DUBLIN – How It Works: Internet Networking
Monday, October 19, 2015 – 12:15 to 13:45 IST
ICANN54 | Dublin, Ireland

ALLAN TURIN: So, my name is Allan Turin. My office address is 801 17th Street Northwest, Washington, DC, 20006, USA. So that's my address. And the route I took to come here is a join flight from Western Dallas to Dublin airport and then taxi to the hotel and another taxi to come here.

All of this is an introduction to say we're going to talk about names, address, and routing. And always people ask me questions about those things. I sense some confusion, what is an address? What's an IP address? What's a name? What's the difference between the two, and that's what we're going to try to clarify today. So next slide please.

Remember maybe painting by numbers, we're going to do network by numbers who try to make it easy. This tutorial is aiming at people who are new to the field, this is not aimed at experts, so if you already know about everything network related and TCP/IP and all of this, you are wasting your time here, you might as well go and have lunch.

Next. So you might remember the seven layer of the OSI model. And I'm going to extend it from not one to seven, but from zero to nine. And so layer zero is really the physical realm, and there are two types of networks, they are fundamental and different properties. One is the weird network, but usually a copper of fiber, and the other one is a wireless type of network with antennas. So if you lay out copper or fiber, you need to dig trenches or you need to put wires on poles, and this is pretty heavy investment.

You need to get all kinds of legal authorization to do that. The plus side is you get a lot of bandwidth. You can get gigabyte, tens of gigabytes of bandwidth for a single fiber. This is essentially point to point, so if I have a fiber from let's say here to [inaudible], well it goes from here to [inaudible] and nowhere else.

Wireless is exactly the opposite. Usually the bandwidth are much, much lower. We are talking about megabytes, not gigabytes. When you have one antenna there is an area that you can cover, so you can have one antenna in the city, a small city, and cover the entire city. So in the whole area, for example, this is much more interesting, because the cost of infrastructure is much lower.

But in other places, people don't like antennas. In the neighborhood where I live, it's in a suburb, people don't like to see antennas next to their house. It doesn't look good. I would like to have people digging trenches and bringing fiber to the houses. So those are different [inaudible] that you have to deal with.

Also antenna, since you're using airwaves so there is a spectrum issue, and spectrum maybe regulated or unregulated depending on the type of network you're using. If it is regulated, you may have to participate in some option, or you may have to work with a government to get a slice of the spectrum. So it's a very, very different process all together.

So that means [inaudible] of a low layer protocols for wired networks and for wireless networks, and so the policies that goes around then are very different. Next slide. So people tell me sometimes, "Okay, I need a fiber to go from let's say my office in Washington to here. Am I going to have to dig a trench all the way from Washington DC through the ocean to here?" Well the answer is obviously no. But if I need another fiber from New York to here, am I going to have to really get two fibers? No, we're going to share fibers.

So the first way to share fibers is in a fiber, you can send light with different colors. You can send red light or green light, or

orange light. It's light color is associated with a wavelength, and you can send different wavelengths, so essentially tune your transmitter and receptor to a particular wavelength, and you have sufficiently spread apart, you can have multiple intervals.

So you can share fiber among different persons. But remember, it's always point to point. Next. So back to my example, if I need a fiber from my Washington DC office to here, I doubt there is a single fiber that goes from there to here. There might be bits and pieces of fibers, for example, a fiber that goes from my office to some kind of central office in Washington DC.

Maybe some fiber that goes from DC to New York, there might be a Trans-Atlantic fiber that goes to New York to let's say, somewhere along the east coast of Ireland. And then some more fibers, etc. So what we're going to do is to connect those fibers. Those fibers are after some switches.

And we are going to create a path by connecting both those fibers together. And it might be that we start sending light in a specific wavelength, but arriving over places in the network, we cannot use that wavelength, because it's already used by somebody else, that's okay, because it's going to be retransmitted, and you can be retransmitted in a different color.

So by creating all of those connections, we essentially in fact connect multiple tier one network, to create a layer two

network, which is this fiber path. Another example of a layer two network is when you have a wireless hotspot in a room. Like a 802 11 hotspot. Essentially what it does is create virtual wireless links between your laptop and your hotspot, and each of your laptop is actually connected to the hotspot, so there is a mesh network that is concentrated, especially meshes, so hub and spoke network, that's concentrated to the hotspot.

And from your hotspot on can go somewhere else. The point is, when you're on this network, and when you're on a fiber path, all of the nodes on that network can talk directly to all the other nodes. So if you send just a communication to one, from one to the next, that's fine. But if you start to send, for example, discovery packets that are typically broadcast or multicast packet pending, I'm going to send a packet to every node in the network.

Well if I have 20 nodes in the network, I'm going to essentially replicate the packet 20 times. Or if there is an access point, I'm going to send the access point, and then the access point is going to replicate it 19 times. 19 is not so bad a number, but think about if we add 10 billion computers, and we need to replicate 10 billion packets.

That will be completely ridiculous. So, those type of networks don't really scale very much. We need to aggregate them in

small packets and then reconnect those packets together, but make sure that when [inaudible] broadcast or try to discover things around me, that they don't propagate too far. Next slide.

