HYDERABAD – How It Works: Internet Networking
Thursday, November 03, 2016 – 11:00 to 12:30 IST
ICANN57 | Hyderabad, India

STEVE CONTE: Alright, we're going to start the next session. This is on naming, addressing, and routing. I appreciate you all coming back or coming into this session. I am going to introduce Alain Durand, who works in our research arm of ICANN's office of the CTO. Just another quick pitch: he has been working on a project called ITHI, Internet Technology Health Indicators. Do I have that acronym right?

ALAIN DURAND: Identifier Technology Health Indicators.

STEVE CONTE: Thank you. It's quite an interesting topic, and you have a session on that on Monday.

ALAIN DURAND: On Monday afternoon.

STEVE CONTE: Okay, please check the schedule. It is something worth noting, and it's something of interest. If you guys have any interest in following this health indicators project that Alain is working on, I urge you to go to his Monday session as well, but that’s not why you're here right now.

Alain is going to talk about naming, addressing, and routing. With that, I'm going to pass it over to Alain, and I'm am going to be your go-to guy [for] questions.

ALAIN DURAND: Thank you very much, Steve. Just a word about the Monday session: it's about defining health for things that ICANN does, so I can help coordinate the unique identifiers of the Internet. We would define health as being free from disease, and we talk about a number of diseases that we can have on the Internet. It could be of interest. I encourage you to come and participate.

But the topic today is naming, addressing, and routing. As I mentioned earlier, when Steve was queuing for this presentation, sometimes people confuse naming, addressing, and routing as all the same thing, and they are very different concepts. We’ll try to go through that. So please, next slide. Thank you.

That’s the order of the presentation. We will talk about the OSI model, and something that dates back from the '70s and '80s. It's old, but it's good to conceptualize networks, and the OSI model is on seven layers.

You may have heard about this, the seven layers of the OSI model, and then we will go through each of them from one through nine, and we'll see why we added two layers there. Then, we'll talk about naming, addressing, and routing. Next slide. Thank you.

I call this networking by numbers. It's like painting by numbers, just follow the different numbers and try to understand what is happening at each of the different layers. Next, please.

Alright, so when you build network, it's started with a real, physical layer. There are two types of network: there are wired networks and wireless networks, and they have very different characteristics.

Wired networks, we talk about copper lines or fiber lines. Wireless, we talk about antennas. A wired network, when you want to lay out fiber, you need to dig a trench. In some places it could be fairly expensive to do that, especially if you have to put a lot of it.

That's why people like wireless solution because all you have to do is to put an antenna, and then you can connect a lot of people very quickly. However, the bandwidth you can get on fiber is much higher than anything you could get on wireless.

If you want to put a lot of people, well, you need fiber. Wireless will not do it. We talked about the cost of putting fiber in the ground, but when you have antennas, sometimes people don't like antennas.

If you are in an upscale neighborhood, people don't like to have an antenna in there, or maybe you have old buildings that are buildings that have some historical characteristic to them and you don't want to put an antenna there. You may not want to put an antenna on top of the Taj Mahal, for example, or on top of the White House. It's just not something that you do.

So different characteristics, and what we see is some kind of a convergence between the two where people lay out wired network when they can and wireless network to essentially go to the last mile and where wireless antennas will be connected with fiber to the rest of the network.

When we talk about really physical networking, those are the two technologies that we need to keep in mind. Next slide, please.

For now, I'm just going to talk mostly about fibers, but most of this applies to wireless also. I want to communicate over a fiber. How do I do that? If it's a copper line, I send electricity over the copper. Fiber, I have to send light. Initially, what we were doing was sending white light over the fiber, and we received the light on the other side.

That's nice and well, but if I want to have multiple communication over the same fiber or multiple people, multiple tenants sharing the same fiber, that doesn't work. So what we do is instead of sending white light, we send colored light at a specific frequency. Kathy, what's your favorite color?

UNIDENTIFIED FEMALE: Blue.

ALAIN DURAND: Blue, thank you. If I want to send a communication in between two points that Kathy [can talk], maybe I'm going to send blue. If I want to send traffic for another subscriber, maybe you – what's your favorite color?

UNIDENTIFIED MALE: Purple.

ALAIN DURAND: Purple. So we're going to send your traffic over a purple wavelength. So blue and purple, they are somewhat close together. I have to make sure that they're not too close. If not, the electronics on the other side may not be able to see the difference. Usually, we can send quite a lot of different communication for different people on the same fiber using those different wavelengths.

Now, if I want to send data from, let's say, here in Hyderabad to my office back in D.C. over fiber, I'm not going to get a trench all the way from here to D.C. There's no way I can do that.

But maybe, we will dig up a trench from the hotel to a point nearby, and then I can reconnect to another fiber that goes from a central office to maybe a main office and maybe catch up a submarine cable somewhere. So now we have bits of different fibers that we need to patch together to create a link. How do we do that?

Well, that's relatively simple. There are some switches that simply take a lambda – meaning a wavelength – on one particular fiber and translate it into a lambda – a different wavelength – on the different fibers and create a path. Instead of having just one fiber, what we have is what we call a fiber path from one place to the next.

The speed on those things could be fairly high. We're talking about multiples of gigabits, like 10 gigabits, 10 gig E, 25 gigabit per second, 40 gigabit, 100 gigabit, and then in some cases up to terabits per second now. Next slide, please.

But all of this creates a path which, at the end of the day, is a point-to-point link between two places. When we do networking, we rather like to group things under some kind of administrative domain and not have simply a view of the entire planet is just one gigantic network. It just doesn’t scale.

In the good old days of Ethernet when people were sending messages, when you send a broadcast message, it goes out from one place and goes to all the other nodes, and that defines what's called a broadcast domain.

Now, if you have ten nodes, that’s fine. If you have 15 nodes, that fine. If you have thousands of nodes, that doesn’t work very well. You can kill the network just by sending a broadcast message. So we try to group networks and segment them so that what's happening there has no impact on what's happening there if there's a problem. That’s how segmentation started.

Now, this notion got expanded into the world now, and [inaudible] administrative domain, so I can do whatever I want in one part of a network, and somebody else can do whatever it wants in another part of the network.

