SAN JUAN – How It Works: Internet Networking
Saturday, March 10, 2018 – 15:15 to 16:45 AST
ICANN61 | San Juan, Puerto Rico

UNKNOWN SPEAKER: March 10th, 2018. How It Works: Internet Networking, Room 209BC, 3:15 in the afternoon. [AUDIO BREAK]

CATHY PETERSEN: We're going to be starting in a couple of minutes. Thank you. [AUDIO BREAK]

Good afternoon, everybody. Welcome to our third How It Works session for the day. We have Alain Durand, who will be talking about internet networking. Alain. Kevin, do we have the microphone working up front?

ALAIN DURAND: Well, good afternoon. My name is Alain Durand. I work at ICANN in the office of the CTO. I want to make this session interactive, so please come to the front. Don't stay in the back. We have plenty of room in front, so please -- that will be the exercise of the day. We tend to be sitting a little bit too much during those ICANN meetings, so let's get some exercise and please move to the front. Thank you.

So, this is an introduction tutorial about the foundation of the technologies that are used in the internet. So, we are going to talking about naming, addressing, and routing. And maybe you already know about this, but will be opportunity to ask some questions, or you don't know anything, and please, also, ask questions, and we'll this like a discussion.

Next slide, please. So, the first thing is what I call, networking by numbers. It's like painting by numbers. It's like the very basic thing you start. You may have heard about the OSI seven-layer model, and we're going to go through all of them and add a few more. And then, we will really use this to go into the naming, addressing, and routing part of the tutorial.

Next. So, the OSI model is a conceptual model to talk about networking, and they have defined seven layers. We're going to talk about layers 0 to 9. So, extended a little bit to the left and to the right.

Next. The first layer is what we call "layer 0" -- is a physical layer. It's about what kind of media do we use to carry information. So, there are two fundamental technologies that are used in the internet, and they have different properties, and those properties makes where use case and the way people actually manage it -- very, very different. First one is things that are wired. And the second one are things that are wireless. So, wired, we are either talking about copper or fiber. People say, "Oh, fiber, it has to be fiber because it's much, much faster than copper."

Well, actually, up to a certain amount, it really doesn't make that much of a difference. There are people that have been using copper and a coax cable and provide something like a gigabit per second to homes. You don't necessarily need the fiber to do that. So, it's all about how you multiplex the signal, and what kind of distance you go. When you go to wired, you need to bring your wire to the home. So, in some places, it's relatively easy. In some, it may be a little bit more difficult.

So, for example, in a city, if you need to dig a trench, get your wires there, you may need to ask for a permit, and that may be relatively easy in some place and very difficult in some other places. The flipside is that if you bring your fiber to a big building, where they have, maybe 100 people living in there, well, you have -- potentially, 100 customers. The one fiber, 100 customers, that's actually a pretty good deal. If you want to bring fiber in a rural area, where there's one house every mile, that may be a much more expensive proposition.

Wireless has other properties and constraints. First thing is when you put wireless, we're talking about spectrum, and there's private spectrum and public spectrum -- unregulated spectrum and regulated spectrum, I should actually say. So, unregulated spectrum are things like WiFi in the home. You do, essentially, whatever you want with the frequency bound that are allocated to you. When you do whatever you want, is really nice -- you don't have to ask for anybody's permission, but other people can do the same.

So, you actually compete with other people. Like I live in an apartment complex that's fairly big, and if I open my computer in there, and I look at all the WiFi signals, there are about 50-100 of them. So, the spectrum is quite busy and, at times, well, I have to compete with the other signals and my router has a bit of a problem to send a proper signal without interferences.

So, my solution was easy. I went and I bought a more powerful router, so I have less problem and my neighbor has problems. And that's now a race -- and like whoever has a stronger signal. When you go to regulated spectrum, well, by its name, it's regulated, so you need to participate to process at a certain place by the regulator.

Sometimes, they have auctions. Sometimes, they have contests. Sometimes, they have different things to allocate to the spectrum. So, that's the first thing that you need to think about. Wireless is nice when you want to cover a wide area. But it doesn't provide as much bandwidth as wired, so that's where the balance is.

Now, in cities, you have another problem. You don't have to dig trenches, but you have to put antennas. So, in the picture on the right, here -- is there any way to make the screen a little bit bigger? Thank you. The picture on the right, what you see is an antenna on top of a historical building, and in many places that's kind of a no-go. You have a very nice building, and people think that antennas are not as nice and the competition of the two is not good.

So, in suburban areas, there's actually the same problem. We're not competing with historical buildings, but we are competing, sometimes, with houses that are expensive -- like million-dollar homes. And in those neighborhood, people don't want to see those big antennas because they think they're just ugly. So, is sometimes a paradox, but in some of the wealthiest area, the cellular reception is the worst, because nobody wants to have an antenna in the neighborhood. In other places, we found an interesting compromise -- is to disguise the antenna as a tree.

So, I have seen in one of the places while I've been traveling, there was what looked like a giant palm tree -- like 100-meter tall palm tree. I've never seen a palm tree like this, until I actually realized it was not a palm tree, but it was an antenna that has been disguised as a palm tree -- it just happened to be a giant antenna. So, that's the kinds of things that people do. So, when you have wireless, it's a shared media, right?