So that's really what IP is all about. Seeing that the world is not just one local world. The world is about connecting networks, and Paul was talking about the IETF, and how the IETF is creating an Internet which is a network of network. That's really what this is about. There is a network for [inaudible], there might be another network for the other room downstairs, it may be different networks.

We're going to interconnect them. In fact, interconnected to the rest of the Internet. At the edges of those domains, there will be devices that we call router, that will forward the relevant packet from one network to the other. But if it is something like a broadcast packet discovery constrain into the local network where it started.

So when we talk about IP, IP stands for Internet Protocol, that's where this is. We are going to give devices and address that will enable them to cross different networks, to go from one computer to the final destination. Next.

Oh, I forgot to mention that [inaudible] my presentation. I would like to make this interactive, so if you have any questions please feel free to jump in at any time. Layer four is about not

simply being able to send packets, but about sending data. So a packet is a small chunk of data, it's called sometimes a datagram, typically it's about 1500 byte maximum.

1500 bytes is enough to say, "Hello, my name is…" But that's just about it. If you want to send a big file, that's not enough. You would have to send multiple of them. Now when you send those things, they get lost. Especially if you go from some of a low c networks, something wireless networks, have some interferences or there is some congestion in the network, some packet will be lost.

So you need to have a protocol that will make sure that there are things that are lost, they will be retransmitted. And make sure that you were talking to the person on the other side, and that person is listening to you and hearing you. So an example here, we're going to send a big packet, a big file to Steve. And the first thing I'm going to do is to say hello.

Hi Steve.

STEVE CONTE: Hi Alan.

ALAN TURIN: Nice to meet you.

STEVE CONTE: You too.

ALAN TURIN: So what just happened in this exchange here? I send a packet, in TCP language we call that a synch packet, for synchronization, say hello my name is. Steve got back to me and said, "Hi Alan." By saying that, I know that he has heard me, because if he had not heard me, it will not say hello back to me. Okay?

But at this point of the exchange, if we only do that, Steve doesn't know that I have heard him. He had heard me, he had sent something back to me, but he doesn't know, maybe what he sent back to me was lost. That's why I have to [inaudible] it was nice to meet you. Now he knows that he can hear me, I know that I can hear him. So it's called a three way handshake, NTCP.

And make sure that the two parties in the communication can talk to each other. So now I'm going to send a bunch of data, and I'm going to number them. So there would be a first chunk, second chunk, third chunks. Steve is a first chunk.

STEVE CONTE: I acknowledge that I received the packet full of data.

ALAN TURIN: What just happened here is that he told me that he has received my packet. We call this an ack in TCP language. Steve there is a second chunk.

STEVE CONTE: I received that data as well. Ack.

ALAN TURIN: Steve there is a third chunk. Nothing. I'm going to wait a little bit more. Nothing. That means that probably Steve has not received a packet. Steve there is the third chunk again.

STEVE CONTE: Ack, I received that data.

ALAN TURIN: Good. Now he has it. I can keep going on. And because he has not received a previous one, I might think maybe something happened. Maybe there was a wireless network, and some kind of a plane went through the antenna, or next to the antenna and make some interfaces, or it went through a router with a lot of people connected to it and there was some congestion, like a big intersection in a busy street.

Maybe I could not send my packet so quickly. So what I'm going to do instead of saying ack packet one, ack packet two. I'm going to send packet four.

STEVE CONTE: Ack.

ALAN TURIN: Packet five.

STEVE CONTE: Ack.

ALAN TURIN: Now it seems that communication is working fine. I have slowed down. Maybe I can increase a little bit my speed. Packet six.

STEVE CONTE: Ack.

ALAN TURIN: Packet seven. Now I need to back off again, that was too fast. So this will goal is what is going to be, just a second, is going to be the TCP exchange to find proper size of the windows, and to find proper size of speed that we can transmit. Question.

JASON HINES: Hi. Comment. Jason Hines. His acks are not specific. He's just acking and he's not IDing the packets he received in any way.

ALAN TURIN: You're correct. This is tutorial 101, this is not 301 or 401. We're trying to make things simple here, but you are absolutely correct. There is another protocol, it's called UDP where you can send packet directly, and you don't know if you ack acknowledge or not, don't know if they are received or not.

It's used mostly in application like telephoning, when you really care more about sending the traffic as fast as you can, however making sure it does arrive. Next slide.

However, there are some other protocols for real screen, and we were talking about telephony. And in some cases, it makes more sense to drop some packets, not retransmit them. Because if I'm saying something, I need you to hear some of the sentence, or it's a movie that's going full, we don't want to have movie rewinding and playing it a bit, and rewinding and playing it a bit.

Sometimes it's just better to degrade the quality but keep the flow going, for human perception, it is better objective. So those protocols do exactly that. Next please. So until now, what we

have done is send work data, and there is no semantic that was associated to the data that you have sent.

So that's fine when you are transferring a file, but if you want to transfer more, organize information. For example, you want to go and configure a router, you want… Don't want to send the config file, and when you are outside they'll have to figure out what to do with it.

You have to send something that is, already structured. So there are a number of ways of doing that, and that's what you do at a presentation layer, and at first started by having like some fix with character strings, and it has evolved things like XML have been used for a while, and now the latest thing is called JSON, J-S-O-N. Well what you see is essentially a dictionary of different values and different ideas, sorry, and values associated with them.