That’s when you need to remember that Internet is spelled with upper case I, capital I. It's not just a fancy thing. It means that this is a network of networks. What you do in no network and what you do in your network over there may be different things, but we will use some common protocol so that we can have communication between those networks, even though they may be controlled by different administrative entities. That’s why there's an upper-case I.

So this is layer three. Layer three is this IP protocol or Internet protocol. Above that, when we have this network, we want to establish communications. Next slide, please.

Level four, layer four, transport. There are two protocols in transport that are widely used, and a bunch of others, but the most widely used ones are UDP – for User Datagram Protocol – and TCP – for Transmission Control Protocol.

UDP is fairly simple. You just send packets from one place to the next. TCP is meant to be a reliable protocol. I want to have a communication maybe with somebody that’s far away from me, and I want to make sure that the communication works, that whatever I send is being received.

I'm going to send data to Kathy, and I'm going to pretend that I'm far away. When I send data, I chunk the data into small quantities that we call packets. On the Internet, most packets are about 1,500 bytes. It's not a lot.

ATM days, for those of you who remember those, there were even smaller cells. They were 53 bytes. On the Internet size, we use 1,500 most of the time. I'm going to send a bit of data. Here's my first packet, Kathy.

UNIDENTIFIED FEMALE: Ack.

ALAIN DURAND: She sent me back a message, "Ack" for *acknowledge*, saying that she has received it. Now I know I can send another one. Here's packet number two.

UNIDENTIFIED FEMALE: Ack.

ALAIN DURAND: She got my second packet. I'm going to keep going like that. Packet number three. Nothing. My packet got lost. I have no idea where, but it got lost, so I know that I have to retransmit it again. Kathy, here's packet number three.

UNIDENTIFIED FEMALE: Ack.

ALAIN DURAND: Now I know that she has received it, and I can move on to packet number four. That's how the TCP protocol works to detect that packets get dropped. Why do they get dropped? You never know. Sometimes, it's because there's an electronic glitch, somehow it gets dropped, or it goes into a router and there's some congestion.

When there's congestion, a router has to drop some packets because there are just too many that are coming in. Maybe the traffic was dropped by one of those routers. Maybe there were wireless issues and things were not worked, or some fiber was cut or something like that.

It doesn't matter what the reason is. The point is that she didn't receive my third packet, so I had to transmit it again. This is all well and good when communication is established, so the first question is, how do we establish the communication?

I'm going to start a new communication, and this is how it works: "Hello, Kathy. Do you read me?"

UNIDENTIFIED FEMALE: "Yes, I hear you."

ALAIN DURAND: Now I've established that she hears me, so I can send things to her and she can reply back to me, but she doesn't know if I hear her. I have to do one more thing: I have to say, "Hi, Kathy. I do hear you too." And when we have done that, we call that a three-way exchange.

We have established that I can talk to her, she can hear me, she can talk to me, and I can hear her back. Then, we can start the normal communication. Thank you, Kathy.

When we' are done, we'll say, "Kathy, we're done," and Kathy will simply respond, "I'm done."

UNIDENTIFIED FEMALE: “I'm done.”

ALAIN DURAND: Thank you. Now, we can all hang up. Next slide, please. Thank you.

On top of those transport protocol that were made for data transfer, bulk data transfer, like a big file to transfer from one place to the next, in the last 15-20 years, we have had real-time communications happening on the Internet.

Back in the 1980s-1990s, we had the Internet using phone infrastructure, and now we have the phone using Internet infrastructure with Voice Over IP and streaming of videos and things like that.

They use some other protocols called Real Time Streaming Protocol, which essentially is opening up a session, closing a session, sending bitrate information and all of that.

Next, layer six. In OSI model [where I say] again seven layers from one to seven, we need to format the data so that when I send a bunch of data to Kathy, she knows what it is. It's not just zeroes and ones. She knows that there's some information that starts there and finishes there.

At the beginning, it was just anything goes. I would send something, and there's some kind of prior agreement and she knows what I'm going to say. Most of the time, it doesn't really work because I may be confusing to what I say.

So there are different formats that have been standardized. ASN.1 was standardized a while ago, and it was about binary communication. It's to say from bit number 17 to bit number 23, we are going to have that field. From bit number 35 to bit number 110 is another field. It's all binary.

It's nice because it's very compact. However, when you see those bits, it's really hard to read. That makes programming much harder. Since then, there are other formats that have been described that are more user friendly. XML or markup language is what is used on the web. Also, you see tags on the web to say that, for example, bold, /bold is a markup language, and JSON, which is the more recent one.

This is an example of a JSON format. This is describing a menu, like when you have a Windows system with a menu. You have a couple of things inside of a menu. You have a popup. You have different values and things. This is a human readable format.

Most of the recent protocols defined at IETF – which is the Internet Engineering Taskforce that does define all those lower layer protocols – use formats like this now. Next.

But all of this is essentially useless until you have users. At the end of the day, those are the users that matter, and what they want is application. What they want is to watch videos, they want to watch some music videos, they want to play games, they want to chat with their buddies. That's what really matters.

We build the Internet for that. All the lower layer protocols that are interesting for engineers to understand how it works, but at the end of the day, this is really what matters. When you have a kid on an iPad who can surf the web and chat with their friends on the other side of the country or in another country and watch the same videos, that's a success, and that's where we are now.

Those are the seven layers of the OSI model, and then we have two more layers. Next slide.

Financial: none of this works unless there's money to make this happen, and money only is there if there's a good business plan behind it. If you are a service provider, an Internet service provider or application service provider, a registry or a registrar, whatever it is, unless you have a sound financial model, none of this will actually happen. You always have to remember that. Next slide.

And, of course, the layer on top of that is the political layer, and this is the ICANN meeting this week, so that's why we are here today. This is my nine layers of the OSI seven-layer model. Now that we understand the basic technology of networking, let's try to see how it works on the Internet. Next slide, please.

This all started a little while ago, when I was traveling and I was not home, far away, and I had a tooth problem. One of my teeth in the back was really hurting, and I was in a foreign country. I knew very little people there. Remember what Steve said earlier? It's all about who you know. Same thing here. I needed a dentist. Next slide.

