO and research side and also the people in the policy side are looking at how this is going to change both positive and negative the environment we work in.

There are questions about how it will affect policy development. You know, if we're using these tools in policy development, what is the effect? Is it positive? Is it negative? Is it a little bit of both? When we talk about DNS Abuse, and if you talk about the security sphere in general online, there's a lot of discussion about both the negative and positive aspects of these technologies. I use the term disruptive technologies because AI is just the latest, right? And we always have to be aware of what's emerging and how it's going to change the environment we work in.

So, we are constantly doing like horizon scanning and seeing what the technologies are. And AI is absolutely one of the top sorts of areas that we're looking at today. It's like, how is this going to affect the way that we all work in the realm of identifiers? Because that's ICANN specific

area. And frankly, we don't fully know yet. So, this is still evolving. It's fascinating. Thank you.

SIRANUSH VARDANYAN Taruna.

TARUNA KAUR BAMRAH Taruna Bamrah for the record. My question is with exploding data traffic; how does IANA ensure the IP address allocation resources remain scalable without the fragmentation?

KIM DAVIES So, I think this is where the allocation schemes that we support that really everyone relies upon, we do not maintain a centralized database. There was a time in the 1970s and 1980s where there was a single list of every device on the internet and every time a person joined the internet, they went to a guy who added you to a list. But even in the 1970s was the recognition that this is something that will not scale.

So, today we operate a highly distributed model where there's no centralized place where there's a list of devices or entities. And even I would argue domain names is more centralized than most. Like IP addresses is truly devolved but there's so many layers of allocation.

Again, I mentioned the example where the network operator in this venue is making decisions about which devices are connecting to the

internet in this venue. And it's all hyper localized within this venue. And so, a lot of internet technologies work that way.

So, there's scalability tends to work well and that's why the internet's been so successful because in comparison to other technologies that failed, it is highly dynamic and it supports scalability just due to the way the protocols are designed. So there isn't a need for all the parties to sort of scale up to meet future demand because the network is resilient by design and allows to scale at the lower level without needing entities like myself to fundamentally change the way they work.

JOHN CRANE

The scalability of the availability of IP addresses has been an issue for about as long as I've been in this, like 30 years or so. So, predating ICANN, we moved from a system where everybody got large blocks of addresses when you got a network and we moved to what we call classless inter-domain routing where we were able to give out smaller in the IP version four round.

That wasn't something that ICANN did, that was something that was done at the time in the IETF and with the regional internet registries. And then when we realized that version four was not going to cut it for eternity because it's too small, we expanded the amount of address space by changing the protocol to IP version six and which a lot of the network uses now.

So, a lot of those changes are not necessarily in the policy realm where people are coming with technological solutions for those. And that tends to happen in the internet engineering task force. And then IANA gets to manage those new registries.

SIRANUSH VARDANYAN Thank you. Gentleman behind. Yes, please come closer to the microphone.

BRAJESH JAIN Thank you, this is Brajesh Jain for the record, in my personal capacity. There are certain IPv4 addresses which are not in registry system at all. Does IANA keep the record of those or what efforts it is expecting to make because such addresses which are non-RIH systems become say, I'm speaking from ISP background. A lot of security agencies ask that who is the owner of such internet address and ISP gets lost. Thank you.

KIM DAVIES I guess I'm curious what IP addresses you have in mind when you state that. I will say that by design, there are some IP address allocations that are made to a party, like to an entity to operate. And there are some IP addresses that are made for a purpose. And so, the purpose-based allocations are not by design allocated to a specific party to operate. Perfect example of this is private use IP addresses.

So, if I tell you a number off the top of my head, I'm curious how many people immediately recognize this number. 192.168.1.1. So, that's an

IP address that half the people in this room, if you go home and connect to your internet at home, you probably get an IP address that looks a lot like that because that is a block that's really designed for home networks.

So, that when you use your network inside your house, it uses this block and everyone uses the same numbers because they're reusable in each individual home network. So, that's by design. We don't track who's using that number because that is the way it's been designed so that not everyone has to get individual numbers in your house for your home networks.

That said, I mean, there are other uses of IP addresses where they're meant to be allocated to parties, but they're not. I think those cases are increasingly rare because there's more and more technologies around routing security that prevents that from happening.

The primary technology here is called RPKI where cryptographic assertions are made that a particular block has been assigned to a particular company and that company can make assertions in the routing system that only they control these IP addresses, which prevents other people from claiming them.

With all that said, a lot of the challenges around this, whilst they've been mostly addressed, have historically related to scarcity, that IPv4 is the numbering system that was invented in the early 1980s that we still prominently use today, but there's not enough IPv4 addresses to go around. There's only 4 billion possible numbers in IPv4, but that

problem is being addressed and there's been a lot of progress towards the next system, which is called IPv6.

IPv6 provides so many numbers that running out of them is essentially impractical. Gratified to see the adoption rates, if you look at some of the graphs that researchers publish, IPv6 adoption is around 50% today. So, we're almost at a point where it is the predominant way that computers communicate. So, IPv4, when it was the only game in town and we were starting to run out of numbers, there was a lot of challenges around people stealing blocks or reusing blocks because it was scarce and people needed them.

Thankfully, we're mostly past those challenges because IPv6 is available widely, used a lot. I can guarantee almost every mobile phone network today, even if you don't see it, is using IPv6 behind the scenes. So, it is very commonplace today and it's addressing a lot of those challenges. So, I don't know if that touched on the question, but that's my thoughts.

