MONTREAL – How it Works: Regional Internet Registries
Sunday, November 3, 2019 – 17:00 to 18:30 EDT
ICANN66 | Montréal, Canada

LESLIE NOBILE: Good afternoon everyone, my name is Leslie Nobile, I am with ARIN. My partner in crime today is Paul Wilson, he is with APNIC. We're going to sort of share the presentation and we're going to talk to you about How It Works, The Regional Internet Registry System. We know that most of you at ICANN are very familiar with domain names and domain name registries and registrars, but probably less familiar with the Regional Internet Registry system.

So, the five RIRs are AFRINIC, APNIC, ARIN, LACNIC, and the RIPE NCC, and we'll talk about each of those in turn. We'll start by telling you who we are and what we do. Then we'll do just a really brief Internet Number Resource primer. Some people don't actually know what an IPv4 address is or an autonomous system number, so we have just a very general slide on each of those to familiarize people with those terms. And then we'll talk to you about significant happenings at the RIRs.

There's been a lot going on, particularly since IPv4 depleted in 2011, the global supply of addresses depleted, and the push toward IPv6, which has been happening really with the RIR since 1998. With IPv4 depletion, there has been the emergence of an IPv4 transfer market, a very active transfer market, so we'll talk about that. And with that transfer market and with the depleting supply of IPv4 addresses,

we've all seen an increase in fraudulent activity, so we'll talk about that, as well. And then we'll move into the last section where we talk about some of the tools and technologies that we're developing, some of them to directly deal with some of the security and fraud issues that we're seeing.

Okay, so the RIRs. So, I go backwards to talk about the history. It's a very brief history and it only really is dealing with internet number resource administration, although domain names are actually included in this. So, we'll sort of go back to the 1980s and we'll move through the 1990s. The administration of all domain names, IP numbers, protocols, were contracted by the US Department of Defense, the US DOD was one of the innovators and developers of the internet, as most of you know, and they needed to have this administrative function handled by someone.

So they contracted this to the University of Southern California, the Information Sciences Institute, in particular to Jon Postel. He was one of the early contributors to the development of the internet and internet protocols. And so eventually this became known as IANA, the Internet Assigned Numbers Authority, and today we still all refer to this function as IANA, you'll see in some later slides. So, Jon couldn't handle all of this by himself so eventually, I think in the late 80s, the government decided to add a supporting contract to support what he was doing.

So the registration and support of the administration of names and numbers was contracted first to SRI International and then in 1991 to

Network Solutions, and I actually worked at Network Solutions in 1991, so I'm pretty familiar with this stuff. At that point we were handling all domain names and all IP address issuance for the entire globe.

Around 1992, the Regional Internet Registry system was formed and regionalization began. Some countries decided that the US government shouldn't really be determining the business needs of other governments and other businesses within their own countries, and that it would be best done locally, you know, regionally.

So the RIPE NCC was split off from the rest of that contract and they became a registry for IP numbers and autonomous system numbers within their region, and I'll actually show you a map and tell you a little bit more about that. So, that was the beginning of the IP number resource administration splitting away from the domain name administration, as you know, IP numbers are finite resources and domain names are essentially infinite resources. So there really was no reason for them to be together, particularly because domain names started being, you know, became worth some money.

So that function was split and eventually the US government decided with the growth and explosion of the commercial internet the government split the administration of the commercial internet from the military internet. So the commercial internet in the mid 90s was called the InterNIC, that was a contract that was handling domain names and IP addresses for most of the globe, with the exception of a

couple of the regional internet registries who had already formed, and then the DDN NIC was handling the military portion of the internet.

So, what is an RIR? Well, RIRs manage the allocation and registration of internet number resources in a particular region of the world and we maintain a unique registry of all IP numbers issued. IP numbers include IPv4 and IPv6 addresses and autonomous system numbers. Who are the RIRs? Well, I'm going to start in chronological order.

I already mentioned that the RIPE NCC was formed in 1992, that was the first registry to split off from that government contract. In 1993, APNIC was formed. Let me just explain, sorry, regionally RIPE NCC handles all of Europe, the Middle East and the Ex-Soviet Union countries. APNIC was formed in 1993 and they handle all the Asia Pacific region. At this point in 93, the InterNIC was still handling all domain names and all IP addresses for the rest of the globe.

So in 1997, ARIN was finally put into place and we handled, at that point in 1997, ARIN was responsible for all of the US, Canada, the Caribbean, Central and South America, and Sub-Saharan Africa. But we work together with the three RIRs to split off some of the function, so LACNIC was formed in 1992, they handle all of Central and South America, and they split the Caribbean region with ARIN. Then in 2005, AFRINIC was formed and they are responsible for the African continent.

The core functions of an RIR, we do lots of things and we do lots more than what is there, but really the core functions, what we're known

for, what we have to be doing on a daily basis, is to manage, distribute and register internet number resources. We maintain directory services, and in particular we maintain a registry called WHOIS which most of you know, and we also maintain routing registries, which contain routing information which I will touch on in a later slide.

We also provide and run reverse DNS. So we provide the systems that actually do the translation of an IP number into a domain name. That's reverse; forward would be the opposite, domain name into an IP address. We support internet infrastructure through technical coordination throughout the globe with our industry partners globally and we facilitate community driven policy development process.

It's important to know that all policies within all five of the RIRs are developed by the communities. They propose policies and the staff implements it. We facilitate the process and then we implement the policy. So we don't make policy, the community in a bottom-up process makes the policies. So, in general, the RIRs are independent. There is no government oversight.