So, maybe you will have thousands of people connecting to the same antennas, and that's when you have to put things like priorities and things on the network, to make sure that not one person suck up all the bandwidth, and it's actually shared among others. Sometimes, we talk about 4Gs, 5Gs, and speed of like 1 gigabit per second -- it's actually 1 gigabit per second, shared by all the subscribers that are in the vicinity of this antenna, connected to this antenna.

So, if you tune things so that only maybe 10 subscribers connect to this -- 1 gigabit divided by 10 -- 100MB. If you put 1,000, that's not the same story; it's only 1MB per second. So, don't necessarily believe all the hype about all the wireless stuff. You have to divide it by the number of people that are actually going to use this.

Next slide. So, let's talk a little bit more about fibers, and what's simply my own bias, that's the industry I came from -- so, fibers is a piece of glass, essentially, and you send light through it, so it's point-to-point. So, I can send a fiber from, let's say, here to a data server somewhere in the basement. I send light, okay, so I have one communication. If I want to have multiple communication going through it, I need to multiplex, and the way to multiplex is to use colors. So, let me ask you -- what's your favorite color?

UNKNOWN SPEAKER: Blue.

ALAIN DURAND: Thank you. What's your favorite color?

UNKNOWN SPEAKER: Yellow.

ALAIN DURAND: And what's your favorite color?

UNKNOWN SPEAKER: Red.

ALAIN DURAND: Great. I'm very lucky, because there's no conflict here. Sometimes, two people choose red. Red is probably of a common color. So, if it's free, if you want to communicate with somebody else, you can decide to use the color that you like.

You will red -- oh, sorry -- blue, red, and yellow. So, that way, there's no problem. Your recipient on the other side just has to tune the laser receptor to correct frequency. So, if the frequencies are far enough aside -- there's no problem. If they are too close together, then there might be some problems -- some conflicts. So, that way, we can really share the fiber among multiple communications.

So, now, let's say that I would like to have a fiber between here and my office that is in Washington. I can ask somebody to lay a fiber specifically for me, or I can look if there is a fiber already in place between Washington and here. That may or may not happen -- and if I have to lay a specific fiber, it's going to be extraordinarily expensive, and it's going to take months and months and months to happen. So, if I just want to set up a fiber to bring internet to the ICANN meeting, for example, that's not necessarily the best solution.

I have to reuse fiber that already exists, but the fibers that exist don't necessarily go from my office in Washington, D.C. to here. They go from, let's say, my office in D.C. to Miami, and there might be another fiber from Miami to -- I don't know -- Puerto Rico, San Juan, and another fiber somewhere else. So, we have to put them together. So, that's what the next layer is all about.

Next slide. Thank you. So, we're going to create what we call a "fiber path." Now, when we are connecting those fibers together -- let's say it was three fibers, the one from my office to Miami, Miami to Old San Juan, and Old San Juan to here. The odds that you have decided to use and nobody else is going to use blue in the other fibers -- relatively low.

So, maybe you will use blue between here and Old San Juan, but in-between Old San Juan and Miami, we will have to convert to another frequency. Maybe you will use yellow then, because somebody else was using blue. Maybe you will have to use a variation of red -- maybe a lighter red, not your actual red, or maybe you'll lucky and you can still use yellow -- but it doesn't really matter, as long as the recipient on the other side knows what to tune for.

So, you have essentially to go and for each hub have an agreement between the sender and the recipient about what frequency to tune in. Think of it as a radio, like a two-way radio. Like a walkie-talkie. You have to dial it to a particular channel you want to talk to. So, if you dial it to channel 1, then you can talk to somebody on channel 1. And if that person has to relay the message of somebody else, it will use another channel -- maybe channel 3, and the party on the other side has to listen to on channel 3. That's exactly the same thing.

So, that way, we construct the path. So, on fibers, today, the typical speed that you can get is in the gigabits, so the slowest will be 1 gigabit per second. It's fairly common on a computer to have a bot with a 1 gigabit per second. On a server, you will probably find, today, things like a 10 gigabit per second, or 25 gigabit per second on a data center. Sometimes, you will find a 40 gigabit per second, which is actually 10 gigabit per second times four.

So, we'll have four ports of 10 gigabit per second that are multiplexing to one port of 40 gigabit. 100GB is relatively common on routers. You can get 200GB, 400GB the same way -- by multiplexing things. The standard are defined by the IEEE, and it takes longer and longer to define the standard for the next generation. And it's more and more expensive and the physics is getting really complex. But we'll see more of this coming. The question is, how much is enough? So, if we try to think of it, what does consume bandwidth today?

UNKNOWN SPEAKER: It's consumed by everything, I believe, because we have video and it's uploaded video.

UNKNOWN SPEAKER: Video.

ALAIN DURAND: Video. Anybody has another guess about what takes a lot of bandwidth?

UNKNOWN SPEAKER: Hi. I would like to say multimedia content, which includes, maybe, videos and maybe social media -- let's say, activities like Facebook, Snapchat, whatever.

ALAIN DURAND: So, things like social media posting -- the textile is nothing. Sound, not much. Video is much more. So, video on social media, yes, that's a lot. Now, we had like basic video, HD video, 4K video -- now, moving on to 8K video, virtual reality video -- but all of this is video, right? Now, how much -- how many movies can you watch at the same time? One. Kids can watch two, apparently. Don't know how.