So this is a description of a menu, it's a pop up menu. If you look at things like open stack and a bunch of, some new tools that I've been propose, this is one of the languages that is used previously today. Next slide. And on top of that, of course, we have application. Because if there is no application, you know, and nobody is really interested into this low level counter and detail.

So application on the Internet is really about the web. That's the most dominant application. And http is the protocol to carry the web. Hypertext hyper protocol. There is a variation of it, it's called https, which is a secure version of it.

It's going to use topography and encrypt your data back and forth, over the communication channel. The beauty of this is that, from a user perspective, you don't have to worry about anything, to the point where kids usually are much better at using those applications as adult.

And it's an incredible success of the technology. Sometimes I've scratched my head and I think like, I'm a dinosaur. I have no idea how to use those new apps like my kids do. Next slide. But none of this makes sense if the community that is supposed to serve is not onboard. And sometimes it's about the financial issue.

You need to make sure that when you deploy something and the application, and your infrastructure, and your network, but you have written some investment. Because if you deploy this and nobody is using it when the bank is going to shut you down very quickly.

So those considerations are quite important. Sometimes it's not just about financial, it's also about making sure that you have some users that are interested by which you are doing. There is

always this catchphrase, there is an app for that, but not every single app is successful. Some are successful because they actually respond to a need.

Somehow not because nobody cares. Next slide. And all of this, of course, need to be coordinated for eventual organization, so there is a political layer onto that in how things are organized. By political I mean, our people comes to consensus on decisions that need to be made.

So there are a bunch of organizations and through this presentation we talked about the IETF and a bunch of organizations, so this is essentially an advocacy of different things, so here we are. This is ICANN right now.

That's where we are now. But there are other things like this in the world. So this is a quick walkthrough from layer zero to layer nine, about all of those different technologies. So next slide please. Now let's talk a little bit about, you know, the meat of this presentation which is about real story, started with me having a tooth pain when I was travelling.

And it was really, really bothering me, and I really needed to go see a dentist. So next slide. Steve, my tooth is really bothering me. I know that you have been [inaudible] longer than I have, and maybe you know a dentist. Can you give me the name of your dentist?

STEVE CONTE: Well I would suggest Dr. Guinness, but because you are in such pain, I would, there is, Dr. Jameson might be a little bit more appropriate for you.

ALAN TURIN: Thank you. So, what is a name? Well if you look at the dictionary, name is a word or set of words by which a person, an animal, or place, or something, is known, addressed, or referred to. Like my name is [inaudible]. So, if I know the name, I know who you are. I know Steve's name. I know who he is.

Now I know his dentist's name, not this Mr. Guinness thing, Mr. Jameson, much stronger guy apparently. And I know who he is. But still doesn't help my tooth pain. So I need to dig a little further, keep going. So, I know his name, so we can talk about Dr. Jameson, maybe I can talk to Steve about Dr. Jameson.

So Steve [inaudible] of a name. I can talk to Steve about Dr. Jameson. Naming itself is useless. It's only interesting when I use it either to start a conversation or to talk about somebody else as a referral. So let's use this Dr. Jameson, and maybe I really need to go see him because now it's getting bad. Next slide.

Names and scopes. How many Dr. Jameson are there here in Dublin? When, in my family, and we only want the name Alan. So next week, I will be at a family reunion and when you say, where is Alan? Well you know it is me, no doubt. But when I was in school, in grade school, Alan was a very popular French name for kids my age.

So usually we had like three or four Alan in a typical classroom. So when the teacher said, "Alan go to the blackboard please." Three or four of us would just looking at each other and say, "Which one?" When the teacher, oh, okay, Alan Turin, please go to the blackboard. So if you qualify a name, you can resolve those ambiguities when you have multiple people having the same name.

Next slide. But still, Mr. Jameson is interesting. I have his full name. Actually, what is his first name?

STEVE CONTE: Jim.

ALAN TURIN: Jim? Like Jim Bean?

STEVE CONTE: Yes.

ALAN TURIN: So Jim Bean Jameson? Okay. I need to find out where he lives, because if I go in the street and ask a cab driver to take me to Dr. Jameson, most likely he won't know where to go.

Yeah. So, what do I do? Well, I can take a phone book and go and look for Mr. Jameson with the letter J, or go to dentist and look at that. I can ask somebody. Where Steve? Where does Mr. Jameson live?

STEVE CONTE: I think he's at 125 Root Canal Drive.

ALAN TURIN: Thank you. So I can ask somebody else to tell me where it is. And I have now the information. That I can use to tell the cab driver. This process is called name resolution. So we talk about DNS, we talk about name resolution. That's really what it is, to map out from the name to something that I can execute at a lower level which is an IP address, DNS name.

So in that case, street address. Next slide please. So apparently this year, but different topics that have come up in recent years about names, and especially DNS names. First one was internationalization. So initially names were written with ASCII

character only. And uppercase, there was no lowercase. That's fine in America. It's not as fine in France when we have those weird characters, e with an accent. And in China and Japan, that's absolutely not fine because they don't even have those characters to begin with.

So over the [inaudible] quite important going on. DNS authentication. What is the address of Dr. Jameson again?

STEVE CONTE: 125 Root Canal Rood.

