ICANN82 | CF – How it Works: IROS and IANA
Saturday, March 08, 2025 – 10:30 to 12:00 PST

SIRANUSH VARDANYAN

Thank you. Hello and welcome to the session How It Works. We will talk about IROS and understand what does it mean and IANA. And we have very interesting people here to present about those topics. My name is Siranush Vardanyan and I am the remote participation manager for this session.

Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior and the ICANN Community Anti-Harassment Policy. And those links are posted in the Zoom chat. During this session, questions and comments will only be read aloud if submitted within the Q&A pod. Interpretation for this session will include English, French and Spanish. If you would like to speak during this session, please raise your hand in Zoom. And when called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use a physical microphone to speak. Please state your name for the record and the language you will speak if you speak a language other than English. Remember to speak at a reasonable pace.

And with that, it's my great pleasure to introduce you the team who is behind of technical pillar of ICANN activities. So how technically

ICANN works. And my pleasure to introduce John Crane, the head of identifier research operations and security function in ICANN. And John will introduce his team as well. John, the floor is yours.

JOHN CRANE

Thank you, Siranush. Unfortunately, I can't stay for very long, but I did want to come and say hello. So you can put a picture of my face and put a target on it if you need to find me. We do like to come and talk to the fellows in the NextGen and the folks who are new to ICANN.

I will ask one thing of you, and I will ask this of you throughout your time here at the ICANN meeting, especially once I'm not at the table. Please ask hard questions, all right? So don't be afraid to ask the questions. You will see a lot of my leadership team here. They are the people who actually do most of the work and lead most of the work. I'm just the person that signs pieces of paper and stuff.

So with that, I want to hand over to my deputy here, which is Matt Larson in the middle. Matt, do you want to take the slides? And I will run to my other meeting. Thank you. And find me. I will be around.

MATT LARSON

All right. Thank you, John. And welcome, everyone. So as Siranush said, we're here talking about the IROS function, Identifier Research Operations and Security. So we're going to go over who we are and what we do.

All right. So IROS as a function covers two different teams. One is the Office of the Chief Technology Officer, and the other is the IANA functions. So we're going to start by first giving an overview of the Office of the CTO, or OCTO, and then we'll hand it over to my colleague, Kim Davies, to talk about the IANA functions.

So here is OCTO in a nutshell. These groups also correspond to the managers under John Crane. So we have an operations group that handles functional operations, coordinates different activities that we do, like the ICANN DNS Symposium, the Registry Operators Workshop. The project management for OCTO projects sits in the operations group, and they also manage the money. So they're responsible for the budget. Technical engagement is my colleague, Adiel. So I'm going to hold off on any other description of that and let him go into that in his slides.

I lead the research team. We actually have two research teams. We have a more general research team with a wider focus that I lead, and then my colleague, Samina, leads a group called the Security, Stability, and Resiliency Research Team, and they focus, as you would guess from the name, on a more narrow focus that focuses on security and stability and resiliency. That being said, all the researchers, both teams, collaborate closely, and we tend to work on projects of all different kinds.

All right. So let me hand it over to Adiel to talk about technical engagement.

ADIEL AKPLOGAN

Thank you, Matt. So I will talk about technical engagement, which is-- yeah, okay. So you see technical engagement as the public-facing arm of the IROS department. My group, which is a regional group, by the way, has people working in five different regions of ICANN. And our main goal is, one, in one hand, to take the work of OCTO team, the research team, the two research teams, by the way, and make sure that the community in general are aware of those works but also contribute very deeply in the capacity building activities that come out of IROS and ICANN in general when it comes to the DNS operation.

We work very closely with all the technical organizations globally to ensure that we maintain a very fluid collaboration and cooperation with them on technical matter and also explore possibility and opportunity for collaboration with them.

In our quest to ensure that best practices are adopted by operators globally, we have launched a few years ago KINDNS, which is a framework that we developed with the community to help DNS operators to run their operation focusing on key best practices of the DNS. So it provides the key best practices that we expect and we think that are critical for a secure and reliable DNS operation in general. And our capacity building activities today are done on that framework in general. We also work with all other departments within ICANN to make sure that the technical messaging, the way ICANN portrays technical posture is coherent and help also make the connection where needed.

So I already explained the two aspects of what we do, the capacity building and outreach aspect, which is very original. If you want to understand, know and work a little more practically on the DNS side, you can approach anyone from my team in your region and organize a workshop, a hands-on workshop. We also collaborate very much with the ICANN-learn team to provide them content related to the DNS and technical matter in general.

So besides that, also we engage very actively, broadly with the technical community, as I mentioned before, taking the work that the research team is doing or other work technical within ICANN and helping to do outreach around them. We also work with ICANN constituents. We are, for instance, currently working with the ISPCP group, where we help them deliver some very specific technical training on the DNS. And this is also part of what we do. And we are exploring the way of doing that with other ICANN constituencies as well.

So there is a lot around that. And I know that this group and you may be interested a lot in capacity building and that. We have a few of our team members here today, particularly David that covers North America, who is in the room, and Nicholas that covers Latin America, who will be around as well. So if you have a question around capacity building or activities in your region, feel free to talk to us.

So that's it. This is the distribution of the team globally. I will go back to you, Matt, to talk about the research team.

MATT LARSON

Thank you, Adiel. So I won't read the slide to you, but the goal of the research team basically is to do research on the Internet system of unique identifiers that ICANN is responsible for coordinating. We have access to a lot of interesting data. Some of that is, for example, the data from the ICANN L-Root server. We work very closely with the DNS engineering team within ICANN that operates L-Root, and we're able to get access to that data stream, which is interesting for a lot of different projects.

We also have a partnership with APNIC, and we have access to some of the data they get from their Google ad-based measurement platform. And we've done some interesting projects where we look at the intersection of those two data sets.

We also participate-- it's not on the slide-- we participate actively in the IETF, the Internet Engineering Task Force. We help with DNS standardization efforts, so we help extend the DNS protocol. And then one final thing is we do represent ICANN Org on the RSSAC, the Root Server System Advisory Committee.

And this is a list of some of the projects that we've been working on recently. A long-running project is the Identifier Technology Health Indicators, or ITHI. And so we've developed metrics to measure the health of the Internet's unique identifier system. And we've been managing these for, collecting these metrics for, coming up on 10 years. So the goal of the program was to have long-term data to

watch not just how these metrics are in real time, but also how they're changing over time.