I often get asked, well, does the US government run what you do? No, the US government does not. There is no government oversight of any of the five RIRs. We are not -for-profit organizations. We're 100% community funded, our fees for our services, things like issuing resources, providing reverse DNS, maintaining WHOIS or the WHOIS registry, et cetera. The fees are not actually for the number of resources themselves. They are for the services. We are all

membership based and our memberships are open to all holders of internet number resources.

And typically, that would be internet service providers, telecom organizations, governments, corporations, universities, et cetera. We are all community-regulated, as I mentioned the policies are developed by the communities. We all have member elected Governing Board of Trustees and everything is open and transparent, it's all documented on our websites, all of the Board of Trustees meeting minutes are documented, pretty much everything is public, publicly available.

So, the Number Resource Organization is the five RIRs operating together. We're all distinct organizations, but together we are known as the NRO. We act as a focal point for internet community input into the RIR system and basically we were put into place, the NRO was developed in order to promote and protect that bottom-up policy process that I mentioned, and the unallocated number resource pool.

This is a very old slide from ICANN, I hope they don't mind I'm using this one, but to me it illustrates where we fit into the bigger picture. I always put this up because we all know that ICANN was put into place by the US government, by the way, to divest themselves from running the internet. They started trying to do that way back in the early 90s and finally were successful to stop running and governing the internet.

But ICANN is responsible for the top level technical coordination of all domain names, numbers, internet numbers and root servers. They're

nonprofit, like we all are and self regulatory and global, and then they have a variety of supporting organizations, and I'm sure a lot of you are familiar with the ccNSO and GNSO, the domain name supporting organizations. We are the address supporting organization, the ASO, that is where the NRO comes into play. In this community we're known as the ASO but we are the number resource organization operating as the ASO within the ICANN community. And then there's a variety of advisory committees here.

So, this is the Internet Number Resource primer, I hope I'm not boring you all with this one, but not sure if you know or not what they are. An IP address is a unique numerical address and it's assigned to every device that connects to the internet to a TCP/IP network, and facilitates moving information and data packets across the network.

So there's two types, there's IPv4, that was the original, developed I think in the 70s. It's 32-bit address, there's about 4.4 billion total IPv4 addresses. And that's an example of a typical IPv4 address. In the early 90s, mid 90s, when the technical community, the Internet Engineering Task Force realized that IPv4 was going to deplete in early 2000s at some point, they created and developed the IPv6 protocol, which was to replace IPv4 and it's just a huge amount of space. It has some enhanced features over IPv4.

There are, let's see, two to the 128th total addresses and 50 octillion for each of the roughly 6.5 billion people alive. I can't really fathom that number, maybe you can. But it's a huge amount of space. And that's a typical full IPv6 address at the bottom, that's an example. So,

Autonomous System Numbers are used specifically for routing. And all network operators actually have to have an ASN to control routing within their own networks and to exchange information, routing information, with other network operators and internet service providers. It's a globally unique number. It's used to exchange routing information between neighboring autonomous systems and to identify the AS itself.

An autonomous system is a group of IP networks, they are administered by one organization, something like a university with multiple networks would run their own autonomous system. And routing, of course, is the act of moving information, the data packets, across an internetwork from a source to a destination. So we issue autonomous system numbers strictly for routing.

I do put this slide in when I know that people often use domain names and IP addresses interchangeably. So I just have to always stress this, IP addresses are not domain names. Distinctly different objects, distinctly different functions, distinctly different registries. IP addresses are identifiers. Computers recognize numbers, that's how they talk to each other.

So the unique number identifies a computer on the internet or device, it's used for routing, moving that information across the network from a source to a destination. And every device directly connected to the internet uses a unique IP address. Domain names are reference objects, they are used to make it easy for people so that people can type names, instead of those long numbers. People recognize names

they will map a host name to a unique IP address and it's a hierarchical system, it's a means of storing and retrieving information about host names and IP addresses in a distributed database.

So, this is how IP addresses are issued, and remember, I mentioned IANA, Internet Assigned Numbers Authority, that function still exists. And they manage the global unallocated IP address pool. it all sits with them, and then they will allocate space to the five RIRs. And those are done by global policies, it's all clearly laid out how we get to qualify for additional space. And when I say allocate that means that we can take the numbers and sub-delegate them.

So we don't just use them within our own network, we get to sub-delegate them because they're allocated to us. So we manage our regional unallocated IP address pool. And then we can either allocate space to our ISP customers or we can assign space to our end user customers.

If we assign it to the end users, that means they use it within their own internal network infrastructure and it stays there, it doesn't get further sub-assigned to other customers. When we allocate to an ISP, they get to take that space and then sub-delegate it. They can either reallocate it to their downstream ISP customers or they can assign it to their end user customers who will then just use that space within their own internal network infrastructure.

So, we'll move into some significant happenings at the RIRs. Probably the most significant in the past 15 years has been the global IPv4

depletion at IANA, and that happened in February 2011. Each of the RIRs receive their final /8 and those are photos of the CEOs at that time with the IANA Executive Team.

So, looking at the IPv4 space currently available in each of the RIRs, this is measured in /8s, and you can see AFRINIC has the most space, almost a quarter, 22% basically of a /8. APNIC has 18, 19% of a /8. ARIN actually has no space available to issue. LACNIC has whatever, 10%,m and then RIPE, actually RIPE is just about at zero at this point, this was as of September 30th, I think, and I believe that RIPE is now down to zero, they're completely depleted, so this slide will change soon.

Looking at IPv6 allocations, we show this just to show the sort of slow, steady growth. This is total number of prefixes per RIR, per year. And basically, what it really illustrates is that most of the growth in IPv6 is actually within the RIPE region, that is the tall bar, and we anticipate to probably surpass 2018, because we still have a quarter left to issue space. So it's slow, but steady. It's just trending up.

