ICANN78 | AGM – How it Works: IROS and IANA
Saturday, October 21, 2023 – 1:15 to 2:30 HAM

SIRANUSH VARDANYAN: Hello, everyone, and welcome to the session, which will talk about technical aspect of ICANN, so about IROS and IANA and what stands behind it. My name is Siranush Vardanyan, and I'm 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, as I have also posted the link in the chat. During this session, questions and comments submitted in the chat will only be read aloud if put in the proper form, as I have also noted in the chat. I will read questions and comments aloud during the time set by the chair of the session. Interpretation for this session includes six UN languages. Please take the headsets when you enter to this room. Click on the interpretation icon in the Zoom and select the language you will listen to during the session. If you wish to speak, please raise your hand in the Zoom room, and once the session facilitator calls upon your name, come kindly and mute your microphone and take the floor. Before speaking, ensure you have selected the language you will speak from the interpretation menu. Please state your name for the record and the language you will speak if speaking a language other than English. When speaking, be sure to mute all other devices and notifications. Please speak clearly and at a reasonable pace to allow for accurate interpretation. This session includes automated real-time transcription. Please note that this transcript is not official or authoritative. To view the real-time transcription, click on the closed

caption button in the Zoom toolbar. To ensure transparency of participation in ICANN's multi-stakeholder model, we ask that you sign into Zoom session with your full name. For example, a first name and the last name or surname. You may be removed from the session if you do not sign in using your full name. With that, I would like to introduce you to a great panel we have today. This is a very interesting session, and I'm sure you will get a lot of learning opportunity out of this session. We have a great team here today with me. And to start with, I give the floor to Kim Davies, who is the head of IANA. And he will tell us what is IANA and what it stands for. Kim, the floor is yours.

KIM DAVIES: Thank you, Siranush, and hello, everyone. We just switched all the speaking order because John was not meant to be here yet. So we're going to leave it switched for the moment, and we'll do John's bit a bit later on. My name is Kim Davies. I head up the Internet Assigned Numbers Authority. You might not know what that is. That is fine. That's what we're here to illuminate you on, amongst many other things today.

The session today really deals with the technical operational functions of ICANN. We have a code name for that internally, which is IROS. IROS is the Identifier Research Operations and Security Function of ICANN. And this session is about who we are and what we do.

So IROS is divided into two departments. One is known as the Office of the Chief Technology Officer. John Crain is our Chief Technology Officer. And OCTO has a function that he will get into in a moment with respect to research, analysis, writing papers, training, engagement, and

so forth. I, on the other hand, work with the IANA team to deliver authoritative registration functions for protocol parameters. And I'll explain what that is in just a moment. But first, I'm going to ask to switch to the alternative slide deck, please.

What is the Internet Assigned Numbers Authority? In a nutshell, we're responsible for the global coordination of unique identifier systems. Unique identifiers on the internet can mean a lot of things. You're probably all very familiar with domain names. That's probably the most obvious form of unique identifier that we use every day. But there's many others as well. IP addresses is another form of unique identifier that is in common use. And there's many others beyond that as well.

The Internet Assigned Numbers Authority, that name was given to it in the 1980s. But as a function, it's actually existed since the very earliest days of the internet, when the very first researchers were creating what was to become the internet that we know today. There was a need for someone to be basically the official record keeper to work out who's been assigned what numbers, even when the network was very small, making sure that which computers were connected to the internet were accurately recorded so the research activity could continue successfully.

The person that originally did that was a man by the name of John Postel. You can see him in the photo on this slide in the white shirt. John Postel's sitting there with Vint Cerf, one of the inventors of the internet. And what was originally a one-man job by John Postel has grown over time. Originally, it was him. He then had a team working with him. And

then later, ICANN was created. And one of ICANN's core missions is performing the IANA functions.

So as I just mentioned, ICANN was created to be a home for the IANA functions. We would have all seen the posters around the venue. It's the 25th anniversary of ICANN. So ICANN's been effectively responsible for IANA now for 25 years. And one of the reasons ICANN was created was to be that home. It was one of the primary purposes of creating ICANN as an organization, was to take the IANA functions, which were then performed in an academic setting, and put them inside a purpose-built organization that was capable of running them.

IANA today is a department that sits within ICANN to maintain these definitive records, or registries as we call them, of unique identifiers. Each of the unique identifier types that is maintained by IANA is based on community-developed policies and procedures. The way that they're developed varies depending on the type of unique identifier. Many of them are set in technical standards. The ICANN community here this week will be developing policies, many of which will have an implementation that is performed by IANA. So IANA implements these policies according to those standards and policies that the community has created. When you hear IANA, it often refers to different things. It can refer to the set of functions, like the registries that are maintained. It can refer to the department of staff that conduct those functions.

All right, so why do the IANA functions exist? And I know this is a bit abstract. I'll get into some specific examples in a moment. If there wasn't some kind of central coordination function, the internet as we know it today simply wouldn't work. The internet is a global network.

Most people say that it's free from any centralized coordination. Anyone can do anything and it just works. Well, it does work because there is a baseline level of coordination globally to make sure that when devices connected to the internet talk to one another, they're basically speaking the same language. This is the role of technical standards on the internet. And it's the role of the IANA to make sure that the parts of the internet that need to work the same around the world do so. Without this level of global coordination, it simply wouldn't function as an interoperable network.