I know somebody here. Kathy, I have this raging toothache. Who is your dentist?

UNIDENTIFIED FEMALE: My dentist's name is Dr. Walker, Johnny Walker.

ALAIN DURAND: Thank you. I need this Dr. Johnny Walker. I need it badly. That's the name of a dentist, right? So what's a name? Well, back to the dictionary: a name is a word or a set of words by which a person, an animal, a place, or a thing is known, addressed, or referred to. For example, my name is Parsons in that case. My name is Walker, Johnny Walker.

If I know a name, I know who you are. There are a few things that I will ask that you remember from this session. The first one is this one: a name is who you are. I know now that the dentist I need to go see is Dr. Johnny Walker. Next slide, please. Thank you.

When I have the name of somebody, I can do two things: I can go talk to Dr. Johnny Walker about my toothache, but I can also talk about Dr. Johnny Walker. I can go and talk to Steve. Steve, Kathy has told me about this Dr. Johnny Walker. Do you know him?

STEVE CONTE: Oh, do I know him.

ALAIN DURAND: So we can have a conversation about Dr. Johnny Walker. [So there are either] two functions here: one is talking to someone, and the other one is talking about someone.

And we can pass this information around. Once information is passed to somebody, it's called a referral. So now we have the name of this dentist, Dr. Johnny Walker. Next slide, please.

Names have a scope. My name is "Alan." In French, it's Alain, but really, for most people, I ask to call me "Alan." In my family, I'm the only Alan, in my direct family. But when I was in elementary school, Alan was a fairly popular name. I remember there were like four of us. The teacher would say, "Alan, could you please explain to me this?" And the four of us were looking at each other, like, "Which one?"

And of course, the teacher would get mad because nobody would answer, or even worse, the wrong Alan would answer. So she had to do something like, "Alan Durand, please tell me what is 2+2." Then I could answer, hopefully right.

So this idea of adding two parts – the first name and the last name – to really qualify the name so that there is no ambiguity is something quite important. Next slide, please.

Okay, but I know about this Dr. Johnny Walker, a good dentist. I need some of his prescription here, but where is he? I have a name. I don't know where he is in this city. Is he here, is he on the other side of the city? In the north or the south? I don't know.

So, I have to go back to Kathy. Kathy, I've checked on this Dr. Johnny Walker. He seems to be a good guy. I just don't know where he is. Do you have his address?

UNIDENTIFIED FEMALE: His address is 125 Root Canal Drive.

ALAIN DURAND: Thank you. Root Canal Drive. That's probably very appropriate for a dentist. What we just did is called address resolution. If you follow the tutorial from Steve earlier on, that's what the DNS does. It takes a name and turns it into an IP address. Here, we took a name and took it into an address. That's the exact same process.

There's a directory. Kathy has this directory, and she doesn't use those paper cards anymore. She has a computer to do that, but that's the same thing. It's the place where you have names, and that's the entry, Dr. Johnny Walker, and then you have a bunch of information like contacts, telephone number, and all of those things.

So if you wonder what was [our archetype] in the presentation from Steve earlier on, those are exactly that: one is the contact, one is the telephone, one is a bunch of other notes. That's why you have multiple records in there. So now I have this address: 125 Root Canal Road. Next slide, please.

In the DNS, there are a bunch of issues that came up in the last few years. DNS is a fairly old technology from the '80s. We can mention internationalization because not everything is English. Even I'm not a native English speaker. I'm French, and we have different characters in French than in English. If you're a Chinese speaking person, you have a very different character set. If you're in India, you have very different character sets too, so this internationalization is really important.

The second part is authentication, and that's where DNSSEC comes to play. To Steve's point, this is really about authentication. The way it works is that instead of securing the channel from the source to you about the data, it's about securing the data itself. So even if the channel is compromised, you can still look at the data and see, is it the real data, or is it not the real data? And if it is not the real data, then it is up to you to decide what you do. You can dump it, or you can still try to connect.

Another important issue is the expansion of the root zone. In the beginning, you had things like country code and .com, .org, .edu for universities, .net, and that was about it. Now, you have several hundred top level domains to choose from, and they're not all equivalent.

You can register in almost any of them you like, not all of them. There are some cosmetic issues in the sense that some are nicer than others. But it's a choice about which one you use. I'm not going to expand much on that. This will be the topic for another tutorial by itself. Next.

Now, again, my tooth is really hurting. I know that I need a prescription from this Dr. Johnny Walker, and I know that he's on 125 Root Canal Road. Let's talk a bit more about this address thing. Next slide, please.

What's an address? Again, according to the dictionary, this is a particular place where someone lives or an organization is situated. Essentially, if I know your address, I know where you are. So remember, first thing, if I know your name, I know who you are. If I know your address, I know where you are. Those two things to remember. Next.

Let's make a little detour here about addresses. Sometimes, you will find structure for addresses. Addresses are not just for Internet; you find addresses everywhere. For example, I live in

Washington, D.C., and the most famous address in Washington, D.C., where lots of people want to live is on 1600 Pennsylvania Avenue, Northwest, Washington, D.C., 20500-0003, USA. And this is the small white house that is there on the picture, where everybody wants to live.

When you look at this address here, there is an hierarchical structure, and it reads from the end to the bottom, from the right to the left. It's in the USA. There is a ZIP code. It's in D.C., the District of Columbia, in the city of Washington. Northwest means the northwest quadrant of the city – there's northwest, southwest, northeast, southeast – Pennsylvania Avenue, and 1600. That's how you read this address, from the right to the left.

Not all addressing structures have the same hierarchy, and they're not always geographically organized. For example, in the United States, you have toll free phone numbers you can call for free, and those are 1-800 numbers. We don't know if the phone number is in Washington, in Texas, in California, or maybe forwarded to India or to China or wherever in the world. We don't know that, it's just flat.

If you look at cell phone numbers, same thing in the U.S. You may have a cell phone that is in a specific area code that is moving, roaming into a different area code, and still have the same phone number. So, by simply looking at the area code, you have no idea where the phone is.