This shows the total IPv6 address space allocated in /32s and you can see that RIPE NCC, as illustrated in the previous slide, has allocated the most number of /32s to their ISP customers.

This is a confusing slide, so don't pay attention, I'm just going to tell you, don't look at those lines. So in AFRINIC, we're looking at the percentage of our members who have gotten their IPv6 allocations. In AFRINIC 47% of all the AFRINIC members have gotten an IPv6

allocation. In APNIC it's about 64% in total have gotten theirs. In ARIN, it is about 60% of our members have their IPv6 allocations, in LACNIC it's 95% of their members have IPv6 address space, and I don't know if there's anyone from LACNIC here, I believe it's because they had a policy that basically allowed them to issue IPv6 address base to every single one of their members. So it was essentially handed to every member, which explains that huge Increase there. And then in the RIPE NCC region about 64% of their members have got their IPv allocations.

So, here's some current observations and sort of a summary of some of the slides. I just showed you. Movement to IPv6 has been slow but steady. It is trending up slowly, sometimes there's a little bump, it goes down and bumps up, but it is steady. ISPs are slowly rolling out, IPv6, not all of them are buying into it yet, we hear this all the time.

Sometimes it's hard to sell that to their management. There's a steady increase in IPv6 traffic and we've seen that, we look at some of the statistics that Google and other large providers are putting out and I think what is it roughly 25% of the traffic, David, do you know? About 25% of the traffic being routed, or it may be a little bit more, is IPv6. So it's definitely getting there. And we do see an increase in IPv6 requests, all five of us do.

But what we're also seeing is that there's still a high demand for IPv4. It is not going away anytime soon. We are all receiving significant numbers of IPv4 requests, even ARIN even though people know we have no IPv4 address space, they still come to us in droves, mostly in

the form of IPv4 Transfer Requests, but we're still getting regular before requests, as well. But what we're seeing is that customers are increasingly turning to the IPv4 market for address space.

When IPv4 depleted, an IPv4 marketplace developed immediately and IPv4 brokers started popping up everywhere and really urging people to start using their space, if you're not using it, why don't you give it to somebody who has a need for it. So that is sort of what's happening right now. So we sort of see a variety of things happening. We see people, customers purchasing IP address space from other registrants who have space to give away, they don't need it.

So people are purchasing space and then they're using RIR transfer policies to update the registries. They're saying we got this space from A and we're moving it to B and the RIRs so all their verification embedding and we make that change in the registry. We also are hearing and seeing people purchasing space outside of the registry system. So they're not actually updating the RAR, it's essentially black market transactions happening.

Another thing we're seeing a lot of is leasing, people are actually having leasing agreements to use space, not to permanently purchase it, but to lease it and pay on a monthly basis. And what they do is they take letters of authority to an upstream ISP to get that space routed. So as I mentioned with the depletion of IPv4 there was the emergence of an IPv4 transfer market and that's really where a lot of our work is being done right now. There is this ongoing demand and decreasing

supply of IPv4 addresses and because of that it necessitated RIR policy changes.

I mentioned that there's transfer policies happening that allow this to take place. So our choices, and the community determined this, the choices were either you facilitate IPv4 market transfers and ensure accurate registry data, or you watch a black market emerge with no registry interaction. So that was the choice. The community didn't necessarily want to promote buying and selling of IP addresses, but there essentially was no other choice.

So all five of the RIRs have implemented needs based IPv4 market transfer policies and they allow IPv4 resource registrants who have extra space to transfer space to qualified recipients who have a need and can justify to each of the RIRs. This is all still being done based on justified need. So the RIR's role is to ensure full compliance with the needs based policies and to update and maintain the accuracy of the registry. That's our role in these market based transfers.

The RIRs are not privy to or involved in any of the financial transaction information that happens between transferring policy parties. We actually have no information about any of that. When it comes to us, all we're doing is looking at them and justifying them using existing criteria.

So, this shows each of the RIRs and their market based transfer policies. So there are I mentioned intra-RIR transfer policies at all five of the RIRs. In other words, market based transfers are happening

within each of RIR regions. There are also inter-RIR transfer policies. At this point, APNIC, ARIN, and the RIPE NCC have those inter-RIR transfer policies, those market based transfers, and that's where all of the action and activity is happening.

But LACNIC has recently passed that policy so it's pending right now, but it is supposedly going to be implemented in Q2 of 2020, so they will be the fourth RIR have inter-RIR transfer policies. And then in AFRINIC there are multiple versions of a policy in discussion, it's been going on for a couple of years, and we'll see what happens, I'm not sure where that is going.

So, looking at intra-RIR IPv4 market based transfers, you can see the real growth, particularly since 201 was where it mostly started, but it's really been sort of booming in 2017, 2018, and again we anticipate that to happen, to increase for 2019 and probably surpass 2018, that is more likely the scenario to happen. Most of the intra-RIR transfers are happening within the RIPE region. You can see there's slightly more there than in the ARIN region, ARIN is in blue and RIPE NCC is in that dark green.

So this is the confusing slide. Okay. Basically, inter-RIR market based transfers. Those are the three RIRs that are doing it right now. You can see that most of the space, most of the inter-RIR transfer policy transfers are happening from ARIN to RIPE NCC and to APNIC, you can see the red arrows, which are confusing, so over 200 in each direction, but it's pretty consistent between APNIC and RIPE NCC and ARIN.