Who uses these registries? For the most part, an end user on the internet doesn't need to know about the IANA function, never needs to visit our website. Really, it's just not a factor in day-to-day usage. The people that need to use it are people writing software, internet software. They use the IANA functions, network operators, things like that. Anyone involved in core or critical infrastructure or building applications on the internet typically are the direct users of the IANA functions. But as an end user, you wouldn't necessarily use our services directly, but you'd be a beneficiary of our services. So one of our key functions I'll get to in a moment is managing the DNS root zone. DNS root zone is something that your device is relying upon every day for it to function. And that's one of the functions of our team.

So today, the IANA functions is performed under a number of different contracts that govern how it's performed. Technically, we're operated under an organization called Public Technical Identifiers, or PTI. The specifics of that are not particularly interesting, but we are sort of an affiliate of ICANN. We're not the same organization, but it is an affiliate of ICANN and controlled by ICANN. So whilst it is a separate legal entity,

it does operate in a day-to-day capacity as part of ICANN. The slide is a bit out of date. It says 16 people. I think we're up to 20, maybe even 21 right now. So we're growing a little bit. But generally speaking, the team's been around this size for quite a few years. Beyond the IANA team of 20, we do rely on all the functions of ICANN more broadly. So all the different departments of ICANN contribute to the success of the IANA functions.

Let me talk about the IANA functions in more specificity. Generally speaking, we tend to divide the work we do up into three categories. These categories are somewhat arbitrary, but it makes sense in terms of how we're overseeing and how we structure our work. Firstly is protocol parameters. Let me talk about that. So protocol parameters I would summarize as a big, long list of identifiers that typically you would not see as an end user, but software implementers are very interested in. These are a lot of the internal encodings that happen device to device in a way that's not necessarily visible to an end user. Some are, but for the most part, a lot of the different unique identifiers that we manage here are used purely for internal network transmissions from device to device. Here's a small sampling on the screen here. Maybe some of these you've heard of. Maybe you haven't. There's a complete list of identifier types on our website. There's about 3,000 registries, and each of those registries has a different identifier type. And each of those registries could have anywhere from 10 to-- I think our biggest registry has something like 100,000 registrations in it. So small to big.

Here, most identifier allocations are made directly by IANA. So if you're a software implementer that might need one of these values, you would

come to IANA and ask to be issued with a value in that registry. There's no politics or middlemen, I like to say. It's a direct relationship between implementers and IANA.

Now, when we talk about these protocol parameters, in fact, everything IANA does is a protocol parameter. That pie chart I showed you in the beginning divides it into three different areas. We call out number resources and domain names differently because they're specialized cases of protocol parameters or protocol assignments. They're hierarchically allocated. You don't come directly to IANA. Instead, IANA gives out chunks of namespace or address space to other entities, who in turn allocate it further along the chain. And also, they're very policy-laden and/or political. The vast majority of the work IANA does is not political. It is operational, again, used by software vendors. Domain names and number resources probably don't fit into that categorization. There we go.

So where do these registries come from in the first place? Well, the preeminent place where internet standards are developed is an organization called the Internet Engineering Task Force, or IETF. Internet standards are generally developed in that organization and are published in the form of what's called a request for comment, or an RFC. RFCs serve as the official standards documentation for the internet. We work within that process. When those RFCs are in draft form, they're known as internet drafts. And we're involved in the process of those being developed prior to publication. People that are authoring RFCs will often work with the IANA team to understand how to create registries, how to update registries. And we provide advice on how to

implement internet standards in a way that IANA can operationalize it successfully into the future.

All right. So now, let me talk about number resources specifically. So as I mentioned a moment ago, number resources are just a specific form of protocol parameter. In fact, it really relates to just three types of identifiers. One is IP version 4 addresses. Another is IP version 6 addresses. And the third type is autonomous system numbers or AS numbers. IPv4 and IPv6 addresses do the same function, which is that every device that is connected to the internet needs to have basically a unique number. And so there needs to be some global coordination to make sure two devices aren't issued with the same number. Version 4 was the original address space that's been used on the internet since the 1980s. Unfortunately, IPv4 didn't have enough numbers in it to really issue to every device, given how big the internet is today. So IPv6 was created in the late 1990s to address that problem. IPv6 has a lot more address space for the same purpose.

AS numbers, on the other hand, are network identifiers. Each network provider only needs to be issued with one. It's not something that's issued to each device, but to a network. And these make things like routing efficiently possible.

So I think I just gave this slide without having it on the screen. I think just the point I didn't mention is when we talk about those network identifiers, the AS numbers, I like to conceptualize those as a bit like a postal code. So different regions, suburbs, cities, et cetera, in most countries have postcodes associated with them, which makes sending mail easier. So you use a postcode so that it can go to your local post

office. And then that local post office knows how to get it exactly to the final destination address. AS numbers work the same way. Your computer doesn't need to know, or your network provider doesn't need to know where every single IP address in the world is. It aggregates up to the level of an autonomous system number. And as long as your ISP or your network provider knows how to get traffic from its network to the network of the AS, then that destination network can work out how to get it to its final destination.

I mentioned before that it's hierarchically delegated. So what I mean by that is that you cannot come to IANA and get an IP address. You need to go through another organization. In the case of number allocation, there are five organizations called regional internet registries. We make large allocations to those regional internet registries of these numbers. And they, in turn, make further allocations. You can't actually go to an RIR directly to get a single IP address either. Network providers go to the RIRs and get large allocations for their networks. And in turn, network providers issue the individual IP addresses to their customers. So it's kind of three levels of assignment that happen for IP addresses through this model. IANA does do some direct allocation, but only for very special purposes, not for general allocation purposes, things like private use address space, multicast address space, and so on. All right.