Or IP addresses for what matters. IP addresses are not distributed on geographical boundaries, so simply by looking at the IP address, you don’t know where the IP address is. You can look it up in a database that has geolocation information to say, "Oh, this is an address that is here in Hyderabad in India," but somebody had to build that database. There's nothing in the format of the IP address that tells you that the IP address is here, and maybe tomorrow, it will be moved somewhere else. Next slide, please.

Same thing as with names, there's a scope associated to an address. When I'm in D.C. and I'm asking somebody where the White House is, they're simply going to give me that part of the information, 1600 Pennsylvania Avenue, maybe Northwest, but they're not going to give me the full address. And that’s enough, because I'm in Washington already, I know where it is.

If you ask somebody where is Paris, well, if you live in Europe, most probably people will say, "Oh, Paris, that’s in France." If you're in the United States, there are 29 cities that are named Paris. I live actually just outside of Washington, D.C., in Virginia. There's one small city – maybe 100 people live there – that is named Paris. Sometimes, my kids come to visit and say a joke, "Let's go to Paris. We'll drive there, it's easy." So half-hour drive to get there.

The point is, same thing as with names, you sometimes need to qualify the address to make sure that this address is globally unique. If not, there could be ambiguities. Next.

Also, similar to names, I can use an address directly, for example to send a post card or to go somewhere. I can also pass it to somebody as a reference or talk about it. Steve, I've heard that the dentist lives in 125 Root Canal Road. Is it a safe neighborhood?

STEVE CONTE: Yes, it's safe enough, if you go before 8:00 PM. Alright, so we can have a conversation about this address. The same way that you have reputation database that says, "Oh, this address has been used to send spam." So that's another usage of an address that has nothing to do with actually sending packets to that address or going to visit it. We can have a conversation about it.

Going back, I need to go to 125 Root Canal Road. Next slide, please. When I have an address, I know where it is, but I have no idea how to get there. If I want to send a post cart to the White House at 1600 Pennsylvania Avenue from here, from this hotel here or conference center, if I write the address down on a post card, drop it in the mailbox, a few days later it will arrive there.

It's not magic. Why does it work? It works because there's a postal system that works here in India that has some collaboration with the postal system in the U.S. and that knows how to forward this. In India, the fact that it is an address in Washington, D.C., District of Columbia or in Texas is completely irrelevant. All they need to know is that it's for America, for the U.S. They send this over to the U.S. and ask the U.S. to deal with that.

This system is a collaboration system that's critical for the post office. It's also the same, similar system that we have on the Internet that's based on collaboration that's really critical. So, we're going now to explore how to build a similar system for the internet. Next.

But before we go there, I need to remind you about those Internet addresses. There are two protocols that now coexist on the Internet. The mostly used one is IP version 4. It was defined in 1981 and still dominant on the Internet. The new one, the new kid on the block is IP version 6. It was defined in '94, redefined in '98, and that's what we're starting to use now.

If you're asking about IP version 5, what happened in between 4 and 6, IP version 5 was an experimental protocol to send real time traffic, so it was never really used widely. But when the document went through IETF and somebody had to assign a version in the table where we keep all the assigned numbers, we realized, "Oh, 5 has already been registered, so we have to go to 6." That’s why there's IPv6. Nobody knows what happened to IPv1, 2, and 3.

So those protocols are fundamentally the same, except for the length of the IP address. IPv4 is 32-bit, and that makes for 3.2 billion usable addresses. Actually, that number is not correct. It's 3.7 billion usable addresses.

IPv6 is 128-bit, so going from 32-bit to 128-bit is not multiplying the address space by four. It's actually squaring and squaring again. So the number is what you have at the bottom of the screen. I have no idea how to pronounce it, but the point is it was meant to be large enough so that we will not have to worry for the next 50 years or 100 years. Next slide, please.

You have heard of IPv4 exhaustion. It started to happen four or five years ago, and now we really are there. This has nothing to do with the IP addresses getting tired. An IP address is a number. It doesn’t get tired. When we say exhaustion, it doesn’t mean that we cannot use this thing anymore or that it's like a use by date on a product and it gets spoiled after a while. That’s not the case.

But the thing is, most, if not almost all of the addresses have been allocated, more or less efficiently. But the Internet is still growing. It used to be that people go to their regional Internet registry to get more addresses. Now, they go there, and in most cases, the pool is empty or almost empty. What's happening next? Next slide.

IPv6 was supposed to be the solution to this problem; however, it doesn't really work as planned. The issue here is IPv4 and IPv6 are not directly compatible, so you cannot have somebody who speaks IPv4 and somebody who speaks IPv6 [talk] together.

It's like I speak French and English, but I don't speak Chinese. So if I speak to a Chinese person that doesn't speak English or French, we can't talk together. Exactly the same problem. If both of us speak English, that's fine. It's a technical limitation. That's why we have this problem. Next slide, please.

You cannot simply say as a service provider, "I'm going to deploy IPv6 and forget about IPv4," because all your customers will not be able to use it. Maybe they will have computers at home that are a little bit older and they cannot use IPv6, so if you only offer them IPv6 service, they're unable to use it.

Or they want to go somewhere on the Internet that hasn't yet adopted IPv6. Of course, they cannot reach it. So, it's kind of a chicken and egg problem, and essentially, this is a game where you have a last mover advantage. Most of the successful technologies have a first mover advantage, meaning if you're the first to adopt it, you have an advantage over your competitor.

In IPv6, it’s kind of the opposite. If you do it first, yes, it's nice, but you cannot really talk to many people. If you do it last, yes, you will be able to talk to everybody. So that’s why it's a little slow to be deployed, and it's still happening. I have a study that I presented at a RIPE meeting last week that shows how IPv6 is deployed in different countries in the world.

So this is really happening, but not in a uniform way. Some of the countries that you would think would be leading the charge are not, and some of the countries that we'll never think about are actually quite far along on this path, for different reasons. But the bottom line is that IPv4 and IPv6 – both of them – are going to be coexisting for a long while. And by a long while, I'm not talking about six months or two years. I'm talking about five, ten, maybe more, maybe 20 years. Next slide.

If you are a service provider or an enterprise and you need to get on the Internet and you need more addresses, you can get those v6 addresses, but that will not help you to talk to your v4 customers. So you need more addresses. How do you do that?