ALAN TURIN: How do I know that's a for real address? How do I know that he's not sending me to wild goose chase? Well, there are two issues. The first one would be, am I really sure that I'm talking to Steve? And the second one is, does Steve actually have the correct knowledge of where the address is? Maybe this Dr. Jameson has moved in the last six months and Steve doesn't know about it.

So DNSSEC is trying to address essentially the first part of a problem. Making sure that this is transferred securely, and there is a key associated to it, and that key has been signed by a person that owns the data, that is by Dr. Jameson. So I know that, yeah, this information is going to be correct.

Third aspect that has been discussed in the recent years is the expansion of a root zone. So in the beginning of the Internet, there are a number of country code TLDs, like dot UK, dot FR, dot JP, dot IE. And the global names were dot com, dot gov, dot org, dot net, there are others.

In recent years, we have seen many other TLDs being created, and you have a full list down, and a bunch of people who are deploying in this market now. So without reading the three more recent issues that have changed with technology that was created long, long, time ago. Next slide. Now let's go to this address, like 125 Root Canal Road. Address, going back to the dictionary, is a particular of a place where someone lives or an organization is situation.

So, essentially if I know your address, I know where you are. So sometimes I borrow money from some friends and they said, we know where you are. Okay? Meaning that, don't pay the money back, meaning they know how to find you. That's what an address is. Next.

So little [inaudible], so as I mention, I live in Washington, DC. So that's one of the most famous building in Washington, DC, just next to, around the corner from my office actually. And the address is 1600 Pennsylvania Avenue Northwest, Washington,

DC 20500-003, USA. That's the complete address. The first thing that you can notice is there is a hierarchy here.

You have to read it from the right hand side, all the way to the left. So the hierarchy is USA, this is not in Ireland, that's not in the UK, that's in the USA. This is DC, well this is not exactly a state, but for this example that we work. This is in Washington, so that's a city, and W is northwest, [inaudible] northwest quadrant of DC, because some streets are on different parts of the city. Some Pennsylvania Avenue and on number 1600.

So this very nice hierarchy. Not all addresses have the same hierarchy. For example, in the US, there is a toll free number system, 1-800, so they say dial 1-800 and then go buy something. You have no idea where the end of a call is. Sometimes it's not even in the US, it could be rerouted somewhere in India, in Ireland, in China, anywhere.

Also cell phone numbers. I have no idea where the cell phone is. Example in [inaudible] in France, all cell phone numbers start with 06, but I have no idea if that person is in Paris, in Nice, in [inaudible]. No idea. There is no hierarchy whatsoever.

IP addresses, there is nothing in an IP address that tells you, we have this IP address is. You can look at an address like, for example, 128.30.8.1, there is nothing there that tells me if it's in US, in France, in Ireland, in China. Nothing. Now what has been

done is that people have collected statistics about addresses, and have been some database that says, oh this particular address is, you're located in this particular country, in this particular street address.

Some people have built those things, but as an afterthought. It's not part of the address structure. Next. Same way as names, address have scopes. So if you live in DC, and you ask for the address of this little white mansion, people will tell you it's 1600 Pennsylvania Avenue Northwest. They will stop where it makes sense.

They don't have to tell you it's in Washington, they know that. They don't have to tell you this is in the USA, they know that. But they have to tell you that Pennsylvania Avenue, and they have to specify that it is northwest, because there is also a Pennsylvania Avenue Southeast or Southwest. So, there is a logical point where you have enough information given your local scope to figure out where it is.

Another example. If I ask anybody here where is Paris, they will say it's in France. Couple hour flight. Where I live, in DC, there is in Virginia, a small village called Paris. There might be like 50 houses in there. It's on the mountainside, I like to go there sometimes on my motorcycle, just on a nice ride in the countryside.

So when I tell my friends, "Oh, I'm going to Paris and I'll be back in an hour," well, you know, that's what I mean. I'm going to this local Paris, not to the other Paris. So there could be some confusion regarding this, and we have to make sure that you qualify with a wide scope. Next slide.

And so we are mentioning with names, that name can be handled, and we can use it quietly, or as a reference, that's the same thing. I can put the name, I can put an address on a postcard and send the postcard into [inaudible], or I can ask my friend Steve, again, the same question, what is the address of Dr. Jameson?

STEVE CONTE: 125 Root Canal Road.

ALAN TURIN: Thank you. That's a reference that he is giving me and I can use it. So this notion of talking to somebody or talking about somebody using [inaudible] or as a reference, is always in this Internet architecture. Next slide.

Now that I have this address and I can put an address, for example, I want to send a postcard to the White House, I can put the complete address I gave you, but that doesn't guarantee that the postcard will make it. What guarantees that a postcard

will make it is essentially a system that has been put in place, that's going to take this postcard, and move it from places to places, maybe on a plane, or on a boat, or on a ship, on a car, on a track so that it arrives.

The system is, a postcard is a post office. They will find a round to send a parcel, in that case my postcard, to arrive at the right place. So we rely on those systems, underlying system, to actually move things around to the correct address.

So on the Internet, what type of address do exist? We have two protocols on the Internet that have been defined. There are three of them, the first one was ncp, but it's long, long time forgotten. The current one is IP version four. It was defined in 1981 and it's still used now. An IPv4 address is described using 32 bites. So you have two to the power of 32 possibilities, that's about 4.9 billion.