So, let's think about the typical household, where you will have a classic household -- let's say, two adults and two kids -- that makes for six videos, right? Okay. Maybe seven, because there will be a TV in the kitchen and nobody's watching it, right? Make it 10, all right? So, 10 video. One flow for video in a single definition was about 2MB/second. On high definition, it went to about 20MB/second. 4K, yeah it's about twice as much -- like, let's say, 50MB/second. 8K, 100MB/second. All right? 100MB/second times 10 -- it could be 1GB per second. Okay?

So, we can scenario where we can draw 1 gigabit per second from a home. So, if you multiply this by the number of homes, yes, you can actually go into the speed of 100GBs and maybe more than 100GBs -- like terabits, coming from big aggregation routers. That's why we still need to improve this, right?

So, how much is enough? 100GB is not enough. 1 terabit, 10 terabits, that's probably in the near future; something that we will need, one way or another. Not necessarily on a single fiber, maybe it will be split on different fibers, but that's the kind of speed when you are thinking about heavy environment or like -- heavy environment -- I mean, high-concentrated environments with a lot of people watching high-definition content as 4K, 8K content -- that's where we are going to.

But still, that was a finite number, because four people in the home cannot watch more than 10 movies at a time, right? But if we are talking about machine-to-machine communication, then sky's the limit. There's absolutely no limit on how much bandwidth there could be. So, all the calculation that have been made on the internet, until about two or three years ago, this is about the max that we're going to get. Those calculations are completely off, now, when you are talking about machine-to-machine.

Let's say if I have to have a machine with like -- I don't know -- a hard drive of 2 terabyte. I want to back it up -- well, I'm going to send 2 terabyte, so 2 terabytes time 8, so that's 16 terabit -- if I want to send that in 60 seconds, so I will have to have to have really, really high bandwidth, right? And if a lot of people are trying to do that, then we are talking about, potentially, infinite bandwidth.

So, that's what we need to think about. When you are going to machine-to-machine communication and not human-to-human communication, or human-to-machine, then we can have potentially speed that will be -- a required speed that would be unlimited. And that's a fundamental change on the internet.

Next. So, that's when we talk about, now, networking. So, we have fibers. I have fiber path connecting different places, but that doesn't make a network, right? It just makes a media. When I want to connect all of those, that's when I get a network. And why do we do networking, because -- well, the world is not flat. We want to simply segregate traffic.

So, this is my network. This is your network. We are going to interconnect, but what is in your network, I don't want all of it. I want simply a part of the bit that's for me. I don't need to use all of your traffic, consumer flow. Keep it -- just send me what I need. So, this is about isolation. I want my traffic; I don't want your traffic. So, the internet protocol is a protocol that does just this. So, IP is the intellectual property -- no -- IP is the internet protocol, okay? So, if you are to talk about IPR, intellectual property right, there's some discussion about that this week, I'm sure, in another room. Today, we are going to talk only about the internet protocol.

So, next slide. All right. Now, I have a network, and I need to use the network, so there's a layer that's called Transport. So, let's say that I want to send data to Cathy. And I want to make sure that she has received the data. So, on the internet, everything is a packet, right? It's 1500 byte, most of the time. So, I'm going to send a little bit of it to her. So, I'm going to send a packet.

CATHY PETERSON: ACK.

ALAIN DURAND: I send some packet, some data, to Cathy, and she responded to me. She said, "ACK." That means, acknowledge. I heard it, so I know that she has actually received my packet. I can send another one.

CATHY PETERSON: ACK.

ALAIN DURAND: Good. That works. Keep going. I send another one. Nothing, right? What happened? Disconnect, possible. There could have been congestion on the network. There could have been -- it's a wireless network, if there was a plane going through it, and it just disturbed the signal. It could have been -- I don't know -- cosmic ray going through something. Plenty of possibilities. Most likely, it was congestion. Like packet losses is a thing of the 1970s and 1980s.

Now, the media [inaudible] reliable, so most of the time it's going to be congestion. Too much traffic, somewhere. And, for example, if there's a router that have five networks coming, five networks coming out -- and everybody's talking at the same time -- just too much traffic and processes all drops packet. Packet drops, congestion -- same thing. That's why she didn't get my message. So, what am I going to do? I'm going to send the same thing again. Cathy.

CATHY PETERSON: ACK.

ALAIN DURAND: This time she got it, okay. That's how the TCP, transport control protocol, actually works. But before we do that, I need to make sure that she can listen to me, right? So, I'm going to say, "Hi, Cathy."

CATHY PETERSON: Hello.

ALAIN DURAND: She responds to me. I know, at this point, that she can hear me, right? But does she know that I can hear her? No. She sent it. She sent the message, hello; she doesn't know that I received it, right? So, I need to respond back to her. Thank you, Cathy.

CATHY PETERSON: ACK.

ALAIN DURAND: So, at that point, we have done what is called the three-way handshake in TCP. I said, "Hello" --

CATHY PETERSON: Hello.

ALAIN DURAND: Hello. So, when we do this, it's a three-way handshake, like this. Thank you. Both of us are sure that the other one can listen. And I remember in my early days, one class about communication -- it was human communication, and they told me, "You can be as good as you want in your speech, but if a person on the other side doesn't understand, that's useless."

On the internet, that's the same thing. I can send out as much data to Cathy -- if she doesn't get it, that's useless. I need to make sure she actually understand the message; she received it, right? That's why we need to do all this layer of TCP. So, there's another protocol in the internet called UDP. It's a protocol that is getting quite a bad reputation these days, because there's no authentication possible on it; you just send thing and expect it works. It has been used, and abused, and really, really abused recently.