Domain names is probably what most of you in the room are, and online, are most familiar with. Our role here is, again, to administer the high-level identifier space, in this case, the domain name space. So this diagram conceptualizes the domain name space in a tree. We have, at the highest level, what's called the root zone. The root zone contains information about what exists the next level down, which is what we

call top-level domains. Here on the diagram, we have .org as one example, .us as another example, and .bel which is the Cyrillic domain name for Belarus.

In turn, those top-level domains allocate the next level down. Wikipedia.org, ICANN.Org, etc. Hierarchically assigned. IANA is only responsible for the highest level. You can't come to IANA directly to register a .com domain for example. You need to go to the right level down the domain name tree.

Our primary responsibility when it comes to the domain name space is the administration of that DNS root zone. It is the authoritative record of what is a top-level domain, which ones exist, which ones don't. And then for those that exist, we maintain all the technical data needed for them to operate, and also operational data, like who operates them, what are the points of contact, things like that. For those who know what a WHOIS server is we maintain the official list of WHOIS servers for TLDs. A lot of that operational data, we maintain.

We work with top-level domain operators to do this function. So the managers or .us, .com, .de for example, they work directly with IANA to administer their TLDs. We review those requests to make sure they comply with the relevant policies, and then we implement based on that. If it's a gTLD, it's largely defined by the contractual relationship the gTLD has with ICANN. for ccTLDs, it's very different. CcTLDs are country code top-level domains. Each country has one or more ccTLD. We perform due diligence on those, make sure they're being operated in a way that's compliant with whatever the policy is within the country.

Once we've received those requests and processed them, we send updates for propagation in the root servers.

Another important function that we have relates to cryptography in the DNS. This is a way of securing DNS data so that it can't be tampered with. And this is one of the security mechanisms that was added back in 2010 to the DNS that we play a critical role for. Much like the root zone is sort of the most critical part of the DNS because it's the top of that tree that you need to follow down, the root key signing key plays the same role on a security level for how DNS is validated cryptographically. Because the key that protects the highest level of DNS is very important to the way that the technology works, it is critically important that we keep the cryptographic key for that particular role very secure. The way that it is operated, it's very difficult to change the key. So it's essential that we keep it secure and keep it trusted. And we need to make sure that the operational community has confidence that it is operated in a proper way and that it hasn't been used in an unauthorized fashion or disclosed and so forth.

And the way we do it, this is quite different from the way security and cryptography is usually done. Usually security is about keeping things secret, keeping things hidden. We have a sort of a radically different approach, which is we try to make the process of managing our key as transparent as possible so that everyone can look into what we do and satisfy themselves that we are managing the key in an appropriate way. One of the key ways we do this is through a process of holding key signing ceremonies. We hold these typically four times a year. And we invite people to come in and participate in these ceremonies. Also to view them online. And everyone can look at exactly how we're using our

equipment, conducting these events to satisfy themselves that we're doing them in an appropriate way and they believe that the key is secure. We also have a very deliberate process that is recorded. There are witnesses in the room. Everything we use to conduct these events is open so you can download the source code, the software, any aspects of what we do is all available.

So the purpose of doing this in this very transparent way is much like ICANN generally and the multi-stakeholder model that you'll witness here this week, is to engender trust in the process. Rather than us saying it's secure, trust us, it's super secret, we keep it in this fortified bunker that just trust us, it's really secure, you can see for yourself. You can follow along. Certainly, most critically, security experts can follow along. And security experts can then vouch for, yes, we've seen what ICANN has done. We're satisfied that the way ICANN is operating this key is trustworthy. We've seen it with our own eyes. We've reviewed all the procedures. We've looked at the source code, etc., and can confirm that the way we do it is secure.

All right. So that wraps up. And this is a very high-level presentation. I was trying to hit the highlights, but sort of that wraps up a summary of what the IANA functions are comprised of. I just want to touch on a couple more slides. I think I'm almost done. Security at IANA is more than just what I described, the key ceremonies that we use for DNSSEC. We have all sorts of other security elements going on, including we have third-party security audits. We have separation of duties in our systems, limited IANA roles, etc.

Other sort of cultural aspects of IANA that I wanted to share is that we provide a function to the community. And above and beyond the security dimension of the key ceremonies I just mentioned, it's important that the community at large has confidence in the services we deliver generally so we can do our job effectively. So we want it to be known that we do our job impartially and neutrally. And we have a lot of checks and balances in place to make sure that that is true. We have a lot of performance metrics that govern how we do our work. We report on those performance metrics to the community. There's a lot of reporting that happens about how the IANA functions are done. As an example, just tomorrow there is a meeting of the Customer Standing Committee. The Customer Standing Committee is one of the oversight bodies of IANA. They validate that we're doing our job in a way that the community finds acceptable. We have a culture of continuous improvement, always looking at what we're doing, always identifying what we can do better next time around. I mentioned before we have audits, third-party auditors that come in, validate our systems and processes, and then various other forms of accountability as well.

All right. So to summarize, our job is essentially to maintain those definitive global records of unique identifiers that are used on the Internet. Why? So the Internet just works. So that when you use a device on the Internet, it speaks the same technical languages as every other device on the network. Most of the IANA registries we maintain are very straightforward. You wouldn't know about them. You wouldn't have a need to know about them. It's generally software implementers that would typically need to know the role of the IANA functions. When it comes to those high-profile, hierarchically delegated registries, domain

names, IP addresses, we maintain sort of the global root for those, but we don't make direct allocations to everyone. Instead, there is this hierarchical assignment model that exists for those.