There are number of addresses that are reserved for some special purposes. I talked about multicast earlier. That's used for TV distribution, for example. So in factors there are only 3.2 billion addresses that are usable. IPv6 has been designed in the early '90s because we realize that 3 billion addresses will never be enough. There are more people on earth than that.

So we need something larger. IPv6 has 128 bit, so it's two to the power of 128, so it's not just four times bigger than 3 billion. If

somebody knows how to pronounce that name, please let me know because it's a huge number. Essentially, that's… If you want to put an IP address on every grain of sand on the planet, we still have a lot more. So there is no worries about running out of IPv6 address, at least in the next few years. Question.

Very good question. So I'm going to repeat it. Why is it IPv6 and not IPv5? Why did we skip over IPv5? So, Paul told you about the registries that IANA maintain for the IETF. And that's a perfect example of that. There is a registry for IP version numbers. So it's a 16 bit, sorry. It's a four bit field so you can have 16 values from zero to 50.

When IP the next generation was defined, somebody went into that registry and said, what is the next value that is available? Until now, that somebody had already reserved a five. And was a protocol called ST1 and ST2, which was for a streaming protocol, that was the next [inaudible] protocol, but that already reserved the value.

So we don't want to duplicate value in the registry, so we went into the next value that was available, that was six. When IP next generation was defined, there were actually three candidates. There was one called [SIP], one called, one was [Tubar], and the other one was [PIP]. And some of them merged, but some of the other were documented for historical purposes.

So what they did is that they went and allocated version seven, version eight, and version nine, I think, in the registry. So if you go in that registry, you will see those things. There is not that they are going to be used, but we want to be able to refer to them in the future.

So that's why those values have been reserved, so if there is a new version of IP that will be defined by the IETF in the future, it will not be IPv7. It will actually probably be IPv10 or IPv11. Next slide please.

So I was mentioning those 3.2 billion addresses. And I'm sure you have heard about IPv4 exhaustion, so it's not an IP address get tired by working too long or being too long on the plane. What is happening is that all of them have been allocated, except just a few of them. But the Internet is still growing, so what do we do?

Next slide. First answer that most people have is, okay, use IPv6, you know? It has been standardized 20 years ago, just go ahead and use it. The problem is that they are not compatible. It's another type of VHS versus beta max type of thing, for those of you who are old enough to remember this.

They're simply not the same protocol, they're not compatible on the wire. It's like this is an American plug, but an Australian plug

doesn't fit. Sorry. It has a technical limitation. There is nothing to do with policy or anything. It's not compatible. Next slide.

So, the equipment on the Internet like modern computers, like laptop or Steve's laptop, I'm sure your laptop too, that support both protocol, v4 and v6. But there are devices that don't. Very often the home router, home gateway does not support v6, or I bought a TV set not that long ago, it's supposed to be a smart TV.

I can stream things like Netflix, Hulu, and a couple of others, but that TV has a firmware that supports only IPv4. I try to replace the firmware once and almost broke my TV. I had to be on the phone for two days with customer service from the TV vendor, and they helped me to solve that situation. And honestly, I think I know what I'm doing, and I failed miserable.

So I don't expect, for example, my parents to be able to change this low level software on their TV, which means that all of those devices, consumer devices, it's very important now when you're talking about Internet of things. All of those devices that are shipped will never, ever be reprogrammed. The software, the low level software, we call the firmware on those devices, once it has left the factory, can't touch it.

There are a few cases where it can, but it's very, very rare. So if those devices are shipped today with something, so only IPv4,

for the lifetime of this device, it's going to only support IPv4 and no IPv6. So if you start to see a lot of gizmos, [inaudible] gizmos today, for Internet of things that have IPv4, it will never do IPv6, end of story.

Which means that from a service provider perspective, you cannot simply say, "I am going to offer you IPv6 and forget about IPv4." We need to provide both as a service. Now it doesn't mean that every single device has to be configured with both, but it means that as a service to your customers, you need to offer both.

So it's going to be like that for a very, very long time. Because we still see new gizmos that are turning devices, popping up at [inaudible] in the US, Best Buy, or whatever the local is in different countries, that are IPv4 only. And the reason why the IPv4 only is simply because the people who designed it says, "I want you [inaudible] in people's home today."

Around the world, most places have only IPv4 not IPv6, so if you have IPv6, so why should I bother? So this is one of those case of late mover advantage or first mover advantage. That's why it takes so long to deploy this technology. Next.

So back to IPv4. All the efforts that we do to deploy IPv6 are really [inaudible]. But as I mentioned, we still need to carry on IPv4. Now all the addresses have been allocated, what do you

do? Well, there are two things that can be done. The first thing is, most of the RIRs now, I think the last one was working on it, have what's called a transfer policy that enable people who have network numbers that they don't use, to transfer them to people who don't have those network numbers but would like to have them to use it.

That's a euphemism for selling IP addresses. And this market has started. We're going to talk about it a little bit. Terms and conditions vary. Different regions have different policies on how you can transfer things or not. But this has, when you've left the station now, and there is a fairly healthy market around that.

The second technology that is going to help you is network address translation. This is something that everybody has been using in the home, later home gateway that connect home to the Internet, actually create a local space inside of home. And as only one global IPv4 address from the outside. So you're going to translate the address from the inside, onto that address from the outside.