Well, you can go to APNIC if you are in this region of the world, and if you are a new entrant, they will give you a /22, which is, if

I'm correct, about 4,000 addresses. Well, a couple thousand addresses anyway. That's not enough to run a full network, but maybe enough to start. Or maybe you can use this block using NAT or [carrier-grade] NAT and share those addresses among multiple customers. That's one way of doing it.

Another way of doing it is to get addresses from somebody else. Maybe there's somebody who has a large quantity of addresses but doesn't use them and is willing to transfer to you some of them in compensation for some money. Essentially, you're buying addresses.

Terms and conditions may vary. You can find addresses in the same region with interregional transfer, or you can find addresses from another region. Historically, many of the addresses were allocated in the [inaudible] back in the early '80s, early '90s. So, that's where you will find quite a lot of what we call legacy addresses, and those are the ones that are being transferred in fairly large quantities. Next slide, please.

Let's look at those statistics next. This graph is a little bit dated. It goes back to May 2015, but what you see is every month, more and more addresses are being transferred. Next slide, please.

You can look at this in terms of number of blocks or in terms of addresses. It paints slightly different pictures, but the bottom line is there is currently a global market for IP addresses, and the going price for an address is about $10 if you buy it in very large quantities, like millions of addresses at a time. If you buy it in small quantities, the price can go to $15, maybe $20 in some cases, depending on the quality of the addresses.

Where are those numbers going? Is the price going up or down? When transfer was introduced, for a while prices went down because lots of people eager to sell the blocks. Buyers were not quite confident yet into the process, so there was too much demand so price went down. But now, people get more confident in the system and there's not always enough supply, so prices are going back up and we see them trending up. Maybe in a year or two, we don’t know where it will end, but now, prices are going up.

But still, this is for a large service provider or a large enterprise, the way today to get IP addresses is to buy them in the market. This market, as I mentioned, is worldwide, so it's relatively easy to move addresses around. Next slide.

Alright, enough about addresses now. I need to go see this dentist, because I have to remember, the reason I have to go there is tooth is hurting. I need to go to Root Canal Road. Okay, I'm here in this conference center, and I want to go to Root Canal Road. How do I do that?

Well, if you look on the roads, you will see signs on the road that says, "Want to go to this direction? Take this road. This road goes to another direction." Somebody has created this infrastructure with directions, with signs, that tells me if I want to go there, this is how I should go. This has been done way before I tried to drive there.

On the Internet, that's the same thing. we need to establish those routes before I can even send traffic. This is a job that's done by my service provider, not by me. I have to rely on them. Again, this is about cooperation. If there was not this cooperation between service providers, nothing would work. Next slide.

The route – or route, depending on how you pronounce it in English – is a way or course taken from getting from a starting point to a destination point. If you have a route, I know where to go. If I have your name, I know who you are. If I have your address, I know where you are. And if I have a route for you, I know where to go to get to you. Those are the three things I would like you to remember today. Next slide.

Let's see. If this is a diagram of my network, I'm near the source on the left-hand side, and I want to go to the destination, to Root Canal Road, to see this Dr. Johnny Walker. So we have to build this map.

The way the service provider builds this map is through route announcements that are going to flow from the destination to the source. The service provider that is the closest to the destination, actually the service provider that is servicing the destination, is going to say, "I know where Root Canal Road is," and is going to announce this to all his neighbors that he has a connection with, a peering agreement with.

If you are the next one down, you will hear – no, stay on the same slide, please. Thank you. If you are the next service provider down, you will hear, "Oh, that service provider up there knows how to get to Root Canal Road." If you are the second down, the one at the bottom of the diagram here, you will hear, "Oh, this service provider on my right knows about another service provider that knows how to get there."

This is all about "I know a guy who knows how to get there, or I know a guy who knows a guy [who knows] how to get there." This is a chain that is built like this. Again, this is a cooperative system. If it was not for this cooperation between service providers that are going to do a good job at announcing this, there would be no Internet. Next slide.

When I send packet, I send packet from the source to the destination, and what I do is simply follow the directions. Somebody will say, "Oh, I know a guy who knows a guy who knows a guy who knows how to go to Root Canal Road," so I go there. Then it will send the packet to the next stop and to the next stop and to the next stop until we arrive to the place that says, "Oh, yes, I have a direct link there."

This is called next stop routing, or hot potato routing sometimes, when you try to pass it to the closest destination. It's all about relying on somebody, not only to have built a map properly, but also to pass the packet to the right place and to do this job correctly. Those are the two facets of this cooperation that is really important. Next, please.

UNIDENTIFIED FEMALE: Alain, we have a question.

ALAIN DURAND: Yes.

UNIDENTIFIED FEMALE: From Miguel Munoz on the chat room. "Where does the money go when an ISP buys IP addresses? To the RIR, to ICANN, or to an intermediary?"

ALAIN DURAND: It certainly doesn't go in my pocket. That, I know. When a party A buys addresses from a party B, that monetary aspect is in a private contract between party A and party B. It doesn't go to the IR, it doesn't go to ICANN, it doesn't go to another service provider. It's simply between party A and party B.

Now, sometimes, they use a broker, somebody who's going to help the transaction, and that broker may take a fee, a broker fee. It might be a couple percent, it might be up to ten percent, and they will take some of the money from that transaction.

Once that transaction is actually made, the addresses have to be registered on the name of party B. So you have to go to the IR and say, "Please update the record now that block address X now is under the control of party B." You need to pay a fee to the IR to actually do that, and the fees are usually nominal fees.

On top of that, you will have to pay membership fees, and you have to pay every year some maintenance fees, as you would do for any other addresses. But to specifically answer the question, money doesn't go to ICANN, and it's only the update of the record fees that goes to the IR, nothing to do with the actual value. And the IRs do not know how much you have been paying. This is not information they have.

When I was talking about $10 or more sometimes, this is an information that has been disclosed by the brokers because they know about those transactions. The IRs do not know about this. So, back to routing, next slide please.

I was talking about this cooperation, but what happens if there's a bad guy in between? There's somebody who's trying to send me not to Dr. Johnny Walker, but he wants to send me to Dr. Jameson. There have been examples in the past where this has happened, where traffic to Google was redirected to the wrong place.