So that completes my tour of the IANA functions. I did want to finish on this slide. This is something sort of new but sort of not. We always do an annual survey of our customers and of the community to make sure that we're focusing on the right things. We're just starting our survey for this year, this time. What we're doing differently this time is using this QR code in every time we're presenting this week. And we also have some little business cards with them printed out as well. If you use the IANA functions or want to provide feedback on the IANA functions, we invite you to fill out just a couple of minutes survey so we can get that feedback. I think one of the questions relates to things we can present about that is of interest to you. So that kind of feedback is useful so we can target or have more specific presentations on areas of particular interest. So with that, I don't know if now's a good time to pause and take IANA questions before I hand it over to OCTO. So if you have any questions about IANA or anything you just heard, please.

SIRANUSH VARDANYAN: Thank you, Kim. There is a question. I will take the question from our virtual fellow who is asking what specific opportunities, if any, does IANA currently offer to ICANN fellows to gain expertise in this field? Is there any specific opportunities which you can mention?

KIM DAVIES: It's a good question. I think directly we don't have much because IANA is a specific operational function of ICANN. ICANN as a whole provides a number of different engagement and outreach opportunities. The fellowship program is a good example of that. So we're beneficiaries of the broader activities that ICANN does. The only thing we do directly that might be specifically of interest to those particularly with a security interest is these key signing ceremonies. We do have room for people to attend them in person. You can also attend them online. So those that have a particular interest in those and might want to attend, there's an ability to submit a statement of interest that you might want to attend one in person on our website. And so that's one way to get directly involved. But if someone has a specific interest, I'm happy to talk to them about it and find out if there's a way we can help facilitate that.

SIRANUSH VARDANYAN: Thank you. And we'll take a couple of questions now. Houda, you are the first.

HOUDA CHIHI: Hello, everyone. Thank you so much for this great presentation. My name is Houda Chihi. I am Tunisia Telecom staff operator. So the question is, due to there is always a lot of threats, is there any different version of the DNSSEC protocol? Thank you

KIM DAVIES: So there's no different version of the DNSSEC protocol, but the DNSSEC protocol itself is designed to be adaptable in terms of things like the security algorithm that is used. So cryptographically, there's different

forms of algorithms that are used to protect communications. The algorithm that's used in the root zone is known as RSA SHA-256, which is a mouthful. It works well today, but one of the activities we have going on right now is we recognize that whilst it is a good protocol today, it might not be in 5, 10, 20 years from now. So we need to have the ability to switch to a different algorithm in the future. To that end, there is actually an active work that has just concluded very recently called the Algorithm Rollover Study Group. And they just published a draft report just a couple of days ago on what would need to happen for us to change that cryptographic algorithm. So that is now available for public comment on ICANN's website. We're hoping that security researchers, security experts will look at that and confirm the findings are correct or provide additional findings. And the final report there should give us a good roadmap on how to evolve the security algorithm into the future. So that's one area where I think DNSSEC, we have some work to do and we have opportunity there. But in terms of broadly DNSSEC in a general sense, no, there's not much more than that.

SIRANUSH VARDANYAN: If I can request to ask your questions in a reasonable pace for accurate interpretation, please. Shita, and then we'll come to you, gentlemen, on the second floor.

SHITA LAKSMI: Thank you very much, Siranush. My name is Sita Laksmi. I'm a first-time fellow from Indonesia. This is a basic question, but because I'm a first-time fellow, it's justifiable to ask these basic questions. The first one is why is PTI or IANA is an affiliation from ICANN? Why not together? When

I see the diagram from the courses, you have your own board. And why not being part of the ICANN? The second thing is also very basic. Should we call PTI or should we call IANA? The third thing is all the IANA functions, is there any IANA functions that's not being done by PTI or ICANN and being done by IANA itself? So that's a bit of clarification of the terms and confusion. Thank you very much.

SIRANUSH VARDANYAN: For the first-time fellow, I mean, it's perfect.

KIM DAVIES: These are very good questions, so please don't apologize for asking them. So why is there PTI? Why does it exist? Why is it sort of separate from ICANN but not really? You'll notice outside when you come in to get your badge on the first day, there's a big poster that says seven years since the IANA stewardship transition. The IANA stewardship transition was a process where up until 2016, the IANA functions were actually done under contract with the US government. The US government has always had a role in overseeing the internet's development since 1969 when the first experiments happened, all the way up until 2016. Leading up to 2016, there was a process called the IANA stewardship transition that recognized that the US government's role should be superseded by the global multi-stakeholder model that's here this week to oversee the IANA functions. And so there was a process that went on for about two years where all the stakeholders came together and they devised an approach on how IANA should be overseen in an effective way by the multi-stakeholder community. Up until 2016, as you just mentioned, IANA was directly part of ICANN. It wasn't even separable in any way.