But coming out, as far as transferring out space, it's coming mostly from ARIN. The reason for this is that ARIN has, so we inherited the internet database. And that meant that we had all of the space that was issued by Jon Postel, as I showed you, and the InterNIC and the DOD NIC, all of that space was inherited by ARIN and we have the majority of what we call Legacy space, space that was issued prior to the establishment of the RIRs with no contracts, no terms and conditions.

We actually had a request form and there was actually justification. We asked lots of questions and then we would give people big blocks of space. Those people still have that space, and many of them are choosing to sell that right now. So, that is why the large numbers of space is going out of the ARIN region to the other regions.

So looking at our current challenges, many of them brought on by IPv4 depletion and this IPv4 transfer market, we're all seeing more fraudulent requests to obtain and/or transfer IPv4 addresses. Some of the RIRs you saw still have IPv4 addresses and we in ARIN are doing a lot of transfers in the market based transfers. So, lots of fraudulent requests, people are very creative. And as the market value increases, we're seeing more fraudulent activity. So it's becoming really interesting, it's very challenging for all of us.

Actually, some of us are working quite heavily with law enforcement on this because obviously there's fraud happening on a daily basis. So there's lots of hijacking of IPv4 addresses and much less so of the autonomous system numbers and they come in, as I said, they target

those dormant out of date records, mostly the Legacy space, things that have registration dates from 1995 or 1992, they're looking for space that's not routed it's not announced, it hasn't been updated in many years and they go right after it, and they try to pretend that they are that registrant, which we all know they're not.

So, they go to some pretty drastic extremes in the RIPE NCC region, they submit a lot of fraudulent documents, but in the RIPE NCC region they asked for passports as identification for points of contact, and we're actually seeing forged passports, I mean they're really going to extremes.

And then in all the regions, particularly in the ARIN region, we see a lot of these companies setting up shell companies. So they're looking for domain names that have expired and company names that are out of business and then they'll set up that exact same company name and they'll register that exact same domain so that they can try to acquire the space that was registered to those older companies. So there's lots of shell company activity.

And then we're seeing a lot of route hijackings. Although the RIRs are not directly involved in routing. if it is reported to us that someone has taken over space and is routing it for malicious purposes, we will actually contact an upstream provider and let them know that they need to check the registry to ensure that they're routing space for the proper registrant.

We don't go to too many extremes but we do let them know. But it is unauthorized use of abandoned or unrouted IPv4 addresses. And I work heavily with law enforcement and this is one of the biggest things they're seeing and one of the things they're very worried about, is this BGP hijackings as they call them, the route hijackings, because even if someone takes over a route and gets to route the space for a week they can do a lot of damage and cause a lot of damage for people. So it's a big criminal loophole there.

As I mentioned, we're seeing lots of leasing buying, selling of IPv4 address space outside of the registry system. We have people who will not update and validate their contact information in WHOIS. All five of the RIRs have tried to maintain an accurate registry and we have policies or procedures where we try to validate every single point of contact in our database. And some people just won't comply. So that becomes an issue.

And then Carrier Grade NAT is a problem. So, Carrier Grade NAT, an essentially simplistic version of this is it's using one IP address for multiple users. Then it becomes difficult to identify the individual subscriber who is actually using the space. And it really is a problem for law enforcement more so than us as registries.

So, I get to turn it over now to Paul Wilson. However, if there any questions from anyone of that previous section, we're happy to answer anything. Does anybody have a question?

UNKNOWN SPEAKER: Thank you. [Inaudible], I am asking in a personal capacity. You mentioned that there are certain Legacy addresses. Should there be an offer to bring them to the registry? And should it be RIR, should it be ICANN? That is something which remains and some of these addresses probably are used for what we hear, dark net or some such. And you also mentioned in your slide. So any comment on that? Thank you.

LESLIE NOBILE: I will, and then Paul will, I think. I know in the ARIN community there was a lot of discussion about separating the Legacy space, having a separate registry, forcing Legacy registrants to sign contracts. What we've done in the ARIN region we have a registration services agreement, a contract that they can sign, that is actually a light version of our regular registration services agreement, to bring them into the system, to protect the resources, and we found that we've had a certain percentage of them that have agreed to.

But quite honestly, in the ARIN region, 45, 47% of all of our networks are actually Legacy networks, they were issued prior to the establishment of ARIN. And right now, that's where they stand. I'm not sure, Paul, do you have anything to add?

PAUL WILSON: I think the most important thing is that the Legacy address space has been redistributed well over 10, 15 years ago to the responsible RIR. So each of the RIRs has got a geographic region for which it's

responsible. The Legacy addresses that were allocated to organizations in each of the RIR regions, the registration records for those address blocks were transferred too or taken over and are under the authority and responsibility of the applicable RIR at the moment. And so that's been done for quite a long time.

As Leslie said, the vast majority of the Legacy space was allocated in the ARIN region and it's under ARIN's authority, and that's why these days we see the vast majority of transfers coming from the ARIN region to other regions.

But the other agents are responsible, for instance APNIC is responsible for Legacy records, Legacy blocks that were allocated to organizations in our region and we basically apply the same general policy, which is the address space which is routed should be registered in the registry and as Leslie said, if it's not registered and it is routed then it's a hijack or a bug on sometimes referred to, and that's something that tends to get pursued by security folks and other operational folks are monitoring the routing table and exactly looking for unregistered space that is routed.

So I think it might be a claim or a perception that the Legacy address space is managed, but it is managed under the RIR, under the authorities of each of the RIRs, it's just that in some cases the early registrations were made and still exist with little or no contractual obligations associated with them.

So back in the old days, those allocations were made very loosely with almost no or literally no terms and conditions or understandings. And so there's an, I guess you'd say a legally ambiguous or legally complicated issue of whether an RIR is really able to enforce policies for instance, for the maintenance of registration of those blocks. Sorry, that's a very long answer.