In general, a lot of our research focuses on DNS authoritative servers and resolvers, all aspects of that, of their behavior, how they're concentrated and distributed. We don't work exclusively with the naming identifiers. We also work on the Internet address side of things. We have someone who works a lot on RPKI, which is how Internet address space can be cryptographically signed to know who owns it. And so that's beginning to be deployed, and so we also work on that.

And another interesting thing we've done recently, you may be familiar with DNSSEC, the DNS Security Extensions. That's taken a long time to deploy, and it's still being worked on and actively changed, and one of those things is error reporting. So we've made it possible for DNSSEC errors to be sent to a central location so that it can more easily troubleshoot those errors.

And then I don't have a bullet on here, but we're also looking into alternative names. For example, based on blockchain, those names are getting a lot of interest right now in the ICANN community, and so we've looked into how they work, how the different systems that are popping up, how they operate.

All right, and my colleague Samaneh is not able to be here, so I will go through her slides quickly and talk about her team. So as I said a moment ago, her team focuses typically more narrowly on

research topics that deal with security, stability, and resiliency of the Internet's unique identifiers.

Let me skip ahead to some of the activities that they're working on. Most recently, her team has been involved in developing a new service called ICANN Domain Metrica, which is a way to collect metadata about domain names. And the first major chunk of data that's in Domain Metrica now deals with how domain names appear on reputation block lists, or RBLs. This was just recently opened to the public back in February, and in Tech Day there's going to be a presentation about it, and everyone can register and try out the platform.

The SSR team, their research, they tend to write academic papers that are suitable for peer-reviewed conferences and journals. So you can see here some of the recent topics that they've been working on and that have been submitted and accepted at different events. And in general, they tend to focus a lot right now, recently, on DNS abuse. So detecting domain names that are being used for abuse and working on ways to prevent that.

And then the final thing I want to talk about is something called the OCTO Document Series. We realized a few years ago that we had documents that we wanted to publish and no good, consistent way to do that. So the OCTO Document Series we created, it's similar to the document series that SSAC has, that RSSAC has, and we thought, well, why can't we have a document series?

So indeed, at this point, we've published 41 documents and counting, and these vary in length and in complexity. Some are

more overview of a topic, some are more in-depth. The topics vary considerably. Here, for example, are the most recent five documents that we've published. I guess I'll specifically point you to OCTO-040 and 039. Those talk about blockchain technologies and then name systems that use blockchain technologies.

As I said a moment ago, we realized that there was a lot of interest in the community about this topic, but not a lot of information available. So Paul Hoffman on my team decided that really there should be information on just how blockchain technologies in general work, as well as some specific characteristics of popular blockchains. And then he also went into detail about how different name systems use blockchain technology. So I would point you to those documents if you're at all interested in that topic, and there's the URL for how to find them.

So you can find us in addition to the OCTO document series. We do publish on the ICANN blog, and then you can see us all over, particularly the folks on IDL's technical engagement team at various speaking engagements. And I believe that's the last slide. So would you like to hand it over to Kim and then do questions at the end or?

SIRANUSH VARDANYAN I think we can take some questions for you, for your team, and then I will hand it over to Kim. Any questions? Yes, please, Rolla. Go ahead.

ROLLA HAMZA: Thank you and greetings all. I'm wondering about how to engage or participate in the research activity or technical activity. Is that possible? Thank you.

MATT LARSON: I think probably doing just what you're doing, coming here, and I'd be happy to talk afterward, probably engaging with individual members of OCTO and the research team. So talk to me afterwards, I guess.

SIRANUSH VARDANYAN: Excellent. Any questions? Yes, please, Saiid. Go.

SAIIDNAJIB SAIDISLOMZODA: Hello there. My name is Saiidnajib. Could you tell me more about DNSSEC?

MATT LARSON: Sure. How much time do you have? So very briefly, when DNS was designed, there really wasn't a lot of thought given to security. And so plain old DNS, you send a query to a resolver, and that query is in one packet, and an answer eventually comes back, and that's also in one packet. And so we realized pretty quickly that that allows for all kinds of bad things to happen, that how do you know that the answer that you got back is really the answer to the question you asked? Is it possible that someone else has sent you an incorrect answer trying to fool you?

So that started a process that took years and years and is still being worked on. The DNSSEC protocol adds security to DNS based on cryptography. So information in DNS is cryptographically signed by the owner of the data, and then anyone who looks it up can cryptographically verify the signature on the data to know that it's valid data or not. That's the really, really simple version of DNSSEC.

SIRANUSH VARDANYAN Yes, please, Alex, and then we'll come to you.

ALEX DANS Can I just add that we have a dedicated page on ICANN.org on DNSSEC with infographic, Octo-publication, many resources, and a link to blogs and announcements that we published. It's available in English, in Spanish, in Portuguese, in Chinese, in Arabic. So just put DNSSEC in the search bar on ICANN.org, and you have all those resources available.

SIRANUSH VARDANYAN Excellent. Thank you. Yes, please.

MABDA HAERUNNISA FAJRILLA SIDIQ Hi, this is Mabda for the record. I was interested in your research on blockchain-based naming system. I understand that you have a report out that you mentioned before, and I was wondering how has that informed ICANN's position on blockchain-based naming system. I understand that this is something that ICANN must be threading very carefully around. So I was just wondering how has

your research informed ICANN's position, and going forward, how will that inform ICANN's position down the line? Thank you.

MATT LARSON: Well, I think the best answer to that is that we did research just to understand and make that information available. ICANN org as a whole, I think, is still looking at information from all places, the research we did as well as what's happening in the community to develop what the org thinks about it. So I would say that our research is just part of the foundation for the decisions that the org will make.

SIRANUSH VARDANYAN: And I have two people who joined online and who are asking about how they can get engaged in the research. Is there an email they can send, a request, or how?

ADIEL AKPLOGAN: Yes, I mean, to answer that last question, you can send any request to octo@icann.org related to any need or interest in engaging with our team, and somebody will reply and get in touch with you if you need something. On the blockchain thing, I think that question also comes very often, and as you mentioned, it's something that ICANN needs and has to trade very carefully.

But something interesting that I keep saying when the questions come to us is that we need to understand the naming system in the blockchain compared to the naming system in the DNS. The two

have two different purposes, and the purposes are not the same. And when you start understanding from that premise, you can clearly decouple this from ICANN perspective, because what we are interested in mostly is the security and stability of the DNS, using the naming and the number for accessing the Internet. Blockchain names have other purposes, which are naming, but not directly related to Internet access information.