If one of those actors was receiving this announcement and forwarding to his neighbors, he’s not playing the game. Bad things happen when we have bad actors. Typically, what will happen is a bad guy will say, "I'm here, and I know how to talk to Dr. Johnny Walker," even though he may not have this link, and if I'm gullible, I may believe it. If he's close to me and he talks loud, I may believe it.

That’s a bit of a problem. To avoid this problem, there's a solution called RPKI, it's for Routing Public Key Infrastructure. What we're going to do is not unlike what is done in DNSSEC, is to validate the data, authenticate the data.

When somebody is going to say, "I'm Dr. Johnny Walker," you're going to sign this with a certificate. If you make an announcement saying that "I am a service provider for Johnny Walker," but that certificate doesn’t compute, then you're simply going to drop the announcement.

That's a very interesting solution to validate the origin of the announcement, but it has a number of issues. It only validates the origin. It doesn't validate all the elements of the path, and for a number of service providers, this is not enough. They would like to have a full validation.

Another issue is about the root of that particular certificate tree. Is it centralized, or is it decentralized? There have been a lot of discussions, and disagreement about how the system should work. So that is somewhat preventing a full uptake of this, but there's still lots of effort happening there.

If you're interested in this, all the different IRs have tutorials and a lot of information on this particular system. Next.

Now, if you recap: I had a toothache, I needed to find a dentist. There are three steps: first one is asking my dear friend Kathy, "Who is your dentist?" And your dentist is...?

UNIDENTIFIED FEMALE: Dr. Johnny Walker.

ALAIN DURAND: I did a little bit of background check to make sure that Dr. Johnny Walker was a good guy. Right, Steve?

STEVE CONTE: Yes, he's the best.

ALAIN DURAND: Thank you. Now, I have a name. Remember, I have a name, I know who you are. Next step is I need a directory that is going to tell me where Dr. Johnny Walker is. Dr. Johnny Walker is...?

UNIDENTIFIED FEMALE: 125 Root Canal Drive.

ALAIN DURAND: Thank you very much. Now, I have this address. I have the address, I know where you are. Now, in order to get there, I need a route. I need to find directions on how to get there. This is what routing is all about. So you have naming, addressing, and routing. Those are three functions that are critical on the Internet, and they're all different. Next slide.

Now, I can have my tooth being taken care of by Dr. Johnny Walker, and I think that’s the end of this presentation. If you have any questions, I would be happy to answer them.

STEVE CONTE: Any questions from the floor? Here we go, a question.

UNIDENTIFIED MALE: We were talking about IP version 4 and version 6. Version 6 again talks about the Internet of Things, IoT. In that, we have a very big number of telecom mobile numbers [as well there]. Is there any way to translate them and talk among each other? Is it possible to?

ALAIN DURAND: I'm not sure I understand the question. Let me try to rephrase it.

UNIDENTIFIED MALE: In this addressing and numbering, we just are referring only computers, handheld devices, smartphone, which are directly connected with the Internet. Is there any way to include telephone numbers also? Telecom.

ALAIN DURAND: Let me try to rephrase the question and see if I get it right. Traditionally, what we have connected to the Internet were computers, right? Even cell phones. Now, we're talking about all kinds of Internet of Things objects, like a webcam, it could be a fridge, it could be a lightbulb, and can we do that? Is that your question?

UNIDENTIFIED MALE: Yes.

ALAIN DURAND: There are different opinions that exist on this topic. Some people will say that's just too many objects. We're talking about potentially billions of objects, and there's no way we can do that in IPv4, so we need to use IPv6. However, if I go to the electronic store today and I want to buy an Internet of Things gadget, most of the time, it will be an IPv4-only device, so that's a bit of a problem.

Now, that being said, the lightbulb in my house needs to communicate with maybe the controller of a lightbulb in my house. Maybe I may have a secure communication between my phone and the controller. But it's unclear that the lightbulb in my house needs to communicate directly with the lightbulb in Steve's house. Very unlikely.

So a number of people have said, "Well, you can use your private address space network, [then] you can have 16 million devices. You can have 16 million lightbulbs in your house if you want to, and that's just fine. So we can keep using IPv4 to address the lightbulbs.

We had the same discussion about 15 years ago with telephones. People were saying, "Oh, we have GSM phones. We're going to have a billion phones. We need one IP address per telephone, and they're simply not enough." Well, we found NAT, carrier-grade NAT, and we find ways to sometimes have 5,000 people be on a single IP address and it works.

Now, we have reduced those ratio to maybe 100. But if you buy a /16, which is 65,000 IP addresses and you use a carrier-grade NAT where you can have 100 customers per IP address, your 65,000 becomes 6 million. And all of a sudden, 6 million that you paid essentially $500,000 for, for 6 million customers, it becomes much more reasonable.

That discussion was exactly the same argument: we have too many devices. In the end, we find ways to get them with IPv4, and I think it will be the same with IoT.

UNIDENTIFIED MALE: A small correction because the answer is going in another direction. My question was not about the number of things in IP version 4 domain, version 6 domain, and the number of people who are actually using so many mobile devices in PSTN and telecom domain. Both are there. [inaudible] they are developing their own [inaudible] there. My specific point was, because [inaudible] are there, IP version 4, version 6 are actually in the same domain. They talk among each other seamlessly. How about including the telecom domain?

ALAIN DURAND: Oh.

UNIDENTIFIED MALE: Specifically, my [inaudible] E.164 protocol [is there]. It was specifically designed for SAP, H.323, VoIP services. Is there any development going in generating that sort of solution? All protocols, because we have WebRTC, our latest development. So is it possible to address mobile number directly from an IP address?

ALAIN DURAND: Yes, it is, and you don't necessarily need E.164 to do that. if I want to connect to a mobile device that is behind a NAT, I may not be able to communicate to it directly, but there are some new protocols. There's one in particular called PCP in the IETF that enables to punch a hole in the NAT and resolve a port.

So, when you're on a device, let's say your phone behind a NAT, carrier-grade NAT, you will say, "I want to be reachable from the outside so that somebody can connect to me," and you reserve a port. Then you can publish in another directory – sometimes a peer to peer directory – where you are. By sending communication to that particular TCP port number, you will reach directly your mobile. So I think this is a way to address what you're trying to do here.