Now service provider using exactly the same technology in the call of the network, and share addresses among different subscriber. So, in my previous jobs, I worked for one of the equipment vendors that builds some of this stuff. And we could very easily put 100 customers behind one IP address. So we

could multiply the space by 100. In some cases, we have seen some extreme ratio of up to 5,000.

So, [inaudible] very, very conservative number. Let's say that somebody gets a block on the transfer market of slash 16 in IPv4 world, which means actually 65,000 addresses to a power of 16. You can say, 65,000, I don't have enough to really make a business plan to serve customers with 65 user, that's nothing. But if I multiple the 65,000 by 100, now I'm talking about 6.5 million.

6.5 million is enough for a decent sized service provider in the country. Okay? So multiplying the space by 100 is a big deal. Next slide. So statistics about his market, this is a number of transfer [inaudible]… I'm sure you can get the slide somewhere. We see several hundred transfers being made every month, but not all transfers are the same size, so next slide.

So if I look at each transfer, and look at what each transfer is, and additional all of that, what it shows us is essentially we have been transferring about 25,000 slash 24s, all class C's, and networks that can have 255 addresses. If you multiply that by the number of addresses in each of those networks, essentially it shows that we are transferring about five to six million addresses a month. If you multiply that by the average price of

an IP address on those market, which is about $6, it's a very, very large market. It's a very large number. Next slide.

UNKNOWN SPEAKER: We have remote participants too, thank you.

UNKNOWN SPEAKER: I just noticed a huge peak in February 2015, especially in ARIN. What happened there that all at once it just stuck out?

ALAN TURIN: There was a very last transfer that was made. Let me reword this. It was a very last transfer that was recorded in ARIN. It's very different from being made.

Now, okay. And we talk about addresses and things. Now remember, this whole story began because I have a toothache. I really need to go in the taxi to see Mr. Jameson. So I need a route to go visit Mr. Jameson.

So, maybe the taxi driver doesn't really know how to get there, and he's going to drive, and he's going to look at signs on the road. And that will take him to the next city where Dr. Jameson is. But the signs on the road, they don't show up spontaneously when a taxi driver gets on the road. They have been put in place there by some of the long ago.

That means that if you rely on a system to essentially take you somewhere, that system has to be built, and those routes have to be computed way before you start them in traffic. So that's essentially the same thing that we're going to see here. So route, is a way of course of taking and getting from point A to point B.

So if I have a route, I know where to go. Remember, if I have a name, I know who you are. I have an address, I know where you are. I have a route, I know where to go. Next. If you only remember one thing from my tutorial this morning, this is this. You have a name, you know who you are. You have an address, you know where you are. You have a route, you know how to go there.

So, we are talking about building this infrastructure and those direction signs on the intersections. We do the same with Internet networks. So, this is usually build very way, so the destination is going to say, I'm here. If you want to reach me, send me a packet. The route, what it says before, is going to receive that announcement and it's going to say, okay, I know that to get there, I need to send the traffic to this link.

Okay, fine. Now, this router is going to advertise for all of these links, to all its neighbors. I know how to go to Dr. Jameson. I know how to go to that destination. So now, both adjust and

routers knows this, if I want to go to Dr. Jameson, I need to talk to this guy. This is, oh, I know a guy who knows how to go to Dr. Jameson.

Same thing we will propagate, and this one will say, I know a guy who knows a guy who knows how to go to Dr. Jameson. And that's really what the Internet is about, is I know a guy who knows a guy who knows how to get there. And remember, have your address.

So at some point, we're going to be [inaudible] route. Next slide. So now, those routes have been built, ICANN is directing traffic. So this route building happened in the background, and it's being refreshed every 30 seconds, every minute or so, but it has been long ago. And whenever somebody wants to send traffic, those routes have already been calculated.

This is not something that's going to be made just before the transmission start. It's something that is happening in the background, all of the time. So when I'm going to send traffic, I'm going to read those signs, like a go to an intersection and says, Dr. Jameson take left, Dr. Jameson take right.

That's all it is. So I'm going to look in the table. It says, I want to talk to Dr. Jameson, I need to send a packet to that guy because he told me that he knew a guy that knew a guy who knew how to

get there. And this one is going to do the right thing, which means, sending the packet to the next hop.

So you see here, there is a collaboration process that is really at the core of a management and the function of the Internet. There is a collaboration process, when we build those routes. When I hear announcements from one side of a network, and I'm going to propagate them to all my neighbors. And there is also a collaboration process, which is the reverse, exact reverse of what I described, when I'm sending traffic now.

So remember, first one was operational, building the routes. Now this is sending traffic. I'm going to forward the traffic to somebody else. And this is a collaboration that is at the core of the Internet. That's why, for example, exchange points have been build, where you simply try to transfer of a packet to the nearest possible place.

So there used to be examples where in the country, two universities could not exchange traffic directly. The had to send traffic offshore, overseas, maybe to New York, and there is an exchange point in New York and then it's going over the ocean again, back to the country. This was totally inefficient.

But it was done for a number of complex financial reasons and political reasons. But what's important is that they decided that it makes no sense. And what they did is they put a node very

close to both universities, so when their traffic that was going between the universities, it would not have to go through the fiber. So they found a better route.