KEVIN BLUMBERG: Kevin Bloomberg, I just wanted to add to that. We just came back from the ARIN meeting a couple days ago in Austin, Texas. And one of the slides that ARIN as the organization shows us is a consistent conversion of Legacy space into a space that is under a registration services agreement, it's not 80% suddenly, you know, this big number, but it is a consistent move from Legacy over towards.

So, that's actually very good to see because then it is under that I believe historically what Paul said, and I can sort of expand on it as being somebody from North America, having a very ambiguous contract from many, many years ago is just a recipe for a lawsuit today. So, that is why you have to tread very carefully in the North American region, unfortunately, because things can be very litigious here and I think the community has said let's move to IPv6, let's get this done. If we can get stuff in from the transfer market into registration, that's great, but the cost, not just financial, but the liability of it was not worth it.

STEVE CONTE: Were there any other questions?

PAUL WILSON: Could I just ask, my name is Paul, I'm the head of APNIC, the IP address registry RIR for Asia Pacific region. Can I just ask about the audience? Do we have we have folks here who are fellows to ICANN, is that right? Okay, that's great. Welcome. So I guess this is mostly for you. So, if you do have any questions, please ask at anytime. So I'm going to go on and talk just for a few minutes, as I think we're a good two thirds or three quarters of the way through, you'll be glad to know. But I'm going to talk for a few minutes about some of the, as it says, the tools and technologies which are used by and implemented by the RIRs.

So the first one is WHOIS, and if you've been studying the ICANN DNS world, then you'll know about the WHOIS database as the registry database for domain names. We also use the same term because it's the same basic technology which is used for the same purpose, which is registration information, but in relation to IP addresses. So the idea of WHOIS for IP addresses is that you go to the right WHOIS server, you can enter a query that relates to a particular IP address and hopefully you find the authoritative information about who actually holds that IP address.

And the reason for that is in the same way as originally with names, the reason that you might want to do that is because, particularly from a network security or maintenance perspective you might have

an attack or you might have some traffic, or you might have a transaction, like an email that's associated with an IP address at its source, and you want to know where it's come from. And so you can do that by looking up the WHOIS for IP addresses and you can find the registered holder of that address which can be quite useful for different sort of diagnostic purposes.

So what does the IP address WHOIS, what does it contain? Well it's both IP addresses v4, v6, and autonomous system numbers for space that's primarily issued by the RIR running that WHOIS. So there are five RIRs, each of us actually runs a separate unique WHOIS server for the address space that we register. So the RIR will put into that WHOIS the records for IP and AS numbers which we have allocated or which are classified as Legacy, as we said before, we all do administer some Legacy space for our region.

We register information quite similar to the DNS, so it'll be the organization, who is responsible for that particular block of addresses or autonomous system number and points of contact or contacts or person objects, depending on the RIR WHOIS terminology which are the literally the people who are responsible, administratively or technically for those things.

Now that data is referred to here is as public registration information, that's quite important from the RIR's perspective in that this WHOIS is really critically important for technical maintenance, diagnostic and security purposes on the internet. And so the information about who

is actually the registered holder of particular IP addresses is very important as public information. Is there a question back there?

ABDEALI SAHERWALA: So, I just wanted to know whether it can be accessed. Abdeali, I'm from Toronto. I just wanted to know whether this can be accessed by the public, but you already kind of said that, so would it be like public information or do I need to get some sort of like a law type access to information, or can I just like find it in a heartbeat?

PAUL WILSON: The registry itself is by definition public and everything in the public registry is public, by definition, so yes, anyone can query, the registry of any of the RIRs and anyone and everyone does. So we're answering literally millions of queries a day and anyone can do that. There might be other information that's held by the registry for instance that pertains to the membership of the body who has received IP addresses from us.

So there may be other information that is held privately by the registry and that's not available publicly. But if there happened to be some legal reason why that was of interest to law enforcement or something then that's only just subject to normal sort of subpoena type processes, I guess. But the registry itself, what we call the registry, is public. And it is intended and the purpose of it is intended to allow you to identify specifically the person or the organization responsible.

So that's quite important in the context of the GDPR, for instance, because you might think that this is private personal Information therefore shouldn't be accessible but it is actually one of the conditions of the allocation of IP addresses that the organization themselves receiving those addresses be contactable because it's really a recipe for disaster if that is not possible because it opens up all sorts of potential for undiagnosable or untrackable security issues on the internet.

Okay. Something that is available in the WHOIS database, as I said, there are five different databases for five different registries in it, and in some cases you'll find reference information. If you query one WHOIS database for a resource that doesn't exist there, then you will find some sort of reference pointed to the database where you can find more information.

Okay, RDAP is something that also you might have heard about in relation to DNS related WHOIS, and RDAP is a much more modern version, effectively a modern version of WHOIS. It stands as you see for Registration Data Access Protocol.

And RDAP is designed specifically to overcome some of the difficulties, some of the limitations that exists with WHOIS. WHOIS was invented and released, I'm ashamed to say, in the 1980s and it's still the primary vehicle for registration data that's available for names and numbers, and RDAP is much more recent and is designed to do a much better job. So there's quite a few things that RDAP overcomes and brings that WHOIS lacks, and those are referenced here.

A standardized command structure data structure for responses and error behavior. So one of the problems with WHOIS at the moment is that you can go to different WHOIS servers, even different RIR or different DNS WHOIS servers and you actually get information back in different formats.