SIRANUSH VARDANYAN Yes, please. Go.

MOHAMMAD ATIF Hello, everyone. So my name is Mohammad Atif. I was just curious to know about how do you amplify your research work? Like you have mentioned about these many technical papers or thought papers being published, but how do you collaborate with a larger ICANN body, and how does that research work get implemented as maybe policies, or if you can shed some light on that, then I would be grateful. Thank you.

MATT LARSON Sure. Well, one way is coming here, coming to ICANN meetings and presenting on them. We have a periodic IDL. I'm under pressure. I'm forgetting the name. The series we do about the evolving-- Yes, thank you. We have a series that we do at some ICANN meetings called Emerging Identifier Technologies, where we've talked about just that, identifier technologies, like, for example, blockchain-

based ones. So sometimes I guess what I'm saying is that we present our research at ICANN meetings.

The OCTO document series that I mentioned has been a huge way to make our work visible. Also, Samaneh's team, my colleague who has the SSR research team, as I mentioned, they do a lot of work in more formal academic venues, peer-reviewed conferences and journals. And then I guess you'd ask how it engages-- What was the second part of your question, please?

MOHAMMAD ATIF: How it engages into policy outcomes, maybe of different forums, or maybe you'll find some vulnerability assessment reports. So does ICANN separate bodies that have-- Do they take your research into consideration and set standards and policies on the basis of that?

MATT LARSON: I think the answer is yes, to a certain extent. Certainly, that's what we aim to do, is to make information available so that other parts of ICANN org and the community can have our, hopefully, objectively-based information. We try to just say, here are the facts, here's what we've seen. And then really anyone within the org and in the community can use that for what they need to do to understand more.

ADIEL AKPLOGAN: Yeah, and that is also very important. The work that Matt's team and Samaneh's team does is to actually dig the fact out, do the

research and put it out there for the community to know. So talking about the policy, as you know, ICANN is a bottom-up organization. So the policy development part is not the org. It's for the community to look at what comes out from those research and identify if there is any need for policy or discussion or change, and come up with them, discuss them within the community, and go with that. It's not the org.

And one of the works that my team does as well, in collaboration with the research team, is to take the outcome of those work and help outreach, explain them and talk about them in different regions. Some research outcomes may have different interpretations in a specific region or have highlighted challenges in a specific region. And that's why the regional team is important, because they help amplify that and work with the regional community to understand some of those issues.

SIRANUSH VARDANYAN Thank you, Adiel. And we will come back to you with more questions by the end. So let's give an opportunity to talk about the next topic, IANA.

And it's my pleasure to introduce Kim Davis, who is Vice President for IANA Services and President for Public Technical Identifiers. And what does it mean? Kim, having the extraordinary ability to explain the complicated technical stuff in a very simple way, I will give the floor to him to explain what does it mean and how IANA

works to people like me, who has no technical background. So, Kim, the floor is yours.

KIM DAVIES

Thank you, Siranush. Good morning, everyone. So I'm going to talk to you about the other half of the identifier research operations and security function of ICANN, which is IANA.

So what is IANA? Let's start with what it stands for. It is the Internet Assigned Numbers Authority. It's a term that's been around for many decades. Whilst we do more than numbers, certainly numbers specifically is our bread and butter, and I'll get to that in a moment. Essentially what we are is the official record keeper for unique names and numbers that are used on the Internet.

Now, it's not every single unique name and number used on the Internet. It's those that are relied upon by core Internet technologies to interoperate. So when we connect a device to the Internet, we want it to work. We want it to work with other devices on the Internet. Having that common language that devices talk to one another is really what the IANA function is about.

The IANA functions predate ICANN. ICANN was actually established in the late 1990s to be the institutional home for IANA. Prior to that, it was operated under various contracts with the U.S. government in academic circles. But throughout the 1990s, with the growth of the Internet, the rapid expansion, the commercialization of the Internet, it was recognized there needed to be a better kind of organization that took in multi-stakeholder inputs from around the

world to oversee IANA, and that's really the thrust behind what caused ICANN to be created.

Now, IANA is essentially a maintenance organization, very operational, and we maintain large sets of records, but according to policies that are set by the community. These policies come from standards organizations, they come from essentially the kinds of people who are at this conference this week. The policies that are decided through our multi-stakeholder process will in turn inform how IANA operationalizes its day-to-day work.

So I talked a little bit about unique identifiers, and I think before I get any further, it's probably worthwhile just to step back a moment and think about, well, what is a unique identifier? And I think the best way to do this is just to conceptualize something we do every day, probably most of you in the room have done in the last few minutes, which is opening up a webpage. It's something that we do without thinking. We type in an address and we hit enter, and then within the blink of an eye, we have the information that we requested. But underlying that, as many of you know, is a very complex series of transactions with transmissions happening around the world often to request and then retrieve information, process that information, and that all needs to work in a coordinated manner for the end result to work as we expect.

Some of the identifiers that you see when we do a simple transaction like that, you can see with your own eyes, but many of them you don't. They're hidden behind the surface. So let's start with what you can see. The most common identifier that we're

involved with, in fact, many of the sessions this week on this very topic is a domain name. A domain name is a human-readable identifier that is a memorable way of remembering where something is on the Internet. It can form part of an email address. In this case, it can form part of a web address. And it's basically what you would type into a web browser to get to a place.

But that's not the only unique identifier you see on the screen here. The words in blue is the domain name, but at the start of that address is actually something else. It says HTTPS. This is another kind of identifier scheme called a URI scheme. And it's one of many ways to define the method of communication with a remote server. HTTP or HTTPS, which the S stands for secure, is the most common way of retrieving things on the web today. But it hasn't always been the case. And there's other different combinations there that speak to older technologies, perhaps not in more common usage, or different technologies. For example, the kinds of URL or URI that's used to make a connection for a real-time application would be different.

Now, when you send off the request for a web page, it needs to go somewhere that actually has the information. And ultimately, it's going to go to some computer, some device, some server, somewhere on the internet. You've typed in a memorable address, like a domain name. How does it actually know where to get from that domain name to the actual server that has the information that you require? In essence, that domain name is converted into a unique identifier that applies to each individual device connected to the internet. This is called an IP address. An IP address allows

data to be transmitted across the internet and reach a very specific endpoint, in this case, to retrieve the contents of the web page.