They decided to collaborate to make the situation better. And really, that collaboration is at the essence of the Internet. Next slide. Now, what happened if there is a bad guy that abuses this collaboration? What happens if there is a bad guy that inserts himself in the system, and says, I know how to go to Dr. Jameson? Come here. I'm going to speak louder than the real guy on how to go to Dr. Jameson.

Maybe that's a plot from Dr. Guinness because all of his customers have been sent to Dr. Jameson and he doesn't like it. He wants to get the customers back. So, if we don't pay attention, traffic might get there because somebody is subverting the system. So the same way as DNSSEC, it was there to help us to guarantee that a name could actually translate to the real IP address.

Or IP SEC, that means when we have an IP level communication, we can have cryptography, and make sure we are sending the packet, we know that's the right endpoint on the other side. There is a technology here that's called [inaudible] for Resource Public Infrastructure. That's going to be used by the networks, not by the endpoints but by the networks.

When they hear an announcement of somebody saying, I know how to go to Dr. Jameson, that's going to be cryptographically signed. There is a public infrastructure and we can check this announcement to see if the signature is verified. If it does, great. If it doesn't, then we drop the announcement so that will prevent wild hijacking, for example, or wild insertion.

There is still a lot of debate around this technology. It's not widely deployed yet. The main topics of discussions are, you have a public key infrastructure, great. Is there one route to that key infrastructure? Or are there multiple routes? Is it one per region for example, one per RIR, one for Europe and for America, one for Asia Pacific, one for Africa, one for Latin American and Caribbean?

Or is there one single route that is an IANA route for it? Or is it country based? Or is it region based? Or is it city based? So multiple discussion on that. The second item of discussion is, we can verify the origination of a route, but can we also verify the path and make sure that it doesn't go through [inaudible] third party? Because somebody could be simply inserting himself into the route and eavesdropping on the traffic, and we may or may not want that.

So we still work is ongoing around that. Next slide. So, after all of this, I ask my friend Steve, who is your dentist?

STEVE CONTE: Dr. Jameson.

ALAN TURIN: Interesting. Where does Mr. Jameson live?

STEVE CONTE: He's at 125 Root Canal Road.

ALAN TURIN: And which street do I need to go to, go to Dr. Jameson?

STEVE CONTE: Well I know that way.

ALAN TURIN: There you go. That's why the basic technology of the Internet is all about. I can go to his dentist, next slide please, and then he will help me with my toothache.

That's the end of my presentation.

If there is any further questions, I would be happy to answer them now.

Question in the back.

UNKNOWN SPEAKER: Sorry, it's not just a question, but any chance to get that slide, a copy of that slide?

UNKNOWN SPEAKER: We're working… This is actually a very thick deck with pictures, so we're working on getting this slide deck onto the agenda. If you drill through the agenda to this session, it should be up by probably the end of today or early tomorrow.

ALAN TURIN: Another question?

NIGEL: Thank you. Nigel [inaudible], Caribbean Telecoms Union. The RPI thing, is it still in development or to what extent is it deployed in service right now. And what you see as the prospects for its ultimate full adoption?

ALAN TURIN: So there is software that has been developed, it has been running for a while now. You can get some of the keys, signatures, certificate, from the different IR. So if you want to go to ARIN, or RIPE, or AfriNIC, or LACNIC, or APNIC, you can get one of those our way, as they call it. Is it widely deployed? Not.

And that is essentially because there is still a lot of discussion about this [inaudible] infrastructure. Is it going to be routed in a single place? Or is it going to be routed in multiple places? And to answer the last part of your question, what I think the prospects of it is, until this discussion is really settled, it's difficult to think that there will be massive deployment. That's really hindering the process here.

The question is, are we at risk for hijacks? And I'm going to give you a yes/no answer, yes and no answer, to be more specific. Yes we are at risk of hijack. And hijack have happened in the past. We remember there was some incident in Pakistan, for example, it has been well publicized. [Inaudible] have learned, and because of this collaboration among service providers, incidents like that are fixed very quickly. So, are we at theoretical risk of hijack? Yes.

Are we at practical risk of hijack? And we say not as much. And if you look at the headlines in the last couple of years, we have heard a lot about breaches in credit cards and Sony Movie Pictures and all of that, but we have not heard a lot about root hijacking. So because of this cooperation between service providers, the risk is much lower, it has been mitigated.

Question in the back.

UNKNOWN SPEAKER: Hi. A very practical question if I may. If you think about the Internet of things, where everything is going to be hooked on the Internet, not only the computer but the washing machine, the dryers, the fridge, the TV, and so on, and so on. This could mean something like 50 devices in the home. How will this work?

Does the home have to be rewired or will this go wirelessly to a hub or something like that? Could you explain a little bit more on this practical question?

ALAN TURIN: Yes. In my house, I think I have something like 10, 15 device already. And I seriously think that in the next year or two, I may have 50 devices or more. Do I have to change anything in my infrastructure in my house? No. Because all of those devices are essentially wireless, 82 11 something.

I may have to change my home gateway, because some of home gateways have some limitation on the number of devices that they can support. And this is, again, in the NAT table, they have some older algorithm on how to do that. But that's essentially the only thing that need to be changed, from an infrastructure perspective.

Questions?