You may heard in the president's discussion about amplification attacks, things that have been used the DNS, by -- the latest one was memcached that amplified things and are not filtered -- it was kind of a problem child to the on the internet.

So, let's just move on to the next layer. So, now I can send data, but if I also want to stream movies -- remember, we said that video was the important thing, in terms of bandwidth. Well, you need to describe what you're going to send, like -- What is the movie title? What's the length of the movie? What is the encoding? Is it one encoding or another encoding? How is the voice coming? What are the languages? Are there subtitles? Are there some extra content that you can go and fetch somewhere else.

So, all this is described into a session that's called the RTSP, real-time streaming protocol, where you have the parameters that you need to actually understand what type of content this is.

Next. For other media and other type of content, sometimes, we need, what's called a presentation layer. So, where I'm going to send some data and this data has to be interpreted. So, I need to put some kind of a syntax around my data.

So, at first, it was like, just put random things and you and I know what it is and that's fine. When people tried to use binaries and say, "Okay, first I'm going to send 42 bits, and then another 12 bits, and another 5 bits, and another 10 bits -- and that's very compact, and if you know what it is, that's fine. But if you look simply at the traffic, you may not understand, easily, what it is, and it's kind of hard to extend if you want to put more things and define new [inaudible].

So, when the web, it was an offshoot of something called XML or SGML, at the time. It's a markup language where we start to explain, "Okay, we're going to send this. This is what's going next." So, this is going to be some text," Bold Carta, for example. But the thing that the cool kids, now, are using is called JSON, and they use dictionaries and there's an example on the right-hand side, on how do you describe a menu.

So, it's a menu -- there are different things -- so, it's a popup. There's some different menu items and different object-oriented method that can be started when somebody click on them. The nice thing about those dictionaries is that it is very, very easy to add more items, to add more vocabulary into the dictionary, and it's fairly easy to read. So, that makes for a very fast development. That's why all the new stuff is done this way.

Next. All right. Now, we have talked about technology, and we are going to stop talking about usage and people. My kids, when they are visiting me, they want to watch YouTube. That's all they care about, right? They want to go on the web, social media, and all those things. They couldn't care less about fiber, fiber-path network, protocols, UDP/TCP, handshake -- what they want is to get their content. If you put yourself in the shoes of a service provider, it's not until you actually deliver those services that you have delivered anything. If you deliver IP, yeah, fine.

But that's not a revenue. The revenue comes from applications that people are going to use and are willing to pay you to deliver the service to get that. That's really what it is. So, that's actually a good segue to my next slide -- thank you -- which is financial. When you build those networks, if there is no financial stream to actually sustain it, it's not going to go anywhere.

So, there are two ways in the internet to financing and all kinds of creative things. But fundamentally it's either you pay directly, or you pay indirectly. Pay directly is --you have a subscription model -- so, you're going to say, "I have a contract with my service provider. I pay -- I don't know -- $100 a month. I have all those services." Pay indirectly is -- you get the thing for free, but you have advertisements. Well, somebody's going to pay for this advertisement to you, right?

So, is the service free? Yes, kind of. But this advertisement, in order to be more efficient, is going to be targeted, which means that the people behind it, who are sending you this advertisement are going to study you and their business is to know everything about you, so that they can send you the most efficient type of advertisement.

And that's where we start into the privacy discussions, the GDPR stuff -- you open e-privacy discussions about how much of this is actually possible, or is reasonable to you or not. But at the end of the day, if you don't pay directly for the service, it usually means that you are the service. So, there's always a word of caution about what is free versus really free.

So, next slide. And, now, it's time to add the favorite layer, here, which is a political layer. Political actually means, how do people come together to make decisions? So, there are different ways of doing that. Sometimes, you have -- I don't know -- kingdoms, emperors -- you may have oligarchy, you may a technocracy -- you can all kinds of different models. The model we have in ICANN is a multistakeholder model, right?

And that's why we are all here this week, to get together from different perspective, different angle, different opinions, and try to make up some decision. Was that a question?

So, next slide, please. Okay. So, before I go to the next segment, is there any question? Don't be afraid. Yes.

LAYAL JEBRAN: Hello. Layal Jebran, from Lebanon. I wanted to ask about traffic, basically. You mentioned usually loss of a message, let's say, is the traffic, is the congestion -- and then you mentioned how much we actually use in a household, for example, in internet -- and then you mentioned it's to infinity later on -- so, where are we going with the internet, then? And how congested is it going to become? And what are we doing to be able to provide internet that withstands everything?

Because we all know everything is getting automative, like next thing, you can actually open this door from your phone and just walk out. Where are we going with this? Will internet become slower or less used? I don't know.

ALAIN DURAND: That's a very, very good question. Different people have different opinions on that. In my previous life, I worked for a vendor and had many discussions with very large service providers around the world -- in different geographies, all around the world. They were all worried at this race of bandwidth and a new generation of radio and all of that will make it very hard for them to be profitable in the future.

So, that's why they are all going to automation. The cost of buying the equipment keep going up, but the cost of running the network also keep going up. That when you cannot get new customers, your customers are paying the maximum that they can afford. If -- let's say, random number -- $100 -- they cannot go to $200 next year and $300 the year after. There's kind of a maximum that people are willing to pay.