So RDAP is a standardized format. It also standardizes redirection, as I said before, if you go to one database that doesn't have the record that you want, all you get if you're lucky is it points to the other database that you then have to query and RDAP will automatically redirect and find the data that you're looking from the right database without you having to do anything.

It's also got, and this kind of pertains to the question before, if it's needed, it's got the ability for different levels of authorization so an RDAP database can contain information that's confidential not available to the public, but only available for instance with the right authorization. And it also supports internationalization, so you can have character sets other than Latin, English or European type characters.

So, ICANN has issued requirements, I'm not familiar with exactly what they are, but there are some requirements that accredited registrars and registries are needed to implement RDAP these days. The regional internet registries are also agreeing amongst ourselves and coordinating to provide, to use RDAP as our main source, our primary source of registry information and registry information services in the future. So that's a work in progress at the moment.

Okay, another modern registry service is something called RPKI, Resource Public Key Infrastructure. For those of you who know what RPKI is, it's a hierarchical system of digital certificates which are issued by certificate authorities to holders of private keys to allow you to use those keys for security purposes. Now the RPKI is effectively a PKI structure that allows us to issue digitally signed certificates that represent the holdings of IP addresses.

So you might think about them as WHOIS records with a digital authorization and digital validation mechanism. It's based on X.509 which is the same as the normal RPKI and email certificate system. And it's basically a general purpose standards based system based on standards from the IETF that's meant to be a general purpose system for anything that you'd like to use it for. But it was primarily built to support routing security.

So Leslie mentioned before the phenomenon of route hijacking where someone might try to use IP addresses which are either unallocated or which belong to someone else and the idea of RPKI is to is to provide an ability for the actual holder of those addresses to secure those addresses and stop other people from using them or for the RIRs for that matter to stop people from using addresses that aren't actually allocated to them.

So there's a few different applications of RPKI which are in relation to that security area. One is called Route Origin Validation, another is called Resource Tagged Attestation. There's also something called Secure BGP that's in the pipeline as another set of protocols that use

RPKI for security of the routing system. So, that's also something that the five RIRs are collaborating on together to launch those services. And it's a work in progress, so there's a few stats here about what the uptake is of RPKI around the different regions.

We've got about a total of 14,000 organizations around all the RIRs, that's probably about 10% of the RIR customer or member population, which are currently using RPKI. So we've got a fair way to go before RPKI is in very widespread use for routing security, but it's something that's underway. We can see here that the actual adoption rate on a per address space basis is quite variable across the RIRs. RIPE NCC is doing pretty well, actually, with nearly 40% of IPv4 address space now actually being covered by RPKI, 26% of IPv6. It's just a different rate of adoption in different RIRs.

So, a bit more about one of the applications of RPKI which is the use of Route Origin Authorizations or ROAs for what's called route origin validation, and that's the process by which, if you want to use a block of address space, that is, if you want to set up a network and connect that network to the internet, then you use the process of route origin validation as the one which says, are you permitted, are you authorized to actually inject that address space to connect that particular block of addresses to the internet?

So, a ROA, route origin authorization is assigned object assigned statement which simply says that as the holder of a set of IP addresses I am signing to give authority to a particular autonomous system number to originate that address space. So it's a use of RPKI, it's

something that is now actually being quite actively used and it will be more and more used in the future, because we're at a stage now where certain quite significant global providers are requiring route origin authorization to be provided in order to connect with them. So they're listed there.

Amazon is in the process of requiring ROAs from a certain class of their customers, there's also Cloud Player, AT&T, Google, and also quite a few IXPs internet exchange points require ISPs that are connecting to the IXP to have their route origin authorizations registered within the system before traffic will flow.

And what that allows is it allows IP address blocks which someone is trying to bring into the exchange or into the provider to actually be dropped by the by the provider automatically to prevent IP address hijacking or squatting on IP addresses, or even quite a common thing which is simply a typographic error that an engineer might make, they tried to set up the routing of their address space and they accidentally enter someone else's address space by simply mistyping a digit, and that's quite a common source of problems. And so, ROA, route origin validation is a mechanism to help with that.

There's a last slide here on internet routing registry, which is actually an old technology. This kind of belongs with WHOIS because it's an extension of the old WHOIS protocol, just to allow ISPs to enter publicly information additional to the IP address registry records that tells the world what they are doing in terms of IP routing in their networks.

It allows other networks to understand what the routing policies and routing configurations are so it's an older type of technology, but it's something that is actually very important in routing security these days, and it's being upgraded by the RIRs generally to provide a higher grade of authorized and authenticated service and also to be provided consistently with the RPKI ROA system as well.

So there's probably not much more to be said about that, but that's where the internet routing registry sort of fits into this into the scene. And I think that's all that I have. So if there are any questions about that, any comments or remarks, for Leslie and myself, then please go ahead.

JOSH GOLD: Josh Gold with the NextGen program. I'm from Toronto. Just on WHOIS, I should have asked this maybe earlier, but is there really a way to verify the information that's provided on WHOIS in terms of the people being who they say they are and generally if you could maybe touch on maybe two to three sentences on the most common forms of abuse of WHOIS.

LESLIE NOBILE: As far as the information that gets into WHOIS all five of the RIRs are very serious about vetting and verifying all the data that gets in there. But that is a more recent practice. I know previously the RIR system was based on trust, trust, but not verify. We've had to change that, as we've started realizing, and ARIN got hit by this first in 2001.

We realized that people were actually lying, hijacking space, pretending to be someone else, and the way we had hit was someone hijacked a /16 that belong to the Los Angeles County Sheriff's Department And it made like national news and we had no idea. So we started working with Spamhaus we started becoming aware, and we started immediately changing all processes, even though probably only 5% of our customers were doing bad things and giving us bad information, we had to change it across the board.