But the multi-stakeholder community felt that there should be some safeguards that would theoretically allow the IANA functions to be separated away from ICANN. For example, if ICANN stopped being a suitable home for them. And the way those recommendations were implemented was to actually transfer the IANA functions into a separate organization. So today, there is a separate organization called PTI. Me and my team, we're employed by PTI, not by ICANN. So it is legally distinct. But in a practical way, like we all work in the same office, we all work together with ICANN staff. It's not a big distinction, but it is important legal distinction. So that should the community decide in the future that the IANA functions should be separated from ICANN, all the operations are legally distinct. So that's really the main reason PTI exists. Beyond that, you don't really need to use the term PTI for common usage. As you know, we typically call ourselves IANA. We don't call ourselves PTI. The community knows us as IANA. They call us IANA. So we use PTI in very specific contexts where the legal distinction is important. But in our day-to-day operation, we deliver a service called IANA. We're called the IANA staff, the IANA team, et cetera. And that's what we're typically called. So I don't know if that answers all the parts of your question, but that's basically what it means, PTI and IANA.

SIRANUSH VARDANYAN: Thank you. And before we take our last question for Kim Davies, I would like to ask to switch the presentations for the next one. Please go ahead. You. Yes.

UNIDENTIFIED SPEAKER: Okay. Merci, Siranush. So I come from Benin. I [head] the Beninese chapter of ISOC, and I have two questions. First, I would like to go back to what the person who just spoke said. I think in terms of the hierarchy in decision-making, I'm not sure about the difference between OCTO, ICANN, and IANA. So if the IANA team decided to do something, let's say "thing" generically, and if OCTO said, "No, this cannot be done," then what will be done? Who is the one deciding? Who has the final word? So how are decisions made? That is my first question. And then secondly, in your presentation, you were very clear about explaining that IANA is the one that distributes the IP blocks to RIRs. So I would like to know how IANA regards the activities of RIRs, meaning that you delegate those blocks, and then each RIR can do whatever they want with IP addresses? How is that done? How do you delegate those blocks? Do they each do whatever they want with them, or are there certain limitations to that? Thank you.

KIM DAVIES: Thank you for those questions. So the first question in relation to avoiding conflicts between OCTO and IANA, I guess I would summarize it as, generally speaking, yeah, we just get into a room and ... No. IANA and OCTO serve two very different purposes. I don't think there's any risk of overlap, and I think when you hear about what OCTO does in a moment, it will become clearer. Because IANA does a specific operational function of making assignments, and it's not what we think we should do. It's that we follow policies, and the policies generally make it very clear about IANA shall do this when this happens. So objectively, the policies should be straightforward, and it's pretty clear when IANA has a role or when it doesn't have a role. When we need to

sign a country code top-level domain, for example, to a country, like what is a country that is allowed to have a country code, what the country code should be, who's allowed to have it, etc., these are all governed by policies that have been developed by the community, and we as staff are there to implement them operationally. But I am not a king. I don't just decide I'm going to recognize this country. I'm following a policy and a procedure. Our team is following these policies and procedures to make these operational decisions, which leads to the second question, which is how are IP addresses allocated to the regional Internet registries? It's a very similar story. There are global policies that define exactly when we make assignments to regional Internet registries and in turn what the responsibilities of the regional Internet registries are. So within each region, each RIR has its own policy setting bodies that set policies. So in this region there's RIPE, which sets the IP address allocation policy for Europe. In Africa there's AFRINIC that sets the IP address allocation for the Africa region. So we work in a global system. I think generally speaking I would say that one of the main reasons that this ICANN forum exists throughout this week is to continue to evolve those global policies and where those global policies require some changes to the IANA functions, it will be developed in the groups that are meeting here this week. And so if there was a need to evolve how IP addresses are allocated, that would happen in what's called the Address Supporting Organization or ASO. If there was a change to how ccTLDs are allocated, that will happen this week in the ccNSO, the Country Code Named Supporting Organization. So it's through these SOs and advisory committees that policies are made, recommendations are made that we as staff then operationalize.

SIRANUSH VARDANYAN: Thank you, Kim. Thank you very much. And I would like to give the floor now to John Crain to speak about another part of the conflict. What is OCTO, IROS? So, so many abbreviations. Please do explain to us what does it mean.

JOHN CRAIN: Thank you, Siranush. As I said this morning, we are the masters of acronyms in my group. You know, I'm the SVP and CTO of OCTO and IANA, which is together is IROS. And I'm pretty sure there's a few other acronyms thrown in there. So I'm going to talk a little bit about the other side of the group that I'm responsible as the senior vice president. And actually, one thing that is very similar for all of ICANN relating to that previous question, as staff, we don't set policies on any of these areas. Right. Those are all set by the community. We are, if you like, a secretariat or an operational function where we implement policy. So that's on both sides. Now, the office of the CTO is a little bit of a strange one inside ICANN because we're sort of a think tank and an outreach education kind of organization within ICANN. We're very much outward facing towards the community, but also inward facing towards the board and the rest of the organization when it comes to giving advice.

We separate the office of the CTO into four distinct groups, all of whom work with each other. So although they're distinct groups, it's kind of all working in collaboration. We have an operations group. So we have projects and work that we have to do. So we need project managers and we have to do a budget and all that kind of fun stuff that any business unit does. So you can think of that sort of as our management of our

business units where all the fun stuff happens. We have a technical engagement group, which is exactly what it is, is out there engaging about technical things. And I'm going to pass off to Adiel in a minute to talk about that. And then we have a couple of research areas, one that is mostly academics that are focusing a lot on issues around security and abuse of identifiers. A lot of that work is in the namespace. So if you see discussions around DNS abuse, which you will do this week, I guarantee it, you may see references to some of our work and some of the studies we've done or some of the tools that we have to measure these kind of things. And then the other side of the research group are mainly well-known engineers and protocol geeks who, when you talk about the IETF, these are guys that are and girls that are in that world and are either understanding or even sometimes developing protocols and then doing research around how those protocols are used and what it means and forward looking, etc. And what I'm going to do rather than listen to me all day is I'm actually going to pass off to Adiel, who runs our technical engagement group. So, Adiel, why don't you introduce yourself and tell me when to go to the next slide and I can be your secretary here and do that for you.