JASON HINES: Yes. Jason Hines. Back to the topic of the route hijacking. When you said there was collaboration between the service providers, I was wondering if you could detail a bit more how that collaboration worked? Does it depend on them coming into a meeting and being aware of best current practice, or best operational practice documents? Or just out there in their world, commercially and they're operating the automatically adopt whatever they need to…?

ALAN TURIN: That's a good question. So what is really the cooperation mechanism that exist? So first you have NOG meeting, like NA NOG for North America, NOG means Network Operating Group. LAC NOG in Latin Caribbean, AF NOG, and you have, it's called, Asia Pacific it's called [inaudible]. There is a RIPE community also.

So yeah, all of those meetings are where service provider go.

Excuse me?

Yeah, Caribbean NOG, and I think there is [inaudible] NOG, and all of those things for large or small regions. Fabulous opportunities for service provider to exchange information about what are the current threats, what are the state of the art

ways to deal with those threats, and how to they operationally defend against that.

Also, there are some informal communications that are happening. They all reach over from number and the Jabber ID, and there is all kinds of communication that are happening directly from the network operation centers, with those different ISPs. So something shows up, something happen, they will know very, very quickly.

When the very first incident on the Internet happened, it was back in 1989, you may remember the more [inaudible] one, it took about three days for everybody in the community and service provider to react to this. Now we probably take less than three hours. There is a question in the back.

JASON HINES: So my concern is if there is any risk if you're not participating in the NOG, like do the RIRs or something where you have to get the addresses from, are they pushing out information to under…?

ALAN TURIN: Yes, and they often, the IR meetings, are held jointly, they're held jointly with the NOG meeting. For example, last week, no two weeks ago, in Montreal, the first part of the meeting was NA

NOG from Monday to Wednesday, and ARIN was from Thursday to Friday. And many people just attended both meetings, because when you're travelling, three days or five days, it doesn't make that much of a difference.

JASON HINDS: Outside of going to the meetings though, okay. So I'm from the Caribbean, I see low Caribbean being in participation in a lot of these meetings. So I'm wondering, are my service providers more at risk because I don't see them as frequently, or in large numbers, at these meetings?

ALAN TURIN: There are a lot of things that is happening in the Caribbean region. I know that AARON is having some on the road workshops where, NA NOG is the same thing, LACNIC is also very active. There are many, many opportunities. There are also mailing list also for people who don't travel. There are a lot of resources that are available on the web. If you go to NA NOG website, you will find a lot of resources there.

UNKNOWN SPEAKER: Also a lot of cross-pollination within the mailing list too, so even if you might be in the Caribbean region, your ISP might be joining AF NOG or APRICOT, because it's a good collection of

expertise around the world. So they might have more than one avenue. So Alan, you're saying that networking isn't just digital though. In order to make this happen, it's got to have a human element to it.

ALAN TURIN: Yes, absolutely. This is about the human element in the end.

UNKNOWN SPEAKER; Alan, using your metaphor of name, number, and route, where do DNSSEC and DANE map into that?

ALAN TURIN: So DNSSEC is really about making sure that when I ask Steve, Steve, what is the address of your dentist Dr. Jemison?

STEVE CONTE: 125 Root Canal Road.

ALAN TURIN: This is all about knowing that this information is correct. It might be somebody is in between the two of us. And is hijacking the communications and sending back data to me. Or it might be that somebody has corrupted this local information and on

good faith, Steve is telling me some bad data. That's what DNSSEC is all about.

Now, DANE is more work in progress. I don't know, Steve, can you talk more about this?

STEVE CONTE: Actually I haven't been following it in a while, so I don't want to give wrong information because of corrupted data.

ALAN TURIN: So this is work in progress in IETF. My understanding of it is, to make sure that when you ask a question, that the question cannot be intercepted by somebody else. That's privacy in that question, because I may want to go somewhere, and I don't want anybody to know that I'm going there. So I don't want anybody who will be listening to the network to know that I am trying to go to Dr. Jameson, because if I know that I'm trying to go to Dr. Jameson, it's meaning I have too fake, meaning that I'm not exactly receptive to what is happening here, and they can use this opportunity to maybe pass on policies and everyone knows that I will not normally pay attention, I will normally pay more attention to.

You can create more interesting examples about that, but that's really what it is. It's to protect your privacy, when you are asking those questions.

UNKNOWN SPEAKER: And DNSSEC can be sometimes misleading because even though it sounds like DNS security, it's not actually encrypting the data. It's more like DNS auth, authenticating the data. And the DANE and some of the [inaudible] stuff that's happening in the IETF is more about protecting the content that's being transferred, and making sure that it's only the host or the client can see that data. Any other questions?

ALAN TURIN: Any questions on…?

STEVE CONTE: Let me check our remote participants.

I do not have any questions from the remote participants. So Alan, I'd like to thank you for joining us today and giving us an overview of Internet networking.

We are on small, a little bit longer break. We ended up early. We'll be back at 2:00 with registry protocols. And again, selfless plug, please take a second, go to this website, give us some

feedback. We are, we've got a great turnout today, but we don't want to ever run into the diminishing returns aspect. And so having new ideas of what we could talk about that's relevant to you and relevant to the ICANN community is super important.

So please take, it's 10 questions, if that, take a second and give us some feedback. So Alan, thank you.

ALAN TURIN: Thank you very much for your attention today.

[END OF TRANSCRIPTION]