So as each of the RIRs got of hit with this same type of fraudulent activity and falsified information we all started doing a whole process of verification and vetting. So now it's always trust, but verify. So the information that's in there now is heavily verified. But there's a lot of old, bad data in there and a lot of people trying to pretend that those are those old those old registrants. So lots of activity in the older records, the older registration records. Paul, do you want to add to that at all?

PAUL WILSON: No, I think that covers it.

ABDEALI SAHERWALA: Abdeali Saherwala, NextGen program, as well. I'm not familiar with all this, I'm pretty new, but I just wanted to know that each computer and phone have their own IP address, right? Okay, so because of that, can we track bots or trolls on social media to prevent disinformation? Because then you can see like this Facebook account is with me, but

then imagine if I have seven Facebook accounts and why do I need seven, or why does this computer have like 1000 Facebook accounts?

PAUL WILSON: It's kind of interesting because back in those old days when things were more trust based and more transparent IP addresses were very obvious. And if you've ever had a cause to look inside the headers of an email that you received, for instance, and it can be, if you're technically oriented, it can be quite instructive and useful.

You find you've got an email that you received that claims to be from PayPal in the US and you look at the source IP address and you find that it's come from somewhere else entirely that doesn't make sense, then that can be a very strong indicator that you've got a falsified sort of email that you've received and that's really, it was on an understanding that IP address information was generally available that then the WHOIS database also was provided in order to exactly help that sort process of verification. These days, though, that IP address information is available to the CDNs, to the people who are running the apps on your phones or providing the services, but it's really not customary for it to be provided publicly.

So it would be perfectly possible for Facebook, Twitter and all of them to attach with the display of for instance a tweet or a message to attach the source IP address that was used to post that, perfectly possible and for a lot of purposes it would be very useful and very transparent to do that. Some people these days I think particularly

because the current practice is not to do that, some people may well object because they might say well it's exposing me in a way that I didn't expect.

I'm not saying that that's likely to happen, but it is kind of interesting that the practices have changed in a way that actually obscure for the end users what the information that is available to the platforms, it's available to the engineers and it used to be available for other services that we're kind of used to using like email.

ABDEALI SAHERWALA: So they know like, okay, that this IP address has 1000 accounts.

PAUL WILSON: Absolutely, it's not that the IP address has 1000 accounts, but they would know for instance if there were 1000 messages posted from 1000 different Twitter accounts that happened to magically come from one IP address, and that would be an indicator, it might not prove, you know, you have to be careful about what it does, can, and can't tell you. But it could still be an indicator.

If those 1000 posts were apparently from all around the world then, there really is only very specific circumstances in which you'd find them coming from one IP address and that would be so for instance if they all happen to be using one VPN and that VPN happened to have an exit point into the internet from one given IP address. So, like I say,

it doesn't necessarily absolutely prove anything but it can be very instructive to help in a sort of a process of attribution or diagnosis.

UNKNOWN SPEAKER: Is there any estimate, especially when AFRINIC comes into the market of selling IP addresses, is there any kind of estimate on percentage of available unused IP addresses for sale or could possibly be for sale?

LESLIE NOBILE: That's an interesting question. There is no estimate, because what we're seeing is that not only is that Legacy space, that space that hasn't really been used for many years, people sort of held on to it and just forgot they had it. That's being bought and sold and traded, but IP address space that was issued by the registries actually to their customers, there are a lot of people who've decided they don't need all of their space and they're going to sell off some of their space and use some of it for their customers.

So probably with the intention of moving on to IPv6 eventually, but it's become a very tempting market. So there's really no way of knowing who is going to decide to sell their space and get rid of their space in the AFRINIC region, or anywhere else.

What I can tell you is that the IPv4 brokers that are operating globally are literally contacting every single person who has IPv4 address space registered within the five RIR WHOIS databases. They're

actually contacting them and saying did you know that you could sell some of your space and make a lot of money?

And that is tempting a lot of people who you would traditionally not think would get rid of their space to actually sell some of their space. So there's absolutely no way of knowing, but it's becoming so lucrative and worth so much money, that we just don't know. I don't know, Paul, do you have anything additional?

PAUL WILSON: I'd just say I don't think the question relates to AFRINIC. What Leslie said, it's true for all of the RIRs, we all have a significant amount of address space that is allocated and which doesn't appear in the routing tables. That doesn't actually mean it's unused because it is possible under the standard RIR policies to receive IP addresses and not to actually use them on the routing table, but to use them in private networks, in which case you can't tell whether the address space is used or not.

You can't tell whether that address space has been forgotten by the person and is simply unused and in that respect, the brokers are actually doing us all a great service because by actually finding address space that may have been kind of overlooked by its address holders and encouraging them to release that address space to people who will use the addresses and will, by the way, also bring those addresses into current policy and current registration then that's actually a good thing, even if someone's making some money out of it,

it's an option for the buyer to either pay for addresses or to not have access to them and these days it seems that buyers are prepared to pay for them at the rate of about 20 bucks US per address as a sort of going market rate.

So that's an amount of money that adds up a lot, but it's actually a relatively small amount of money compared with the value of that address to an ISP who might be deploying it on a service that that earns them hundreds or thousands of dollars per year. It's great to have these questions, by the way.

JOSH GOLD: Josh Gold once again, NextGen. You mentioned that there's no government oversight of the RIRs, but have any governments tried to have more influence or oversight over these organizations and hopefully not just yes or no, and maybe if you could get a little bit more into that. Thanks.