ADIEL AKPLOGAN: Yeah, thank you. OK, so I'm Adiel Akplogan, vice president for technical engagement and work within the OCTO team. Next slide, please, John. So the technical engagement function, as the name mention it, has its main responsibility to take everything that OCTO does and bring it to the community. And we do that through four key responsibilities. One is to develop and that is for the whole organization wide our relationship with other organizations that are involved in the

coordination of the technical identifier system in general, work with them and maintain with them a viable and healthy relationship. The second element of our responsibility is to build a framework that allow us to also cooperate with other organizations that are not directly involved in the stability of the identifier system. But that has technical role, technical function that we are interested in. So how can we work with them? How can we bring them the evolution in the unique identifier system? How can they work with ICANN overall on technical matters?

The third aspect of what we do is developing and leading initiatives that can help us in our engagement activities. One of them is KINDNS. You may have heard about it. It's an initiative that we have launched recently to help operators implement DNS best practices. We have a few other initiatives, some are internal, some are external as well. We are very much engaged in Africa region, for instance, with the digital coalition for Internet development in Africa where Yazid is here. Yazid is pretty much involved in that initiative as well. So anything that can help operators or any other constituency that are technical oriented are managed and lead through the technical engagement function.

And finally, another area that is very much visible from the technical engagement is capacity building, capacity development and outreach to various stakeholders. Here we interact with DNS operators, we interact with ccTLD operators, we interact with law enforcement and public safety organization, all of those constituencies to help them understand better the Internet identifier system, the DNS security around that and how they can work with ICANN in general.

So the two tracks for that is the capacity building, as I mentioned. This is regional. And the engagement, which beyond capacity is about outreach and explaining what we do. And we do that more on the regional basis. So we have a team in different regions because each region has their own needs, has their own specificity. So we want to make sure that the way we engage is aligned with what each region has. So I also have David Huberman here. David Huberman is managing the North America region, for instance, while Yazid is more focused in Africa. I have Nicolas Antoniello focusing on Latin America. I have [inaudible] in Asia Pacific and Peter [inaudible] for Europe.

So that gives us a very good approach in the way we do our engagement and the way we build collaboration and cooperation with different organizations as well. Another thing that we are working on as well, more actively nowadays, is to help technical constituencies within ICANN to bring their outreach to the community as well. So I know that RSSAC, SSAC, they issue documents or advise regularly. And some of those advice have impacts on operators. So my team helps them outreach and talk to the community and explain what they are doing globally. So that is it globally for the technical engagement function. So it is the public facing part of OCTO. The three of us are going to be here for the rest of the week. If you have questions, you need any support in terms of engagement, capacity building, feel free to reach out to any of us. So back to you, John.

JOHN CRAIN: So I am not Dr. Samaneh Tajalizadehkhoob, but she does work in our group and she's the director of SSR research. SSR meaning security,

stability, resilience. And she's not here at the moment, although I do believe she's coming later in the week.

So the goals are really to support both internal and external functions around things that affect the security, stability or resiliency of the ecosystem. A lot of that is focused on the DNS, but we also look at things like IP addresses and routing, etc. So very much an academic group. They're looking at what's the state of the art. You know, what are the latest things happening out there in the ecosystem? Writing research papers on that, educating and doing outreach to the community about what they find, creating tools. We have visualization tools for certain data sets. We have a lot of data. So we do build visualization tools to help the community understand what we're seeing and facilitating discussions. So there's a lot of discussions around. You'll hear the term all the time, cyber security, which is one of the buzzwords that we've had for quite a while now. We focus in this group generally, and I used to lead this group before I took over the CTO position, which is why I say we, really focus around how the identifiers are used or misused in security situations, either as threats or how they can be used to actually improve the security. And basically, like most of OCTO, they are subject matter experts to the community and to the organization and the board.

So here are some of the key projects that are ongoing at the moment. I'm not going to go through them all. One that you will hear about is DAAR, Domain Abuse Activity Reporting. This is actually taking reputation data. So reputation block lists, looking at what names are resolving in the DNS, see if they are on these block lists, and then building statistics around the reputation of particular parts of the

namespace. So you can do that, obviously, by a top level domain, but you can also do it by a registrar. And of course, you have data, so data you can cut in many different ways. So we can look at that data from many different perspectives.

An interesting follow-up to that was something called DNS TICR. Did I mention we like acronyms? Now, I'm pretty sure that they start with the acronym and then make the name. So Domain Name Security Threat Information Collection and Reporting. So DNS TICR, we call it. This was interesting because this evolved around the start of a pandemic. There was something called COVID that you may have heard of. And abuse on the Internet and the misuse of identifiers like names normally occurs around some kind of significant event. So if there's a significant event, bad guys will use that for a lure to pick on society, right? They're after money or they're after data. So DNS TICR, what we did is we took similar data sets. So what names are resolving, zone files. What are people saying about those names? Reputation feeds. But also, we did what we call string recognition. So we created a database of thousands of strings or labels, or you might just call them words, that we believe could be associated with the pandemic or the cures for the pandemic in multiple languages, multiple character sets. So literally thousands. And when we found registrations of names with those strings in them, we'd look to see if they were causing harm. And then we would gather evidence of that. So it could be a screenshot. It could be a malware sample. And then we would pass that on to the registrar. So, and this is an important thing, if you're following the DNS abuse discussions, we gather evidence, and then we actually provide that evidence to somebody who can take action. And what we found is when you actually go out into the

industry with evidenced arguments about abuse, they take action. When you go out in the industry without any evidence, they don't. Not shocking. Anybody who works in security knows that you get a lot of complaints that are either misled or false. So evidenced reports of what we call DNS abuse. We did this mainly around phish or phishing and malware droppers. So people who are actually dropping malware. And we can do that pretty much around any term. But we did this particular experiment because we're a research department around COVID and the pandemic.