UNIDENTIFIED MALE: What is the status in that direction?

ALAIN DURAND: The PCP protocol has been standardized in the IETF, and there's a number of implementations that exist already for that.

STEVE CONTE: We have a comment from Mike Palage on the web and a question to follow.

ALAIN DURAND: Question on the web?

UNIDENTIFIED FEMALE: We have a couple of questions, but I'll start with a comment first from Michael Palage: "To the question being asked from the floor, it is interesting to look at the efforts of the RIRs and the domain name registration authorities to implement RDS. This allows registries to augment additional data fields associated with these unique identifiers."

The first question is from Ravi Bandi. "Why has it take so long for the IPv6 to be implemented?"

ALAIN DURAND: I think I touched on this earlier on. I've been working on IPv6 since 1993. Back then, we thought, "Yes, there will be a problem with IPv4 in the future, and we need to prepare for something." Until 1996, we thought the number of devices on the Internet is doubling every few months, so the legacy base doesn't matter. As soon as we will put the new code for IPv6 in all the operating systems, like Microsoft Windows, Sun Solaris, IBM OS 400, we thought that's it, we're done.

I think we were naïve, to be honest, and it didn't happen that way. Why is that? Because we were early, and there was no incentive for anybody to actually deploy this stuff. Even in the 2000s when all the routers, all the operating systems of our computers supported it, nobody had any incentive to go and deploy it.

It's not because it's available that people go and deploy it. Why? Because they still had a lot of IPv4 addresses. So why creating more headache? If you have to manage two protocols, it's twice the number of headaches, twice the number of potential security vulnerabilities. Why do that?

That's back to my point about OSI layer eight, financial. If you don't have a financial incentive, it doesn't happen, and that's exactly what happened to IPv6. Nobody paid much attention to it until the IR sent letters to all their members saying, "Oh, we're going to run out of IPv4. Do something." Then we saw a regain of interest in IPv6.

But again, as I mentioned on the previous slide, if you deploy IPv6 and only IPv6, you are not in a good position, because your customer may not be able to use it, because their devices are v4, and the people they want to communicate with may have also IPv4.

A colleague of mine this summer was in a hotel that had gone that path, to provide IPv6-only service. They tried. I don't know why, but they tried. The observation was – my colleague had a fairly recent laptop, a Mac just like mine, that works perfectly fine with IPv6.

So it was okay, and he could go to places like Google and Facebook, but he could not go to some of his news organizations that he wanted to contact. He could not go to the website of the school where his son was. He could not go to his dentist's website. There were tons of places where he could not go.

Is the web reduced to Google and Facebook? No, it's not. There's much more to that, so this idea of deploying only IPv6 as some strict limitation, that's why this is taking so long to be deployed. That's what I mentioned earlier, there is a last mover advantage to the technology. Next question. [inaudible]

UNIDENTIFIED MALE: Hello, good afternoon. During the course of your presentation, you mentioned a word like quality of the IP address. What do you mean by quality of an IP address? The first question, and the second question –

ALAIN DURAND: Hold on to the first – let me answer this one. An IP address is a number, so the same quality of number three as number four, right? However, there are people who have built databases about IP addresses. Like I was mentioning earlier, you can use an IP address to communicate with somebody, or you can have a conversation about that IP address.

Some IP addresses have been put into lists of spam, well known spammers. So when you go and buy an IP address, if that address is mentioned in the list of spammers, if you originate a packet from that IP address, you are going to be filtered out. To you, the quality in terms of the potential for you to use this address is lower than if you use an address that has not been put into a filter.

What you need to do is to work with the people who are maintaining those lists of spammers, and say, "No, I'm now the new holder of that IP address, and I have nothing to do with those bad spammers. Please remove my address from your filters," and essentially clean up the address so that it gets to a better quality. Now, your second question.

UNIDENTIFIED MALE: The second question is, does the price of an IP address vary based on the range, or the fancy numbers or something like that?

ALAIN DURAND: Excuse me?

UNIDENTIFIED MALE: Does the price of an IP address vary based on the range from where you get, or like the fancy numbers you get out of it, or something like that?

ALAIN DURAND: I have never heard of people charging more because the address is fancy.

UNIDENTIFIED MALE: Or something like the range?

ALAIN DURAND: No, the only thing that makes the price different is if the address range is really known to be in a block that has been associated with spam. And even then, there are ways to clean it up. The fact that it is a legacy address versus a non-legacy address sometimes plays a little bit, but not always. More or less, it's the same price. It's just the market price.

STEVE CONTE: We have another question online.

UNIDENTIFIED FEMALE: I have a question from Arshad: "I want to know if the present infrastructure used for IPv4 is compatible for IPv6."

ALAIN DURAND: That’s a good question. Yes and no. Most routers – well, let's step back. Fibers, we don’t care. It could be IPv4, IPv6, could be anything else, it doesn’t matter. Most routers if they are relatively recent – like five, ten years old, no more – there will be no problem. Most servers, there will be no problem.

But there is always something in the infrastructure that is old. It could be, for example, some old DSLAM in the DSL world, or it could be something completely hidden in the back office, like for example a provisioning system that goes in configured cable modems or DSL modems, and in this system, there might be something that is IPv4 only.

So when you deploy IPv6 on a new infrastructure, it is usually much easier than if you do it on an old infrastructure. And a case that was brought to my attention was in Bosnia. Bosnia is one of those countries that are popping up on the chart of countries that are deploying IPv6 that I was not expecting.

The reason is after the war, they had to rebuild the infrastructure, so all the equipment there is fairly recent. So it was much easier for them to do it. So in theory, yes, the infrastructure can support it. In practice, if it's new or newer or newish, it will be much easier than if it's an old infrastructure. Question?

NAVEEN LAKSHMAN: I have a question regarding the black market of addresses.

ALAIN DURAND: What is a black market?

NAVEEN LAKSHMAN: Like when v4 comes in shortage.

ALAIN DURAND: Well, keep going with your question.

NAVEEN LAKSHMAN: The question is, how does ICANN support the process of transfer of addresses?