So, you need to offer better service for the same money. How do you do that? Well, sometimes, you do it because the new generation of network allow you to make economies of scale. So, when you actually divide total bandwidth, total number of users -- you divide the cost by the total number of users, it's actually less expensive than the previous generation.

So, it's how you justify going to the next gen. Sometimes, you have to do it by simply reducing your costs, reducing your operating costs, and that's why people are going toward automation in networks. We have heard about things like SDN, software-defined networks; like NFV, network function virtualization. All of those things are fancy keywords for automation. And automation means they're trying to reduce the costs of operating the network.

So, that's where a number of folks are going. Is this going to be enough to make those people profitable? Hopefully so, because if it's not profitable, they're not going to run their business. If we want an internet to work, we need to realize that, at some point, all the players have to find a way to make a living. It cannot simply be free from the air. I hope this answer your question.

Let's talk now about naming, addressing, and routing. So, this whole started a few ICANN meetings ago, when I had a bad toothache. Really, really, really bad, and I needed to find a dentist. So, that was the beginning of my journey.

Next slide. So, essentially, I need to find the name of a dentist -- so, what is a name? Well, one of my favorite things to do when I talk about something or think about something is to go to the dictionary to find about the exact definition of a term. So, in the Webster dictionary, a name is defined as a noun -- a word or set of word by which a person, animal, or place is known, addressed, or referred to. So, they give example, "My name is Parson, or Köln is the German name for Cologne."

So, a name is your identity for that matter. If I know your name, I know who you are. So, I'm going to ask you to remember three things from all my tutorial today -- this is thing No. 1, okay? If I know your name, I know who you are. All right? So, please -- what's the first thing to remember today?

EUGENE OHU: Your name.

ALAIN DURAND: That's a good. What's your name?

EUGENE OHU: Eugene.

ALAIN DURAND: So, if I know your name, Eugene, I know who you are, okay? Do you remember that? If I know your name, I know who you are. Next slide, please. So, sometimes, can talk to Eugene. I know your name. I know who you are. Or I can talk about Eugene. Eugene is there, sitting in the corner. You see the difference between the two? I talk to him, or I talk about him. So, in this story, I'm going to talk to Cathy. Cathy, that bad tooth in the back, like is really, really, really hurting. I know you have been living here for a while before -- can you recommend me a dentist?

CATHY PETERSON: Sure. My dentist's name is Dr. Piña Colada.

ALAIN DURAND: Thank you. Now, I need to get to the office of this Dr. Piña Colada to help my tooth. So, we just had this discussion about Dr. Pina Colada, right? So, that's the important thing. We can talk about, or we can talk to. So, and we can pass this from party to party, so now, I got a referral -- I can go talk to Eugene and says, "Eugene, help me find Dr. Piña Colada." Right?

Because now I know your name; I know who you are. I know Dr. Piña Colada's name, and I know who he is.

So, next slide. All right. My name is Alain, spells with an i because it's French. When I was young, Alain was a fairly popular name. In my classroom, usually there were four, or five, or six Alain. That created some interesting situation, when the teacher was mad at us, or at one of us. And she would say something like, "Alain! Go to the whiteboard. I'm not happy."

And, we were five of us, and we were looking at each other and says, "Which one?" So, she would say, "Alain! Why are you not at the whiteboard?" and everything -- No idea what to do. And then she would say, "Oh, Alain Durand, go to the whiteboard!"

I knew it was me, not the others. So, you see here that when there's ambiguity in names -- and you really want to refer to somebody, you need to qualify the name. If there's only one person, you're fine. How many other Eugene are in the room, here? Are you the only one? So, that's okay. We can call you Eugene. But if there are a couple of others, we will have to do this disambiguation.

So, next slide. Okay. I know that I need to find Dr. Piña Colada, but where is Dr. Piña Colada? I don't know. So, I'm going back to Cathy. Cathy, could you, please, look up in your rolodex, in your address book, whatever you are using back then, and find the address of Dr. Piña Colada for me?

CATHY PETERSON: Sure. The address is 125 Route Canal Road, D.C.

ALAIN DURAND: So yeah, I think that may help my tooth. So, what we just did here is, by looking into a rolodex, is to map the name, Dr. Piña Colada to an address. This is called name resolution. When you use the DNS, if you have been in the DNS tutorial earlier this morning, that's the exactly what we're doing, right? We use a directory, the DNS, to do a mapping between a name and an IP address. Now, that in the DNS, we can do other type of mappings. It doesn't have to be that, but that's the most common today. So, that's what name resolution is about. Next. All right. So, I have a name. I have, now, an address. So, where is this dentist?

Next slide. Okay. Let's look at the dictionary about an address. So, an address are "the particulars of a place where someone lives or an organization is situated." It's where you are. That's really what it is. So, remember the first thing I want you guys to remember today -- I know your name. I know who you are. Okay?

Now, I know your address. I know where you are. Remember the movies -- I know where you are. Don't mess with me; I know where you are. That's what it is. I know your address. So, I know where you are. First hall, here. So, remember those two things. I know your name; I know who you are. I know your address; I know where you are.

Next. So, when you look at addresses, most addresses have some form of structure. So, I told you I live in Washington, D.C. The most famous address in Washington, D.C. is 1600 Pennsylvania Avenue NW; Washington, D.C. 20500-0003, USA. That's a complete address. If you want to know what this is, that's a picture there -- it's a small house painted white, okay?