Now, there are literally billions of devices on the internet. How you could possibly track where every single device is on the internet at any one time, it's just not practical. It's not feasible. So the way internet routing works is that big groups of devices are organized into what you can think of as neighborhoods. For example, this hotel, this conference, is one neighborhood. So rather than working out where each individual device is here physically in this conference, in this hotel, any data that's being transmitted to one of the devices on the hotel is actually just sent to another identifier called an autonomous system.

Once that traffic gets here, locally the decision is made how to get it to the individual device. But the device at the other end of the world that's trying to connect to your device does not know exactly where your device is. It just knows that you're within an autonomous system, the one that you happen to be in right now. So an autonomous system could be an ISP, could be your company's network, could be a data center, and certainly at a conference like this it could be the hotel that you're in.

Now, each individual device or server on the computer can do more than one thing. It could serve web pages, it could receive emails, maybe it takes video calls, maybe it has a file server on it. There's any number of things that could sit on one individual computer. How do we decide what the transmission is designed to do? We do this with another form of unique identifier. Well, two actually. One

are called service names and another are called port numbers. So when you make a connection to a server requesting a web page, you're connecting to port number 80, which port 80 is the port over which web pages are transmitted. However, if you were sending an email to that exact same server, it would go to a different port. In this case it would be port 25, which is where emails are sent to.

Okay, so we've made a connection to a server, we've gotten back a web page, all sorts of data is inside a web page, right? There's some text, maybe images, maybe a bunch of advertising that we don't really want to see, but there's all sorts of mixed media that is coming through in one individual web page. How does our web browser know, well, this part here is an image, but this part over here is an animation, this part over here is audio, this part here is an interactive form?

Essentially, little chunks of the information that are sent back are all tagged with a different kind of unique identifier called a media type. The media type then instructs the receiving computer on how to handle that transmission. And it's not just used for web pages. For example, when you put an attachment on an email, a media type is used to signal to the receiver what kind of content is in that attachment so it can work out how to decipher what it's receiving.

And there's other parameters that are used as well. Just two quick examples. You might have noticed that if you contact a web server and your preferred language perhaps is in English, you get back the content in your preferred language. How does that work? Well, usually it's because your browser will transmit its preferred set of

languages in order of preference to the server. And the server, if it has that translation available, will adhere to your preferences.

So if you ask for, for example, Russian as your first preference and the server knows, hey, I've got this content in Russian, it's going to send it to you in Russian. If it doesn't have Russian, it has your second choice, your third choice, it can do that. And it does that through another kind of unique identifier by sending what's called a language tag in a header. And the header is defined through what are called header fields.

So I've just raced through some of the unique identifiers used for just that one thing, requesting a web page. But already we've gone through all the things on the screen here. We've used a domain name, we've used a URI scheme, IP address, AS number, service name, port number, media type, header field, language tag. And that's not an exhaustive list. There's all sorts of identifiers that you use and little facets of the transmission that we just described.

So what do these all have in common? If any of the definition of what these identifiers weren't globally understood and recognized, then the internet just wouldn't work. You would send a transmission to another device on the internet and it would, frankly, be perplexed. If port 80 didn't always mean this is where web servers live and you tried to send an email to port 80, it's not going to work. If you typed in a domain name here and you got to one company, but you typed that same domain name in when you're at the local coffee shop and it took you to a completely different company, then the internet wouldn't work.

All these things require some level of global coordination to make sure they have a common meaning that means the same no matter where you are in the world. And that's essentially what IANA is about. IANA's responsibility is to manage unique identifiers globally. And I'll talk a little bit about how that happens as I go on through the presentation.

And here's just a flavor of some of them. I think, by last count, we had about 3,500 different types of identifiers that we manage. One is domain names. That's a big one. IP addresses is another big one. But there's this long tail of lots of different identifiers that are used in day-to-day transmissions.

Now, most of the kinds of identifiers we manage work by direct allocation. This means that parties come directly to IANA and we issue them identifiers based on their need. For example, if you're an inventor, you've invented some new technology and you need identifiers to support that technology, you can come to us. Now, the way you come to us is twofold. One is, if you're creating a brand-new technology, you actually standardize that technology first. And you do it in a forum called the Internet Engineering Task Force, or IETF. This is the primary forum in which internet standards are developed.

And in the course of writing your technical standard, you write inside that document instructions for IANA, how these identifiers will be used, how they'll be assigned, and there'll be an explicit section in your technical standard that explains what IANA's role is in the administration of the standard. If it's an already existing

standard, something that everyone uses day-to-day and you just want to augment it, often you can just come to IANA directly and ask us for an allocation and we'll just give it straight to you.

Now, there are a lot of domain names and there are a lot of IP addresses. Many people have domain names, everyone has at least one IP address. If you have a phone, a laptop, a smart TV, all sorts of things, you probably have lots of IP addresses. It frankly would be impossible for IANA to directly allocate individual IP addresses to everyone in the world, or to allocate domain names. So long ago, these technologies were designed with a hierarchical allocation mechanism. In both cases, domain names, IP addresses, and those autonomous system numbers for that matter, IANA's just responsible for the highest level of allocation. So we never do direct allocation to an individual end user of the internet. Instead, we allocate, at a very high level, chunks of namespace or chunks of address space to other organizations, who then in turn allocate them below themselves.

So domain names is an obvious one. We manage the root zone, which is the highest level of the domain name system. We only allocate top-level domains. Then in turn, a registry that operates a top-level domain will make sub-assignments within that namespace. On the same token for IP addresses, we allocate these big blocks of IP addresses to other organizations called regional internet registries. You can see the map on the screen. There's five of them, one for each continent-sized region. And they in turn have policies and procedures where if you want an IP address here in

North America, you would go to ARIN. But if you're in Africa, you would go to AFRINIC, for example.

We've given them these big blocks of IP addresses, and they then divide those up into smaller blocks for regional allocation. And how we make those allocations is not for us to decide. Again, those multi-stakeholder policies developed here at ICANN conferences and at other like-minded organizations then instruct IANA how to do those kinds of allocations.