We have a bunch of papers. We write a lot of papers on both research sites. An interesting one, at least I think it's interesting, is we developed a classifier. We're looking at unsolicited emails, spam you would call this, that has been reported as being abusive. If you work in the security industry, you'll know there's a lot of places where people report abuse. One of the feeds that we had access to was an email feed. So people would send the email to this particular feed and say, this is a phish. And we were looking at statistics based on this feed, and we didn't quite believe the data we were seeing. And we realized that nobody had built a classifier to look at an email and decide what kind of abuse it was. So using machine learning and lots of CPUs and lots of data, we built a classifier that can look at hundreds of thousands, millions of emails and classify what kind of abuse is actually occurring there. And it turned out, indeed, it was not all phish. There was all kinds of different things in there, and some of them weren't even classified as what we would classify as abuse. So that was an interesting one, because we found something that we thought was a problem and a question, and then we tried to answer it. And there will be a paper coming out on that soon.

We actually did present it at TrustCom. If you're in the security world, you'll see there's a lot of academic conferences, and you can go and present your academic papers. So we presented an academic paper. And hopefully later this year or early next year, we will come out with what we call an OCTO document, and we'll probably come back to those, which is our own papers. And so we will be publishing that. Looking at those reputation block lists and deciding what is a suitable block list for use for your purpose of measurements. So another thing that we've done some work on. Once again, presented it to a conference, and will soon be coming out as an OCTO paper. And we do a lot of that. We're also looking at blockchain-based names in the root. Are they appearing in the DNS? You'll hear a lot about blockchains and names, but one of the questions was, do we actually see them in the public DNS? So there's been some research on that. And many, many other research projects. So these guys are very busy. So I'm going to hand over to who is not Matt Larson, just like I am not Samaneh Tajalizadehkhoob. And I'm going to hand over to my colleague, Alain, and I'm actually going to give you the pointer.

ALAIN DURAND: Thank you. So I'm not Matt. I'm Alain Durand. I'm a technologist, distinguished technologist at ICANN. And technologist means, well, I look at technology and scratch my head. That's really what it is. So we have a research group, and we can think of us as a think tank, as John was mentioning earlier. And we look at many different technologies, but are not really focusing on DNS abuse, essentially everything else. The idea behind the group is to provide trusted, verifiable information

to the community. So when policy decisions are being made in this forum, it can be based on actual facts that are verifiable facts.

So for this, we engage with the community. We look at, as I said, different technologies that are brought to our attention, try to analyze them and write papers about them. We also have access to data directly from the root servers that are managed by ICANN, the IMRS, formerly known as L root. And we're really lucky to have that because it's the trove of information we can get in there. And we do some research on this, trying to understand how the DNS actually works.

And one of the things that really made our head scratching is the DNS doesn't work the way we teach it in textbooks. It's different. Overall, again, that's the same. But when you look at the details, it's very surprising. And we're trying to understand this and we are making progress. So also facilitate the NCAP discussions on named collisions assessment project and represent ICANN on RSSAC, the Root Server Advisory Community.

We have also a number of project activities. One of them is called Internet Technology Health Indicators or ITHI. It's a project that I've been running for a number of years. We do a number of measurements with what you call metrics. And I'm just going to give you an example as how many resolvers does it take to really cover 50% of the users? And we had no idea. We tried to do that when there was a conversation about concentration of DNS resolution. So we came up with a definition of a metric and then we measured something. But we had no idea if it was a good number or a bad number. But it's a number that we can track the evolution of. And if we see this number going down, it means

that some more conversation is happening. And that is, we hope, food for conversation to be held in a community forum like what we have this week.

Another example is we had the metric showing that about 50% of the traffic, DNS traffic to the root, was actually not really DNS traffic. It was some kind of a side effect that was created by a specific software. 50% is a lot. So other people found out this issue and published some papers. We did, too, talk to that specific software vendors and they actually changed the behavior. And what we have been tracking this all the time and now we're down to about 7% from 50%. So we see that there is a clear good trend in this where those requests are going down. They're not going down to zero. So another research project is how is this long tail really shaping? It might be simply a question of this has been distributed to millions of phones, for example, that never update. So it's very interesting for us to see what is going to be the noise level of this.

So there are many other projects. I'm not going to go through all of them, but you can see them and talk to me later this week. John talked about the OCTO documents series earlier. We started this in 2019. We have about 35 documents in it and every time we are adding some more. A few of them on recent documents on DNS key life cycles. Another one that we're going to talk a lot this week is challenges with alternative name systems. It's a shorthand for blockchain names. DNS resolvers in the EU, quantum computing, a bunch of documents edited by OCTO people. There's a link at the bottom of this page. There are some really good documents in there. I'm not talking about the ones

that I've not written, the ones I've written are mine. Please go read the others. And that's it. Back to you, John.