Like every few years, people try to get in there, but the rent is getting more and more expensive, so -- the complete address is structured, and you can read it from the right to the left. From the end to the beginning. The most important part is USA, and then you get more and more localized. Like, forget about the postcode for a second, but D.C. is District of Columbia. D.C. is not a state, but most of the time it will be a state, at the state level. Washington is the city in the state. Northwest is because there are four quadrants in D.C. There's northwest, southwest, southeast, northeast. And Pennsylvania Avenue is the name of a street. 1600 is a number on that street. So, there's a nice structure that we can recognize.

Addresses are not always geographically organized. For example, in the U.S., there's a phone system, which we call a toll-free phones; you can call somebody and you don't pay. That person will pay. We call it the 1-800 numbers. When you dial this number, you have no idea where the person on the other side is, and actually, very often, you will talk to somebody in India, or somebody in another country.

Cellphone numbers. My cellphone number starts with 703 -- that's the area code for part of Virginia, next to Washington, where I live. But right now, I'm here in Puerto Rico. So, the fact that 703 is associated with Virginia, next to Washington, D.C., it doesn't tell you where I am, right? So, there's no structure there. Same way with IP addresses. There is no geographically structure. So, you're going to tell me, "Oh yeah, but I can take your IP address and know where you are." It's called geolocation, right?

The information is not encoded in the address. What people have done is reverse engineer the thing, and they have built tables with a lot of energy, a lot of effort to say, "This address belong to this service provider, and we have found that that street has been deployed there, and we have done some trace routes, we have done some measurements, and we believe it is in that suite, from that house. Geolocation, in the best case, can almost give you the house number. In the worse case, it can give you nothing. So, we have some measurements in another project, and we were seeing some -- we were tracking the deployment of IPv6, and we were looking at IET, and we saw a lot of IPv6 are in IET, and we said, "That's weird."

I talked to people from IET, and they told me, "No, there's absolutely nothing." So, the reason why is that they were using addresses from a service provider and those addresses had been geolocated to IET when it was wrong. They were actually going somewhere else. So, sometimes, information in the geolocation databases is simply wrong. So, it's not 100 percent perfect. You have to be aware of that.

Next. So, I was talking about names having scopes -- addresses have scopes, too. So, if you live in Washington, and you say 1600 Pennsylvania Avenue, it's pretty clear where it is. Everybody knows that you are talking about the White House. You don't even need to say northwest. Pennsylvania Avenue goes to southwest, also, but nobody really cares. We know what it is. If you live in France, and you say, "I need to go to Paris," everybody knows where Paris is. If you live in the U.S., that's a different story.

Sometimes I tell to my kids, “Let’s go to Paris, this afternoon,” and they start laughing and then -- because they know, now -- I’ve done the joke several times -- there is a small village called Paris that’s like a 20-minute drive from where I live. A very small village. Actually, there are 29 cities in the U.S. that are called Paris. One of the famous one was in the title of a movie, is “Paris, Texas.” The movie was in the ‘80s; it was quite famous. So, when you simply say, “I’m going to Paris,” in the U.S. that is a little problematic.

Next one. So, remember with names I could to Eugene. I could talk about Eugene. Same thing with addresses -- you have been living there, too, so -- Route Canal Road -- Is it a good neighborhood, or should I be careful when I go there?

UNKNOWN SPEAKER: Again, please? What?

ALAIN DURAND: Route Canal Road -- is it a good neighborhood?

UNKNOWN SPEAKER: Yes. I think so.

ALAIN DURAND: Okay. Do you agree with him? Is it a good neighborhood?

UNKNOWN SPEAKER: It sounds good. It sounds safe.

ALAIN DURAND: All right. So, I can just walk there; that's fine, right? What happened here is we had a conversation about what Route Canal Road, right? So, the address, which is where you are -- you can use this to go there, or we can use it to talk about it. All right. So, remember two things I want you to remember today. First one, I know your name. I know your address. Thank you.

Next one. All right. But an address is, again, not enough to communicate. I have no clue where this Route Canal Road is. I know the address; that's it. So, let's take a little detour again. Let's assume that I want to send a postcard to whoever lives in this small white house in Washington. So, I'm going to write the address, 1600 Pennsylvania Avenue NW; Washington, D.C. 20500, USA. Anywhere in the world, I can post a letter with this address, and it will arrive there. I don't know if anybody's going to read it. I've heard that actually they do read it, but it will arrive. Why? Because there is a collaboration system.

There's an agreement among all the different post offices in the world, so that they will move the letters -- so that it will arrive at the destination. This thing only works because there is this agreement between the different post offices. On the internet, this is exactly the same thing. Next one. How do I go to my dentist?

Next. All right. My favorite exercise is to go to the dictionary -- what is a route? Or root (route) depending if you're British English or American English -- or mispronounce it, butcher it the way I do -- I was born in France, so you have to excuse me for that. A route is "a way or course taken in getting from starting point to a destination." So, if I have a route for you. I know where to go. All right. No. 1 thing, if I know your name --

UNKNOWN SPEAKER: I know who you are.

ALAIN DURAND: If I know your address --

UNKNOWN SPEAKER: I know where to go?

ALAIN DURAND: And if I have a route for you --

UNKNOWN SPEAKER: I know where to go.

ALAIN DURAND: So, you two are a little bit in conflict.

UNKNOWN SPEAKER: I know where you are. Sorry.