ALAIN DURAND: Two aspects. The first one is, there is a market for IPv4 addresses that is a white market, not a black market. It's a white market because there are policies that have been passed in all the different IRs to allow for transfers.

Somebody can go there and say, "I want to transfer my address block to you." You have to bring the proper documentation to say, "I am the one who used to have this address," and you are who you pretend to be. If you follow this process, this is not a black market. That’s a white market.

Now, there's something else that is happening. In order to do that, you have to follow the policies that are enacted by the different IRs. Sometimes, players don’t like those policies but still want to exchange addresses. Those are not transfers by definition. They are financial transactions where something is exchanged, but this is not a transfer of IP addresses.

Sometimes, this is done through private contract. Sometimes, this is done through letter of authorizations, and those are not documented anywhere in public database. It's all private. So do you call this a black market, a gray market? Everybody has their own terminology, but this is something that is happening that we know about.

What does ICANN do about this? All the number policies are actually delegated to the IR, so those are the ones that are actually looking at this. From a research activity, we look at this because that’s an interesting space, but we don’t enact any policies. Those are the IRs that do their own policies. Does this answer your question? Thank you.

STEVE CONTE: I have two more questions queued up, and then we're going to stop the queue, except maybe – are you asking a question too? No? Okay, two more questions, and then we'll stop the queue and end the session.

UNIDENTIFIED MALE: Can I have both IPv4 and IPv6 deployed [inaudible] system, and may be used as and when required? If IPv4 traffic comes, IPv4 may be used. If IPv6 traffic comes, IPv6 may be used. As a network administrator, can it be deployed in some network both the systems?

ALAIN DURAND: This is actually the way it's deployed in most places. It's called dual stack, meaning that you have an IPv4 and an IPv6 stack. If you talk to an IPv4 device, you use IPv4. If you talk to an IPv6 device, you use IPv6. If you talk to a device that uses both, well, it depends.

There's an algorithm called Happy Eyeballs that you start a communication with both, and you give a little bit of a preference to IPv6, and you look at who is coming back first. You start the communication with IPv6 a little bit earlier than one in IPv4, so you give a few milliseconds' advantage. You try to tilt the balance towards more IPv6.

There are some service providers in the wireless space – T-Mobile was one of those – who went a different route: they went IPv6-only on their handsets, and they have a NAT at the edge of a device that translates from IPv6 to IPv4 so that they can still communicate with IPv4 sites.

That being said, I will observe that if they had gone for a dual-stack system, the size of the pool of IPv4 addresses that we'll have to put on the NAT – being a NAT 6 to 4 or a NAT 4 to 4 –

would be exactly the same. So it doesn't change that. There had some other reasons to do it.

But what you're suggesting, having two at the same time on the server, is actually what most people are doing. One last question?

UNIDENTIFIED MALE: Recently, there is a lot of activity in the ransomware, spam, a lot of activities going on with IP base, and there are a lot of botnet [servers] there. Can ICANN have an [inaudible] number where they can [authorize] these IP blocks which are not related to spam or something for the downstream ISPs? To have a [inaudible] section, and then [drop the external root] or something like that?

Because there are a lot of IP addresses out there. Some of them we know, some of them, these people are like Trend Micro, different companies are giving these are botnet IPs or ransomware IPs. But if ICANN can have [inaudible] everyone can peer and then filter out these [inaudible] so that will be helpful for those ISPs [inaudible] customers.

ALAIN DURAND: That's an interesting question. The short answer is no. The long answer is, ICANN is not the network police. ICANN doesn't [authorize] anything on their network. This could be done by service providers themselves, and there are some forums where service providers talk about issues like spam and malware and botnet and phishing and all that.

This is an issue that should be brought there. For example, there was the MAAWG meeting in Paris last week, which is the malware and abuse and all this discussion there, where people are specifically addressing those issues. This is not an ICANN function. This is outside of the remit of ICANN.

STEVE CONTE: I'm sorry, we do have one more in the web room, and then we're going to wrap it up.

UNIDENTIFIED FEMALE: One last question from the chat room from [Purity]: "Is it true that Africa has plenty of IPv4 addresses compared to large population countries? If true, can Africa make money out of this?"

ALAIN DURAND: I have another presentation that talks about this. If you look at the total number of addresses that they have in the Africa region, no, they have much less than, for example, North America or in

Asia Pacific. If you look at how many IPv4 addresses are left to be allocated by the IR, they have a lot more.

They have about 1.8 /8s available, so that's about 20 million, when there's zero left, or almost zero, in North America, and a small pool in Asia Pacific or in the RIPE region. And the small pool is all the /22s that I was talking about earlier.

So they have this pool of addresses, and there are discussions within the AfriNIC community on what to do with that: either slow down the consumption so that it lasts longer or accelerate the consumption so that they can move to another phase that's similar to what is happening in the other regions.

Those are the discussions happening there, and in ICANN, I am looking at them. I'm interested into the outcome personally, but this is a discussion that is being made and decisions that are being made by the AfriNIC community.

STEVE CONTE: Great, thank you. While you were answering questions, your doctor called. Does anyone want to give a guess on when his appointment is? Do you want to shout it out? Yes, 2:30. "Tooth hurty." Sorry, I know. Let's hear a big "Oh."

ALAIN DURAND: Thank you. It's 5:00 PM somewhere in the world, so I guess I have some of Dr. Johnny Walker's prescription.

STEVE CONTE: Let's thank Alain for the session. I appreciate you doing this, and appreciate you all for coming out here today.

ALAIN DURAND: Thank you very much.

STEVE CONTE: We're going to take a break for lunch. Our next session is at 1:45, back in this room. It's with the registry operators workshop. We'll be talking about a couple of protocols: RDAP and EPP, and we'll have a gentleman, Francisco Arias from ICANN, and Joe Waldron from Verisign here to talk about those protocols.

And then this afternoon, we will have the Root Server Security Advisory Committee, RSSAC to come in, and they'll talk a history of the root servers, a little bit about Anycast, and some other really fascinating things, so I invite you all to come spend the rest of the day with us, but go eat lunch first. Thanks a lot.

[END OF TRANSCRIPTION]