JOHN CRAIN: Apparently, it's back to me. So where do we talk about what we're doing? We have blogs. We have those OCTO documents, of course, and we do speaking engagements. Apparently, we spoke at ICANN 73, but we also spoke at 74, 75, 76. Obviously, we come to all the ICANN meetings, but we go to many others. You will see our staff out and about in your communities, especially if you're in the technical communities. You will bump into ICANN staff from OCTO. Please be nice to them and say hello because we work for you. As a secretariat, as ICANN, we work for the community. So we, especially in our group, we get our kicks. You know, that's what drives us is being out amongst the community and actually helping answer some of these questions. So I want to actually turn over to questions.

SIRANUSH VARDANYAN: Yes. So please go ahead.

NOVECK GOWANDAN: Good afternoon. Thank you for sharing. My name is Noveck Gowandan. I'm from Trinidad and Tobago and I am a first time fellow. So my question is, I noticed in the presentation the office, well, the OCTO does a lot of work in the communities. So you are involved in the communities. But is there a mechanism for community involvement in OCTO? And I say that as a researcher with on my core specialties of research would be governance, organizational behavior, digital

transformation. And I also have a strong technical background. So I was just wondering if there's a mechanism for me as a community member to work with OCTO.

JOHN CRAIN: So there are multiple ways you can work with OCTO. So certain projects and actually with the IANA, too, right, certain projects we need skills from the community for. So a good one would be KINDNS. There was a community discussion about what the norms were, etc. Another one would be the actual algorithm rollover discussion. That was not staff sitting there having those discussions. We bring people from the community in.

Research is a little harder. Now, we did mention this morning that we seem to hire a lot of fellows. So sometimes we have jobs. Yeah, that is a little harder. But if you have things, well, you can always read our OCTO documents and comment on them. And if you have something you just like to get our input on or you have a thought, we're easy to find. You'll find that especially if you have a good research question, our guys love those. That's what they thrive on. So do talk to us. It's not an official way. For example, we don't have ways to bring students in for the summer and things like that. I wish we did, but we just don't. Maybe one day in the future. It is one of our wishes. Thank you.

SIRANUSH VARDANYAN: Before I move to Chuma, there is an online question from [inaudible]. Oh, Adiel, you want to add? Okay.

ADIEL AKPLOGAN: Yeah, I think I think there are other—to add to what John mentioned, some of the research that Alain does also require measurement. And on those measurements, I think you can talk to Alain to how to contribute to some of the metrics that he measure he measures on, which are one way. We also have an initiative that is there that we have created for the community and for people interested in Internet Identifier to engage with us, which is the SIFT, the Special Interest Forum on Technology. The platform is out there and our goal with that platform actually is to allow people who are interested to interact with us and highlight, for instance, things that they want to look into. And that platform is there. It's not actively used, but we are looking for people who are interested to activate that. It's not for us as staff to make it live. It's a community stuff and we would like to hear more from it.

SIRANUSH VARDANYAN: So just to remind, we have only three minutes to go and we will not be able to take too many questions. So the next question is from [inaudible]. What steps is ICANN IROS taking to ensure that the Internet Identifier system are resilient to our new and emerging threats?

JOHN CRAIN: That's what we do every day. Every piece of research is about that in some way. So especially the technologists, the folks over on Matt and Alain's side here are looking at new technologies. And don't forget, it's not just about threats. It's also about opportunities. So we are constantly scanning the horizon and doing environmental scans about what is happening out there. And if it becomes to practical things, if there is an issue somewhere in the ecosystem, we also help. So if you're

a top level domain manager, so that is a part of the ecosystem and you are having practical issues or security issues, many of them know to reach out to us and we will lend our expertise because, as I said, we work for the community. Chuma, please go ahead.

CHUMA AKANA: Thank you. My name is Chuma. I'm a first time fellow. So my question is, what impact does the research conducted by the OCTO, what impact does it have on ICANN? Because you mentioned earlier that you only implement policies, but can that research be also used to drive policies or is it merely academic?

ALAIN DURAND: I would like to provide two examples. The first one is what I mentioned earlier, where we were seeing this traffic at the route, 50% at the route and reduced down to 7%. It was not really an impact on policy, but it was a clear technical impact on the way the Internet operates.

Another one I want to mention is the conversation we had, especially in Europe, and on the paper I wrote on DNS resolution in Europe. Before I wrote the paper, there were random numbers that were floating around in the conversation. Like, oh, this particular player controls 80% of the DNS resolution. While we ran some measurements and we observed that it was not 80%, it was 2%. And we demonstrated this and we published all that. And that provided really clear data to that conversation, which was a policy conversation. So those are two examples of the impact of what we do.

JOHN CRAIN: And one that is relevant in discussions at the moment are the data we've produced around DNS abuse. So there's a lot of discussions. There are both contractual and policy implications of that. We bring data and facts into the policy discussion and that makes for better policy, in my opinion.

SIRANUSH VARDANYAN: Thank you. I think that we will not be able to take more questions for now, but you know these people already. So you know these faces. I really, really love these sessions every time, because for me, being not a technical person, it's very important to understand that we as a non-technical community also have our place and influence in technical aspect of ICANN and how Internet works in general. So with that, I would like to [inaudible] you well. Thank you. Thank you for coming. And our applause is to our panelists. Thank you very much. I thank also our technical support and our interpretation for this session. With that, our meeting is adjourned.

[END OF TRANSCRIPTION]