ALAIN DURAND: Exactly. I have your name; I know who you are. I have your address; I know where you are. I have a route for you; I know where to go. If you forget everything, but remember those three things, that will have been a success for me.

So, next slide. All right. When you drive a car across a country, but you have never been before, and you don't yet have a GPS, because it was like a number of years ago -- you have to rely on traffic signs. There will be signs that will say, "You want to go to San Juan, this is this way." You want to go to another city; this is the other route, okay? Those signs have been put in place by the local authorities, by the transportation department, whoever it is. They have been in place before you actually start your car, right? If there were no signs, it would be much, much harder.

I mean I was Nepal, last week, and there was not a single sign on the road. The local people have no problem, but if I had to drive by myself, I would have been completely lost, right? So, it only works for most people, like me, when there are signs put in place before I start my journey. So, on the internet, that's the same thing. If you want to be able to send packets, you need to create a map, and you need to put those road signs.

So, I have created here on the slide a small network. And when you build those signs, you do it from the destination to the source; that's reverse. The source is going to wait to get to post messages. It all starts from the destination. So, on this, I wish I had a laser pointer, here. The destination is going to be connected by a service provider. The destination is going to tell the service provider, "You can reach me through this link. I have a contract with you. If you receive packets for me, please send them to me through this link. That's part of my contract with you."

So, the service provider, which is the first orange circle, here, is going to tell all his service provider friends that are neighbors that, "You can reach my customer by sending the traffic to me." That's fairly simple. "Oh, perfect. Thank you so much." But this guy is going to tell his two neighbors here, you can reach my customer by sending the traffic to me.

Those guys, in turn, are going to tell their friends that they have a direct connection with, “You can send the traffic to this destination, to me, because I know what to do with it. I’m going to send it this way.” So, this guy says, “I know where to send the traffic.” This guy says, “I know somebody that knows where to send the traffic.”

And then, it moves one step further. This one will say, “I know somebody who knows somebody that knows how to send it. Yeah. It’s the old movies again. I know somebody who knows somebody. That’s the same stuff. All this is a cooperative system. Exact same thing as a post office that I described earlier. Without cooperation from service provider, there will be no internet. There is no government that explain how this is going to work. There’s no contract, global contract. There’s no global organization that define this. It’s all by collaborative cooperation among service providers. That’s how it works.

So, next slide, please. So now, I’m going to start my journey with my car, following the signs on the road. So, I’m going to send the pack from the source, now, flowing to the destination, and all those routes will have been put in place way ahead of time. The protocol to do that on the internet is called BGP, Border

Gateway Protocol. So, I'm going to send this to this guy because he is my service provider.

This one is maybe a backup provider, but I'm going to send this to my first provider. And I'm essentially going to expect that my service provider is going to do the right thing. I have a contract with him. And part of the contract is, he's going to connect me to the internet. I trust him because of his contract. I trust that he's going to do the right thing. The right thing is what? Is sending it to somebody else who is closer.

And then we refine this process: one potato, two potato, three potatoes -- no -- one hop, two hop, three hop -- we call this a hot-potato routing. Essentially, we pass the packet to the newest one as quickly as possible. And all of them are going to get closer and closer to the one that really matters, which is a service provider that has a direct contract with the destination, and then he will send the packet there.

Now, you're going to ask me, "Okay, I remember your layer 8, from the beginning of the presentation. Who is getting paid?" Who is getting paid? My service provider is getting paid by me. This service provider is getting paid by this guy. But what about those guys there? Then it depends on their size. If all of them were equal size and a major player, we call them tier-1 service provider. Top service provider.

And, essentially, they say, “Well, I’m getting about as much traffic this way that you are going to send me the other way, we don’t need to exchange money. We are just going to do it because it makes a lot of sense.”

Sometimes, there is a discrepancy. This on may be, for example, 10 times bigger than this guy. So, there will be 10 times more traffic, essentially, flowing in one direction than the other. And this guy may say, “Okay, I’m willing to take your traffic, but you need to pay me.” And there’s some settlement agreement, usually, that describe how much they pay per megabit. This is how it works.

So, those guys usually go to some kind of pairing forum and figure out who do they need to pair with, or have a settlement agreement with, in order to have coverage of whole internet. There are a relatively small number of players that are large enough and they have the connection to just about everybody -- and those are those tier 1’s that I was mentioning earlier. And almost everybody else has some kind of a financial contract with one of those guys.

Next. All right. So, next slide. I have arrived to Root Canal Road. Now, I can have Mr. Dentist -- Piña Colada -- for him to look at my tooth and maybe do something about it. So, I hope that I’ve kept your interest during this journey from a toothache to a dentist, trying to go through naming, addressing, and routing. Very often those notions are mixed up. I want you to remember that if I know your name, I know who you are. The name is the identity. If I know your address, I know where you are. If there's a map somewhere, I know where you are. I have no clue where to go, but I know where you are. I know the route. I know how to go to you. Those are the three basic concept that I would like you to remember today.

All right. So, that's the end of my presentation, if you have questions, I will take them now. Shoot.

UNKNOWN SPEAKER: Thank you. Very, very interesting presentation. So, you mentioned that with IP addresses, it's not -- you can tell specifically where a person is or, sometimes, you can't if the records are wrong. I just wanted to understand that a bit more. Because, okay -- for instance, I use different services, like Netflix, so when I'm in my home country, I can see particular programs, but when I go somewhere else, like I come to the U.S., I can see totally different programs. But I find that if I go back home, it's a different kind of Netflix.