Now, many of our registries are simple. They're like spreadsheets, like here's a number, here's who it's assigned to, for what purpose. But not all are, and some of them are actually very complex to run. The DNS root zone is one that is particular. It's an operational database that is regularly contact with thousands of queries a second from all around the world. And to support that requires a lot of technology, a lot of infrastructure, a lot of support to make that work. We heard about DNSSEC a few minutes ago.

DNSSEC adds this cryptographic protection on top of the DNS as well. But the secret keys that are used to do that encryption need to be very securely protected. And that's one of IANA's operational roles, is to maintain what's known as the private key for the trust anchor of the DNS, like the very top, the very most central cryptographic key from which all other trust derives.

We're operationally responsible for maintaining that key and maintaining it in a secure way. I don't have time to get into the details on exactly how we do that, but it's very interesting. The way we do it is that we have these public key signing ceremonies every

three months where we bring in security experts from around the world who monitor and advise us on how to operate that key. And just in the interest of transparency, it's something that we put online. You can watch it on YouTube if you want. We have them scheduled in advance. And if you happen to be in town when we're doing one, we also invite people to attend it in person as well.

So without getting into too much detail, some of the IANA registries we maintain do require a lot of operational effort to be applied to them to make sure they're operated in a correct manner.

Accountability is very important to us. The role that we have in IANA, making these allocations, can have serious repercussions if we do it wrong, if we're not meeting the needs of the global Internet community. So we have a lot of mechanisms in place. We have service level agreements to make sure when we do changes, when we process requests, we do them in an appropriate amount of time, and that there's accountability that we're doing them accurately in a timely way. We have all these accountability mechanisms in terms of oversight by the community.

There's various different committees and organizations and processes within ICANN to constantly monitor how IANA does its job and to report to us if changes need to be made. Should a customer be unhappy with how we do our job, there's remediation and escalation procedures so that they can escalate and get the attention they desire to the problem they're experiencing.

We bring in a third-party auditor who does various different tests to make sure our operational controls are effective, and issues reports

so the community can have confidence that we're doing it well, we're doing it appropriately in accordance with those controls. We survey our customers. We're regularly asking, how did we do? We do annual surveys of the community to get feedback as well. And we have all sorts of reporting available on our website, and there's a few screenshots of it there, where if you're so inclined, you can drill down to all the specifics of how we do things, what we did, etc. It's all available on our website.

When we look at accountability, we generally divide it into three main areas. Protocol parameters, number resources, and domain names. And the reason we divide it into these three sections, it's a little arbitrary to do so, is that represents the three key areas of oversight and accountability that we have with the community.

Now, I've gone all this way through without explaining what PTI is. When Siranush introduced me, she said I was the president of public technical identifiers. So PTI is actually an organization that runs the IANA functions. We're part of ICANN, but a legally separate organization. This is the result of a decision by the community back in 2016 that there should be some legal separation between the IANA functions and the remainder of ICANN. The reason for this was to allow certain oversight and accountability mechanisms to be more effective. So what we are today is PTI, a nonprofit organization that is affiliated with ICANN. It's not that we're independent by any stretch of the imagination. There are close ties between PTI and ICANN, but nonetheless, it is a separate legal entity from ICANN.

PTI hires all the IANA staff. The IANA staff are headquartered in Los Angeles, and we currently have 20 staff. Four of us are here on site this week, so Marilia, Tyler, and Selina. If you run into any of us, we're very happy to have a chat with you. And PTI has its own board of directors. It's a five-member board. It's very small compared to ICANN's Board. There's two community members appointed by the ICANN nominating committee and three staff members of ICANN, including myself.

So what's ICANN's role? If PTI runs the IANA functions, what does ICANN do? Well, ICANN is ultimately responsible that the IANA functions are actually done. So what ICANN does is contract PTI to do the work, and then ICANN in turn oversees PTI's performance and makes sure PTI does its job well. Now, PTI is a very small organization compared to ICANN. We don't have all the internal resources we need. Our focus within the PTI team is doing the day-to-day work. But we do need support, things like legal, IT, HR, finance, etc. So we do use ICANN's departments that do those functions as resources that we can utilize.

ICANN also provides all the funding for PTI. So PTI never charges for its services. If you want a top-level domain, you're never paying IANA for it. If RAR wants a dress block, you don't pay for it. If you want a media type, you don't pay for it. All IANA services are provided for free. But obviously we need to be funded from somewhere. So our funding comes generally from ICANN, and IANA funding is budgeted through ICANN's budgetary processes.

ICANN also operates some of those accountability mechanisms I mentioned. So there's a function called the Customer Standing Committee who meets monthly to review all of our statistics and work out if we're doing a good job. And there's also less common review processes. Right now, there's something called the IANA Naming Function Review that's ongoing. It's in its concluding phases. But essentially that convenes, I think, every five years to do a more holistic view of how we do the naming portion of our functions and provide recommendations for change.

So, last slide. I know this is a whirlwind tour of what we do. What's our job? We're an essential records keeper for the internet. We record identifiers, the ones that are needed to keep the internet working, interoperating. We are part of ICANN org, but we're an affiliate organization called PTI. And our team members are here this week, and we look forward to talking with you. With that said, I'm very happy to answer any questions that you have right now.

SIRANUSH VARDANYAN Thank you, Kim. Very interesting. And indeed, in a very simple way, at least, we can understand. Questions. Now we open the floor for questions for both groups, for OCTO and IANA. Yes, I see our Next-Genner wants to ask a question. Please go ahead.

KHADIJA MIAN Hi, thank you so much for that presentation. My name is Khadija. And I wanted to ask, based on the kind of increasing data privacy regulations happening worldwide, has that affected IANA's

function at all, or do you see it affecting IANA's function in the future?

KIM DAVIES

Yes, that's a very good question. I mean, it has affected us a little bit. You know, when data privacy regulations were being promulgated significantly with GDPR, etc., five, seven years ago, we went through a fairly exhaustive exercise to work out our record retention practices and make sure they aligned with best practices in that regard.

That being said, I mean, our understanding is always evolving. And we do have a public registration function. So generally speaking, the kinds of personally identifiable information we typically publish is in relation to operating critical internet infrastructure. So we do publish personally identifiable information, for example, who runs a TLD. And the assessment is that there is a general public interest. You cannot run a TLD anonymously. We need to publish points of contact for operational reasons and so forth to be able to contact people.