SIRANUSH VARDANYAN Suncica, please.

SUNCICA ROSIC Thank you. Suncica Rosic for the Record. I have a question about Okta. So, you mentioned that you developed a machine learning-based tool to classify malicious emails. And I'm assuming that as malicious attacks become more sophisticated, the distribution of the target variable changes and therefore there is a concept drift in the

sense that the relationship between the input data and the target variable changes as well.

So, my question is, how do you handle this shift in relationship as the attacks evolve? Thank you.

JOHN CRANE

Well, I don't. I have people actually understand this kind of stuff much more than I do working on this. So, we are looking at those kinds. That particular one was a sort of one-off project to prove a point. And we know that some people have taken our code because we told them they could. It was part of the purpose and are actually thinking about how do we implement this into mail systems, right?

So, we can actually do this larger. And we all have to deal with the fact that the landscape that we deal with when we're dealing with malicious organizations and people is constantly changing. And we're in this constant battle of technology with them.

So, what I would suggest if you're really interested in that specific project is that I actually introduce you to the scientist that worked on it. And you can actually have that conversation because he will completely geek out with you on it. And I'm sure he would love to like maybe run the experiment again and update it. So, let's talk afterwards and I'll make that introduction.

SIRANUSH VARDANYAN A quick question from virtual participant. Is there an independent body or checks in place to make sure IANA operates without any bias in the recognition?

KIM DAVIES Yes, that's a great question. So, practically everything IANA does is subject to some kind of external oversight or scrutiny because the communities we service are complex. Like there's no singular body but there are all sorts of overlapping mechanisms to ensure that.

In terms of community bodies, one that immediately comes to mind, one of many is there's the customer standing committee, which as the name implies is a committee of our customers that monitor our performance.

I mentioned in my presentation, there's also a role for third party auditors. So, we have a company that we've appointed that is completely neutral. They're not from our industry that their only job is they look at our policies and they look at our decisions and they make sure that our decisions follow the policy. So, there is a complex framework of oversight that I think is quite comprehensive. And frankly, the community continues to discover new ways to oversee us. So, I'm sure there'll be more in the future as well.

SIRANUSH VARDANYAN Thank you and we'll take the last question very briefly, Jia, please.

PETER JIA WEI CUI

Yeah, hello, I'm Peter. And I really appreciate this session and I want to know more about how the office of CTO to reports and all these kinds of practices can really influence the policies because I'm from legal background and I think it's always very challenging for the people from policy side to understand the technical details that is able to be transmitted and translated into the policy practice.

JOHN CRANE

Well, that's our job, right? So, we try and write these in a way that they're actually consumable by non-engineers, right? We're not writing these for the engineers. We're actually writing them for the policy people. And we don't just write a document and publish it.

There is an insane amount of review for every document. The guys writing the documents will tell you there's too much review. They should just be allowed to publish, but they don't. So, we start by having our engineers. I'm going to take an example that comes out of our more academic research. One of the first things we will do is our engineers will check it for actual operational integrity.

Is this the theory or is this the practice of how things actually work on the internet, right? So, they'll start with that and then we will run it through other experts. We will then go maybe for a, well, we always go for a legal review. We go for many of the reviews. Even if it's a legal review, it won't just be about the legality. They will be saying, we don't understand this, right? Because as a non-engineer, they can look at it and go, we don't understand the word.

We understand the word, but we don't know if you do because you've written this and it doesn't mean the same thing for us. So, we have those conversations. We also, before we finally publish, we go through our communications group and our technical writers and we go through various edits to make sure that it's readable.

So, the whole point is to make sure that actually as a policy person, you can read that and have a decent understanding. We cannot make something that is technically highly complex, easy for the lay person, but we can bring it down a level so that then you can come have a conversation with us or you can have a conversation with another engineer and they can explain the nuances.

It's really, really hard to get it all the way down. Just like it's hard to explain the policy nuances sometimes to non-policy people. So, we really put a lot of effort to strive into readability. And when we're doing things that we hope other people will replicate, that tends to be more aimed at the scientists, et cetera, we make sure that we have repeatability of what we're doing. And one of the final reviews that happens is always it lands on my desk.

And my review is always; can I put my layman's hat on and understand this? I do a lot of communication with governments and with researchers and policy people. So, I have to be able to communicate in that manner. And then I go through with every line and I ask myself, is this a fact or is this an opinion? And I try and make sure we remove all the opinions from there that we can, or if we have an opinion that

we state that it's an opinion. And so, it goes through so much rigor that by the time it gets published, hopefully it's very usable.

If you're finding documents that are not usable, reach out, our guys are really friendly, reach out to the author or reach out to me or reach out to Kim if it's one of the things they're doing and we will help you understand it because that's what we're paid to do. That is our job.

SIRANUSH VARDANYAN Thank you, John. And thank you really for your time being here and please reach these people during this week if you have any questions, you will see them all around. We will have the opportunity to talk more about internet identifiers and how domain name system works tomorrow.

We will be having several sessions from our technical team about this tomorrow. But with that, I would like to thank you and let everyone know that all the presentations from today and tomorrow's sessions are already posted and available for you on the meeting schedule. Whenever you go to the specific session, click on the details, you will find the presentations which had been shown during the session. So, they are available for you to download and round of applause to our presenters and panelists today.

Thank you very much. I will see you all back at 13:15 in the same room where we will talk about how policy works in ICANN. And with that, this meeting is adjourned. See you soon.

[END OF TRANSCRIPTION]