PAUL WILSON: Yeah, interest is expressed by different governments in a lot of different ways. Some governments have actually made regulatory measures that refer to IP addresses as critical communications identifiers for instance within their country, that gives them some right or authority to require users of those numbers to give them access to usage information, et cetera. It's quite possible for a government to implement or to pass legislation that is meaningless and unimplementable and that's happened in quite a lot of cases.

We've had governments legislate that they are responsible for the management and distribution of IP addresses in their region in their country, even though they have no IP addresses that they can actually allocate or manage. So it's a kind of a piece of legislation that doesn't really have any effect. But these things are there are as many different approaches and attitudes, as there are governments in the world, I think.

So, it's something where the RIRs are still operating today and a lot of what all of us do actually is governmental outreach and engagement because governments are very interested in this stuff, whether the legislation they're passing is enlightened, meaningful or useless, or whatever, it's still a very legitimate interest that governments have and all the RIRs helped governments to better understand what they're dealing with.

JOSH GOLD: Does that threaten RIR model at some point?

PAUL WILSON: It hasn't significantly. There's been some initiatives through the ITU for instance, for the ITU to take a role in IP address management or to create an IP address registry of some kind, but it's quite difficult to find a way that would fit into the existing framework and although there's been a lot of discussion, there hasn't been a decision or a move that that would challenge what we're actually doing. Anything else on

that? We've got quite a few RIR folks and representatives here too. And we could all probably comment on that one.

MATT JOHNSON: Matt Johnson, I'm with INTA. You mentioned that the WHOIS system is somewhat separate or it's like a slightly different system for IP addressing, and I'm just curious because, sort of piggybacking on that question, have you guys been affected at all by the GDPR because at least as I understand in the DNS side of things, it's become quite political. But it seems like your system doesn't really have any privacy in it, or not as much. I'm just curious.

PAUL WILSON: Well the RIR most affected by the GDPR is RIPE NCC obviously because they're the ones with the European subjects. For most of the other RIRs there are a few if any Europeans who are the subject of data in the database. So I think mainly the answer could probably come from Chris Buckridge, for instance, from RIPE NCC.

CHRIS BUCKRIDGE: Hi, Chris Buckridge, I work for the RIPE NCC. It's been a very interesting question for us and as RIPE NCC, we obviously do a lot of as Paul mentioned, do an awful lot of government engagement. We talk an awful lot with colleagues at European Union institutions and we looked at the situation there very closely. Basically, our community was looking at issues of personal data protection a

number of years before GDPR even came into effect or was being developed. And what the community developed was a very clear understanding of why this was operationally necessary information to have publicly available.

And so then, in that sense, when GDPR came in we didn't work with people and talk to people, but we went back and looked at what the community taskforce had come up with in relation to that and found that actually with that operational necessity, it actually worked with GDPR, there was there was an allowance there for that to still be made public, under the GDPR.

And so that's our understanding today of what the situation is, but we do obviously watch with great interest and very closely how things are evolving. I think the GDPR situation is one that is somewhat dynamic and is evolving as the different data protection authorities determine and understand their role and what this all looks like. So yeah, we're definitely following the conversations.

ORE LESI: Hi, my name is Ore Lesi, I'm an ICANN66 fellow. I just wanted to clarify again about the WHOIS. I think you had mentioned earlier that information about individuals would be private, is that my understanding, but information about organizations, I guess, who are the points of contacts could be made I guess accessible through the WHOIS database. Is that a correct understanding? And then I also wanted to clarify in some cases the admin contact is the same as an

individual, in which case I mean they don't have any kind of privacy or data protection.

CHRIS BUCKRIDGE: I'm only talking about the RIPE database here, I'm not talking about other RIR WHOIS registries, but in the RIPE database most of the holders are organizations, but there are also natural persons who hold resources and obtain resources from the RIPE NCC.

As I said, there was a taskforce that did look quite closely at all of these issues including the issue of the different admin and technical contacts in the database, but it was found that, yes, this did still fit with there is a need for publicly accessible information so that when there are issues with networks, which is the original reason for these WHOIS databases to exist, you can actually contact the contacts for those networks and troubleshoot or work out what the situation is there.

PAUL WILSON: The people that are registered in the WHOIS databases are generally contact people for organizations who are using IP addresses in their business capacity so there are very few if any individuals who are registered in the database in their personal capacity. They're there for the purpose of administration of the IP addresses that they've received and it's quite explicit.

And I know it's very explicit in APNIC's case that if you receive IP addresses from APNIC, then you are obliged to put this information in the database and that even though we're not talking about European citizens here, it's still a basic privacy principle that the information is stored for and used for purposes which are declared, and we're very careful about that as part of general privacy principles, which also satisfies GDPR.

KEVIN BLUMBERG: Kevin Blumberg, I'm going to speak because I run an internet service provider in Canada. And we are under something called PIPEDA, which is the Privacy Act for Canada. Privacy is not new I think to many countries, many countries have different levels of privacy. In the region I have needed to use one of the policies which was related to private residential customers and that redacted the information for those, but we knew who the end contact was, it would just show city, I think, and not even postal code. So we've had to deal with privacy but ultimately, at least in Canada, business and operational need allows for this data to be published. It is not personal data in that view.

But ultimately, it's following the rules and regulations. What happens in Europe happens here as well differently, similarly at the same time. But this is predominantly business data where on the other side of the fence with domains you have no idea whether it's a business that's registering something, or an end user. Predominantly IP addresses are going to be an organization of some kind.

PAUL WILSON: Was there anything else? Anything else from RIR colleagues in the room? Well, I can say it's been a really interactive session, more so than in the past many times, so thanks very much to those of you who raised questions and I hope this has been useful. Thank you. Thanks a lot.

[END OF TRANSCRIPTION]