Yeah, so, I mean, I think privacy regulations it's something we have to adhere to. It's the law. We work within the parameters of the law. We have had to make operational adjustments associated with that. It hasn't really fundamentally changed the way we do our job, though. You know, we still make assignments. And most of the assignments we make, to be clear, are not to individuals or parties. They are for purposes. Like, when we assign port 80 to mean web page transmissions, it's not that we're giving port 80 to

a person or a company. We're saying, for the world, port 80 is reserved for this purpose. So it doesn't really have a privacy impact, per se. Thanks.

SIRANUSH VARDANYAN Yes, thank you. Yes, let's go with you.

VINNY ADJIBI Good morning. Thank you for the presentation. My name is Vinny Adjibi. My question for IANA. I understand that you guys are assigning identifiers. Do you also have a function of retiring some identifiers?

KIM DAVIES Yes, I mean, we essentially manage the full lifecycle of identifiers. The creation, the maintenance, and then the removal, should it be deemed necessary. Every technology is different. You know, many assignments are never retired because for promoting interoperability, even if it's an old technology and it might not be commonly used today, we wouldn't necessarily retire that technology. It's still available.

You know, just to pick one off the top of my head, before the World Wide Web, before HTTP, there was something called Gopher, which was one way of doing information transfer on the Internet. Well, all the allocations for Gopher are still there today. If you want to pull up this 1980s technology and use it, you could. So it doesn't mean

that every kind of assignment needs to be removed, but certainly for those where it makes sense, we definitely do that as well.

SIRANUSH VARDANYAN The lady over there, and then we'll come back to you, and then Ashirwad.

BOLUWATIFE JIDE-OLUGBADE Hi, thank you very much for your presentation. My name is Bolu, and I have two questions, actually. The first is, what kind of challenges do you encounter managing these records? The second is, you talked about collecting customer surveys annually. So I want to know how you integrate feedback from the customers into what you do, and do you bring this feedback to the ICANN community before you implement them? Thank you.

KIM DAVIES So on the feedback, I mean, we do two kinds of surveys. One is immediately after a transaction. This is common. You would have seen it with other kinds of things. After we've performed a service for a customer, we send them a quick survey and say, how did that go? Do we do a good job or not? And if they say no, we say, well, what could we have done better? So we have this immediate feedback channel with individual customers to get that feedback and act upon it.

In terms of the more overarching strategic consultations we do, we do do an annual process, and we publish annual reports. You can

find them on our webpage. In fact, we just published our annual report just a few weeks ago for last year. And then in the course of our interactions with some of the community groups that are directly involved in our oversight, we talk about the results, and we talk about, well, what's this going to mean for the work that we do?

Thankfully, most of the surveys have been highly complementary, so we haven't had cause to look at the surveys and need to see it as a signal that we need to radically shift our approach. But it's tweaks. It's evolution. It's how can we take that and be even better than we've done before? And then apologies, I've forgotten your first question.

BOLUWATIFE JIDE-OLUGBADE I asked about, like, the challenges you encounter managing the records.

KIM DAVIES Right, right. You know what? I mean, I think the-- I could list several, but I think I'll just focus on one. The IANA functions date back to the 1970s, really, when the internet was experimental. And some of our records are just that old. So one big operational challenge we have is we don't necessarily know the circumstances of which some of those early allocations were made. This comes to the earlier question about how do you make changes to something if you don't know the context of how it was originally assigned? So some of the challenge we have is, like, going through the history books to try and work out exactly why things were done the way

they were, like, 30, 40 years ago, so that we can make good decisions today.

And then also a lot of what we do is like, do it once and forget about it. So we might make an allocation on behalf of someone and that someone, we can't contact them anymore because they didn't keep their contact details up to date with us. And even if we could contact them, like, were they doing it on behalf of their employer? Were they doing it in a personal capacity?

So yeah, there's just a lot of record-keeping challenges associated with records that are sometimes decades old to make sure that we can get to the right person that's authorized whenever we want to make a change. If we want to remove something down the road, like, how do we make sure it's properly authorized?

SIRANUSH VARDANYAN Rolla, please go ahead.

ROLLA HAMZA Thank you. I have two questions, actually. No, one comment and one question, sorry. It's a very interesting presentation on how to simplify the technical terms. It's very interesting. The second, my question is, PTI, IANA and ICANN, how they are coordinating, and what is the role of each one, and how they coordinate with each of them, and how they can implement the multi-stakeholder model? Thank you.

KIM DAVIES

So I mentioned in my presentation that ICANN was created to really be the home of IANA. Now, IANA has never been a legal entity. IANA is a concept. IANA is a function, if you will. It has a name, but there's no company called IANA. It's an over-encompassing name for a set of functions that are performed by managing the root zone, issuing IP addresses, things like that.

Now, I mentioned that ICANN was created to be the home of IANA, and that is true. In the 1990s, ICANN was created to be the place that IANA sits, and that's exactly what happened. So ICANN was created in 1998, and for the first 20 years, IANA was just a department within ICANN. Paychecks that we get said ICANN on the top and fully integrated into ICANN, but it operated under a contract with the U.S. government. At that time, the U.S. government asserted its authority to decide how the IANA functions are performed. So through a contractual way, the U.S. government told ICANN, you can do this, but you can't do that. You have to do it this way, not that way.

As ICANN grew and matured, there was a recognition that ICANN was sufficiently capable of making those decisions itself. The multi-stakeholder community had blossomed. You know, these conferences, this framework where everyone comes to participate works well.

So in the mid-2010s, the U.S. government signaled, we want to step away from this role. Like, the international community, you can decide the future of IANA. The structure of ICANN is fit for purpose. So let's work out a way for the U.S. government to step back and

remove its direct involvement. And the U.S. government will now be just like any other world government. Of course, it comes to ICANN conferences and participates, but just at the same level as any other government.

So there was this big process that happened. It was called the IANA Stewardship Transition. And there was a big convening of a lot of discussions over about two years. And ultimately, decisions were made that there needs to be some checks and balances in place so that should ICANN fail to do its job properly, that IANA could be separated out as a separate legal structure so that it could, for example, be rehomed under a new organization. And so that was the impetus to create PTI.

PTI is a new specific vehicle for performing the IANA functions. It's not intended to not operate within ICANN. I mean, even to this day, we're in the same offices as our ICANN peers. We operate very much today like we're still a department of ICANN. But legally, we're a separate organization, and we have different forms of accountability and oversight than the rest of the ICANN organization.