So, obviously, it can tell where I'm coming from, and so it's showing me different programs, depending on where I'm coming from. So, I just wanted to understand a bit more about your IP address showing exactly where you are or not.

ALAIN DURAND: Thank you. So, yes, Netflix is one of the cases where they do geolocation to offer you one type of service or another because they have contracts in different countries with content producers, content distributors, about what they can or cannot do. How do they do that? Well, there are some commercial services who do this geolocation. And those people send years, and years, and years trying to finecomb all the tables exist on the internet to figure out where this address is.

So, the first step is, addresses are located by service provider to their customers. And the service providers get them maybe from another service provider, upstream, or directly from one of the internet registries. They are called the INR; there are five of them in the world. In North America this ARIN. In Africa, there's AFRINIC. In Asia Pacific, there's APNIC. In South America and Caribbean, there's LACNIC, and in Europe, that's RIPE -- five of them.

So, when those allocate addresses to a service provider, there is a table somewhere that says, "This block belongs to this service provider, and it is in this country." Now, that should be enough, right? Except that sometimes service providers span multiple regions and then they may have addresses registered, for example, in Germany, but they are actually using them in Luxembourg, you know. So, that's why this is not enough.

And there are people who investigate the internet and are trying to use techniques as trace routes and they cap it off with their measure of the distance of the time it takes to go from one place to the next there are all counts for the things like that that they put in place, and they manually rebuild those tables to say, "For this particular address, is actually located in this country, in this city."

In the best-case scenario, they can tell you almost the street name and almost part of a block where you are. In some scenarios, where people have not really developed those map very much, then they cannot tell you much. They say, "It's in" -- I don't know -- "It's in Spain. That's all it is, and we're not even sure."

Now, there are cases where you can bypass this. For example, if you use a virtual private network, a VPN. So, you can VPN in from, let's say, Spain to the U.S. and then you arrive in the U.S. with a U.S. address -- and you are seen on the internet with a U.S. address. So, people sometimes do that to try to look at Netflix in the U.S. when they are in another country.

Now, I'm sure that Netflix has other ways to detect what [inaudible] addresses, and made do something about that. But that's how geolocation address are actually calculated and used. Does this answer your question? Thank you. Another question?

UNKNOWN SPEAKER: I want to ask about the UDP. You have mentioned that it's very -- this is the protocol with a lot of vulnerabilities, and we still use it for very critical applications like DNS. Why we still do it? Why we don't use DNS over TCP, for example, or something like that?

ALAIN DURAND: So, the technical properties are at the same time a strength and a weakness. UDP is really cheap. I send a packet to a server. The server does something, look up a table, send it back to me. It does not need to keep anything in memory about this transaction.

So, that's why it's really cheap. That's why people offer DNS resolution services for free, because it takes very, very little resources to do it. When you do the same thing over TCP, you need to establish with TCP, three-way handshake. Remember -- I say, Cathy, we are going to do a three-way handshake -- Hello?

CATHY PETERSON: Hello.

ALAIN DURAND: Hello. All this takes time, but also while we are doing this, both me and Cathy have to remember that we have a connection, and then we need to talk, send packet, and then tear down the connection. Let's say I forget to tear down the connection. Her server is going to have to keep this for a long time -- until they decide to give up, which can be minutes, several minutes -- actually, with documents, it could be 30 minutes or more, if you have millions of connections coming per seconds, and you need to keep it open for 30 minutes, that's very heavy, right?

So, that's why people are a little bit hesitant to move to a protocol like TCP, because it takes a lot more resources. Now, the flipside of it, just from a security perspective -- if I send a packet to a UDP server over UDP, and nobody checks my IP address, I can put a fake IP address, and the server is going to respond to that fake IP address and if the answer is bigger than the question, then that's an amplification attack, right?

The reason it works is that nobody has been filtering out my packet in the first place. If things like BCP 38, best common practice No. 38 document from the IETF on filtering out networks that don't belong to you. And if that filtering had been put in place, when I send my packet to my service provider, he should have noticed that was not my address. It should have dropped the packet. No, no, it's simply a reverse of that filtering. You are my service provider.

I send you a packet to go to Cathy, all right. And the source address I put on my packet is not mine; it's Eugene, all right. Why are you accepting my packet? You should not. So, there's a recommendation to service provider that they should not accept the packet but, unfortunately, that recommendation is not followed in many, many cases. And the result is, you're going to send the packet to Cathy. Cathy's going to respond to this packet, and it's going to go to Eugene. So, I will have attacked you Eugene, indirectly, using Cathy. That's a problem.

So, we are long-term efforts to try to move the DNS on other transport. There are some efforts to move it over HTTPS. There are some efforts to use QUIC, a new protocol that came from Google, initially. So, there are different ways of looking at this.

Next question? Going once, going twice -- all right. So, thank you all very much. Just remember, I know your name; I know who you are. I know your address; I know where you are. I know a route to you; I know how to get to you. Just remember those three little things. Thank you.

CATHY PETERSON: Thank you, Alain. Thank you, everyone. We have our next How It Works at 5:00 pm on root-server operators. It'll be a really interesting session. If you want to hang around, grab coffee, and come back by 5:00, that would be great. Thank you again.

[END OF TRANSCRIPTION]