So in sum, the way it's structured now is what the community wanted. ICANN is the umbrella organization for everything, but the specific staff that do the IANA functions work in a different company called PTI, but that company is closely affiliated with ICANN. As for IANA, it's a trademark. It's a brand name. It's not a legal entity of itself. I don't know if that answers all the questions, but that's how I see it.

ASHIRWAD TRIPATHY

Thank you, Siranush. Ashirwad for the record. So I also have two questions. One is related to inside this. The geopolitics is getting quite complex. So how does PTI balance the geopolitics? Would they leverage the GAC, or is there some different mechanism where a few governments are making it in the risk of Internet being fragmented?

And the second one relates to the interplanetary network as we're trying to move Internet to space as well. So I don't know much about it, but I was seeing on Internet something about the multiverse model and the single model where there was another, let's say, registry kind of thing, and there was this another [inaudible – 01:03:34]. Could you explain a little bit about that as well?

KIM DAVIES

Certainly. I would say, firstly, when it comes to geopolitics, I think we're quite thankful it doesn't really directly impact IANA too much. We just try and do our job as best we can. I mean, our job is to coordinate things globally, and that's what we strive to do. And we try to navigate through any potential impediments, but I think the best insurance for keeping global cohesion and stopping the Internet from fragmenting is making sure all the structures that are in place today just work. And then there's not really an argument where we need something better because the thing that works today works well already.

So I think that's my mindset. I mean, I think it's worth noting without getting too much into the specifics of geopolitics, like, we service every country in the world. Like, every country, no matter the geopolitics, has a country code top-level domain that IANA has issued it. So we find ways of working with entities no matter where they're from. Like, our goal is to keep the Internet working anywhere, everywhere, no matter what the circumstances. So we really try to service everyone, and that's our goal, and I think we've been successful at doing that.

As for interplanetary communications, I think we're not a research function like OCTO, and we're not like-- IANA isn't a thought leader of where the Internet should go, we're a service organization for people that are really doing the inventing on the Internet. So our role in those kinds of things is, insofar as a lot of academics and work is going into those new technologies, when they mature and they're starting to be deployed, just like I mentioned before, technical standards will be devised and we will play a role in administering the records for those technical standards. But it's not something that we lead or proactive on, it's something that we will respond to as those technologies mature and we will put them into operation.

SIRANUSH VARDANYAN We have a question from our remote participant, Seth. Does the work done by the research team affect the policies made by the GNSO? So does the work done by research team affect the policies which GNSO develops?

MATT LARSON I guess I would say I would hope that, as the GNSO is making policy decisions, they find the work that we've done informative and useful to them.

SIRANUSH VARDANYAN Okay, thank you. Yes, please go ahead.

DIPSY DESAI Hi, my name is Dipsy Desai and I have a question for the OCTO team. I was wondering if you have any ongoing research or documentation about how to deal with name collisions in the global DNS and the alternate naming system?

MATT LARSON Ongoing research? Well, we did a lot of research on name collisions. ICANN had a large effort. It was actually an SSAC project that had community members and that OCTO participated in. The name collision, it was only six years, why can't I remember, the name collision something. Yes, thank you, analysis, the A is analysis. Name Collision Analysis Project, thank you. So that work has concluded and has been turned over to the folks within ICANN who are working on the next round project. That project didn't consider collisions between domain names and alternative name systems. We ruled that out of scope.

So I guess the answer to your question is yes, OCTO is looking at it. We're going to go back and update those two documents, or excuse

me, there's one document about alternative naming systems and that's going to be updated on an ongoing basis. So it's something we're aware of and we're following, but it's not intense work at this point, I guess, is a way to answer that.

SIRANUSH VARDANYAN Yes, Enerst, go ahead.

ENERST MAFUTA KATOKA Thank you, my name is Enerst Mafuta, for the record. I have a question for Adiel, KINDNS, yeah. So are there any kind of research around standardization around DNS in relation to MNOs on how you can optimize quite curious levels in terms of quality of service? Because according to the research that I've been doing right now, my interest in research in terms of infrastructure DNS, I've noted that many MNOs lack proper DNS infrastructure as a result there's poor quality of service. So I wanted to know what type of standardization research you're currently doing that can support that, like maybe load over, fill over, load balance, stuff like that?

ADIEL AKPLOGAN Yeah, I mean, MNOs most of the time interact with the DNS from the resolver perspective, right, the wrong resolver. And in terms of standardization, there is nothing specific done for MNO per se. But what we do in our team mostly is that we have a very specific program for mobile network operators on the DNS and engaging

them specifically on those aspects that are important for them and the service that they provide to their customer.

And I mentioned KINDNS in my presentation, and KINDNS is one example of things that we are doing because KINDNS is a framework that defines for a resolver operator, for instance, what are the key best practices that they can implement. And we have realized that many MNOs don't dedicate enough resources to their DNS infrastructure, as you mentioned.

So we do two things. One is to increase our capacity building and target it specifically to mobile network operators. But two, also helping them to use the framework like KINDNS to self-assess themselves and know where they have work to do in terms of running those services. We can talk more about that and introduce you to my team that is doing a very good work on the ground, talking to mobile network operators in one-on-one discussion to understand their challenge and to understand how they can implement those practices from those best practices framework.

SIRANUSH VARDANYAN Thank you. Yes, please, lady in red. I see from here your hand, but I cannot see the face.

DONNETTE SABRINA O'NEIL Okay, good morning. Please go ahead. My name is Donnette O'Neil, and I'm a new fellow. This question is for IANA. Before I proceed, the presentation was very clear. But my question is regarding to new identifiers, about how many applications do you

see for new identifiers per year? And in the last five years, was there an increase or decrease in the applications? And finally, due to the explosion of AI, do you see a lot more identifiers being created for the upcoming years?

KIM DAVIES

Thanks for the question. You're putting me on the spot to recall numbers off the top of my head that I don't have immediately handy. But I think, firstly, I'll say that the specific numbers are in all of our performance reports on our website. That said, I think the volume of identifiers, I would say, is typically increasing over time because there's more internet technologies over time. But the nature of the work varies.

So firstly, I'll say that a lot of what we do is highly optimizable allocations. And where we have very straightforward processes, we've invested a lot in automation so that our staff don't have to laboriously process individual requests. We have self-service portals for a lot of those kinds of identifiers we administer where they're very frequent, a lot of common operations.

So whilst the number of requests has gone up, our ability to operate on them has been okay because we've built tools and systems to support that workload. I would say that new technologies are always coming along from standards bodies. Many will not be market successes. Just because you create a new internet technology doesn't mean that everyone's going to use it. So some technologies have a small flurry of interest maybe over the first few years, but then fizzle out because that technology never

gained any traction or something else came along that is a better fit for the problem.

I think some of the trends we are seeing, firstly, when it comes to the way we administer domain names, obviously the explosion of top-level domains is a big impact. 10, 15 years ago, there was only 300 top-level domains. Now there's 1,500. ICANN's embarking on this big next round of new gTLDs. That means there's probably many more in the near future as well.

I'd say another trend is that there's more complexity to some of the things we have to administer. With new technologies, sometimes it requires more complicated identifier technologies to go with it. A lot of the very basic technologies we rely on were invented in the 1970s and 1980s when, frankly, computers were pretty limited in what they can do. A lot of the core technologies we rely on were designed for computers that didn't have a lot of computing power compared to what we have today. Today, with a lot more computing power, you can create much more elaborate technologies that have a lot more specific and complicated need for how identifiers are deployed.

As for AI, we've not seen anything yet where that's influencing one way or another how we do our work. Whether that's going to result in some kind of explosion of new technologies to support it, we haven't seen any evidence of that yet, but it's not something that I would say is expected or not expected. It's just something into the future. Thank you.

SIRANUSH VARDANYAN

My colleague, Marilia, shared in the chat in the Zoom space that if you want to learn more about IANA surveys, you can go www.iana.org/performance. And since 2023, fellows and NextGenners are also invited to participate in those surveys. So thanks for sharing, Marilia. Any last-minute questions? Yes, please go ahead, and then Saiid, you will go.

MABDA HAERUNNISA FAJRILLA SIDIQ

Thank you for your presentation. This is Mabda for the record. I have a question to IANA, to Kim. You mentioned about the stewardship transition, which, of course, it's been 10 years, so it's been quite a while. But my understanding is that, politics-wise, I think there were observations that the demand for transition was highly supported by democratic parties, for example. And you also have demands from politicians from the other party who supported having the U.S. still being the steward of the IANA function.

It's been 10 years, so I was just wondering if there has been any disruptions or similar demands 10 years into the PTI operating, and if there has been any disruptions that is similar to what happened back in 2014, how would you say has that affected the way the PTI operate? Thank you.

KIM DAVIES

I mean, I think the short answer is no. There's been no disruptions or impacts, like negative impacts. I think fundamentally what

happened in 2016 is we're more directly accountable to the community here. Like, we talk directly to our customers, our customers directly inform what we should prioritize, what we should focus on, and it's no longer mediated in the context through a contract with the U.S. government, where we had to see is this possible under the contract or are there other priorities. So I think it's been a net positive that we now have a direct relationship with the people that matter most, which is the people that benefit from our service.

Stepping back, everyone has opinions on how the internet should be governed and what the best model is, and I'm sure that debate will continue indefinitely. So even if people are speculating, well, maybe you do it this way instead, I think that's natural, that's common, and I don't think it's a negative. I think we should always be looking at are we doing it the best way we can, is there room to evolve, but are there like icebergs in the horizon that we're about to crash into to fundamentally change the dynamic?

I mean, I wouldn't think so, and I come back to the earlier question. I think the best insurance against those kind of outcomes is just keep doing our job well. As long as we do our job well, the community is getting what they need from our service, then I think we're in the best position to navigate into the future.

SIRANUSH VARDANYAN Saiid, go ahead.

SAIIDNAJIB SAIDISLOMZODA: It's here, and my question is for Kim. Is there any information how many persons of IPv6 are using all the world?

KIM DAVIES: Again, don't quote me on the numbers, but it's .0001% or something of that effect. There is so much IPv6 address space that it's not conceivable under the current allocation practices that we could run out. Now, we could run out if there's a decision to have a very poor allocation method, but the allocation method we use today, there's no prospect of it running out for centuries to come. It's a very large address space, which, as implicit in your question, is IPv4, which we managed. We ran out of in 2008 or somewhere thereabouts.

IPv4 was the numbering technology devised in 1980 for what felt like it would last forever, 4 billion possible addresses, but internet advanced such that we actually ran out of that. So IPv6 is the next numbering technology. It's what's called a 128-bit address space. It shouldn't run out as long as reasonable allocation methodologies are used, which I think they are being used. But in terms of IANA's pools that we have available to give to regional internet registries, we have essentially all still there. We've only just given a sliver of a percent out to regional internet registries so far.

SIRANUSH VARDANYAN: A follow-up?

SAIIDNAJIB SAIDISLOMZODA: Yes, I have one question. Are there any predictions for how many generations the IPv6 will have? I know the IPv6 dimension is over 60 million years, but all that starts have the end.

KIM DAVIES: Not that I know of, but I'm sure there's researchers out there that have projected out allocation rates and made estimations, but I'm not directly involved or aware of anything like that.

SIRANUSH VARDANYAN: So just for you to know as well, we tomorrow have another session, How Internet Really Works. So we'll touch base on this topic again tomorrow, and we will have also a session about tech career. But also, I would like to let you know that on Monday we have Tech Day sessions, so if you are interested to participate and learn more about technical aspects, so this is the way to go. Yes, Adiel, you want to add to this?

ADIEL AKPLOGAN: Just that tomorrow also there is a series of How It Works sessions. I don't know if it clashes with some of them.

SIRANUSH VARDANYAN: It does. It does clash with our sessions tomorrow. Unfortunately, we have other sessions planned for Sunday. But anyway, there are more opportunities, and also you know these people now, guys. You can reach them in the corridors. They are very busy people, I

know they have tough schedules here, but please reach them, and they are very happy to respond.

I would like on behalf of our group to thank our presenters here. Thank you for coming, and thank you for interesting presentations. Just for you also to know, all the PowerPoint presentations, which you will see on the screen today and tomorrow are already available in the meeting schedule, so whenever you go to a particular session, click on the details to find out the link to the Zoom room. You will find also the PowerPoint presentations there. You can upload them, download from there, and you can use for future references as well.

With that, thank you very much. And I invite all of you to enjoy some rest during the break, and see you back at 13:15pm here to talk about how ICANN policy works. So see you at 13:15pm. With that, this meeting is adjourned. Thank you very much.

[END OF TRANSCRIPTION]