OLIVIER CRÉPIN-LEBLOND: Good morning, good afternoon, and good evening everyone. This is an ATLAS II At Large summit capacity building webinar, the second one of our number of webinars we are going to have until we all meet face to face in London. And today we have Leo Vegoda, who is the operational excellence manager on IANA staff, who is going to be speaking to us about IPv4/IPv6 transition.

And the capacity building that is need on this worldwide, and all of the different issues of IPv4 to IPv6 transition. As you know, we are running out of IPv4 addresses. ICANN deals with not only domain names, but also with the numbering side of things, and that every computer that gets connected to the Internet has a number.

So since we're running out of those numbers, let's see where we are going and how we are going to do that. Tijani Ben Jemaa is unfortunately not with us, so I'll be your host today. And I guess I can hand the floor over immediately to Leo Vegoda. Leo, you have the floor.

LEO VEGODA: Thank you very much. So, let's move on with a review of the agenda. I want to go through a brief history of IP addressing, take a look at the relative sizes of IPv4 and IPv6, have a look at feature equivalents, some transition technologies to help us move from IPv4 to IPv6. Look at a few graphs showing the current status of IPv6 deployment, and then at the end, a question and answer session in which I will do my best to answer any questions that you have.

So, let's move forward. What we have here is a picture of Vint Cerf. Vint Cerf was one of the two people who created the Internet protocol. Now, the reason I've got Vint Cerf in the picture is because he was the manager of the project, and he owned the budget. And basically, the thing is, we aren't really, according to the original project plan, meant to be connected to the Internet, or at least we're not meant to be connected to this Internet, the one that uses IPv4.

The IPv4 Internet uses a 32 bit address space. And the 32 bit address space was significantly bigger than any network address space had been before. Typically, people used, back in those days, an eight bit or maybe a 16 byte address space. And an eight byte address space gives you 256 hosts, and a 16 byte gives you 65,000.

A 32 byte address space gives you four billion. And Vint thought that that would be big enough for the experiment that he was running, four billion, because no experiment that he envisioned would need more than four billion hosts. But he had a success problem, which is that people really liked it and they deployed it.

So let's take a look at the three phases that we've had up until now, I suppose, with IPv4. The initial plan was that there will just be 256 networks, or up to 256 networks. And each of them could have 16 million addresses. And a little while later, within a year, they realized that wasn't going to work, and they came up with what was called calssful addressing. So it's just like buying clothes, big, medium, and small.

Or small, medium, and big. Either you've got a network that was big enough for 16 million addresses, or you've got a medium sized network big enough for 65,000 addresses. Or you've got a small network, enough for 256. And that was good up until the start of the 90's, at which point people realized that they would soon be running out of the middle sized networks, the class B's.

And two efforts were developed to resolve the problem in the short term, and they were the development of CIDR, which is C-I-D-R, classless inter domain routing, and it's classless because that big, medium, and small is called classful.

That classless domain routing was developed in the early to mid-90's, and also network address translation was developed. The network address translation was something that we all experience even if we are not aware of it, because that's when you have one unique address, and multiple devices sharing that unique address. So generally, that's used to identify a subscriber or a subscriber line.

Now, of course, the longer term solution to all of this was IPv6, but that was the short to medium term development. So, moving on again, let's take another quick look at the history of the significant dates in IPv4. So it was developed officially and launched in 1980. Classes were developed the following year. CIDR was introduced, that's the classless inter domain routing, that was introduced in 1993.

ARIN was formed in 1997. Now, you go and look at that and you think, why is ARIN's formation significant? And the reason is not because ARIN is any more important than any other RIR, the reason is that the formation of ARIN marked the formal move from a government contract for the distribution of IP address, to that function being moved out to the community.

So it was the end of the central registry, and a particular US government contract, and it was noteworthy as the formalization – or the beginning of the formalization of multistakeholderism in this particular aspect of managing Internet resources. The first global policy was developed in 2005. The second global policy for IPv4 address space distribution was developed in 2009.

And the last IPv4 slash eight in the IANA free pool was allocated in 2011. So, those are the highlights, the significant dates in IPv4's history. IPv6, there is quite a bit of overlap as you can see. In this announcement which was made in 1999, the IANA functions, which at the time, although being performed by ICANN were still running on some ISI systems.

They went and allocated IPv6 addresses to the RIRs so that the RIRs could allocate IPv6 addresses to ISPs for use in real production networks. And that was just three years after the initial IPv6 protocol had been standardized. So, once it became available, vendors started the work in supporting it on hardware and in software, initially in software.

So, let's take a look at the relative sizes of IPv4 and IPv6. And here I've got a very famous cartoon from XKCD. And this is from 2006. And everything in green is an unallocated IPv4 block, and everything in white is allocated. And as you go and look at that, you can see that between one third and one quarter of the address space was yet to be allocated, and that's all of the address space there.

So, in just 26 years, almost all of the IPv4 address space had been used up. If we go to the next slide, this is from 2011, and it goes and shows our 128 byte address space for IPv6, and it's not dirt on your screen, those little dots are actually what has been allocated. You can go and download the code that generated this map from v6 stuff dot [?] dot net, and you can run it yourself, pulling in the public information for the address space distributions.

And you can go and see the change since 2011, it's not a huge change. Not because there hasn't been lots of IPv6 deployment, but because the amount of IPv6 address space that we're using is really very tiny in comparison with the entirety of IPv6. Because 128 bytes is much, much more than four times 32 bytes, because it's an exponential growth.

So, when you go and look at that, you go and think, but what does it mean for me? And what are the differences? So I've got this nice Australian road sign here, which focusing us on the slash 64, the LAN segment, because in IPv4 each individual address is precious. There are seven billion people on the planet, there are roughly four billion IPv4 addresses, there is a real mismatch.

There is basically half as many addresses as people, and we know that most people use more than one address. So, we've got a problem. In IPv6, each LAN segment which is a slash 64, has 18 trillion addresses, roughly. So, that's a massive, massive difference. So what we do is we don't count addresses, we count LAN segments.

And the initial guidance from the IETF to the RIRs was to hand out slash 48 to individual subscribers. So a DSL line, a cable connection, that sort of thing. And a slash 48 is 65,000 slash 64s. So even a very large house, one for a massively wealthy person with hundreds of rooms, is probably going to do quite well with a slash 48. It's probably not going to need another.

Nowadays a lot of ISPs are handing out slash 56s, which is 256 LAN segments. And for most people, that's probably enough certainly for the next few years because in a relatively modest house, 256 LAN segments is probably going to be fine. And there is no difficulty in getting a slash 48. So, the ISP needs an allocation to make their slash 48 assignments, or their slash 56 assignments.

And the minimum allocation in most of the world is a slash 32, which is again 65,000 slash 48s. So if you're a small ISP, that will be fine for you. Certainly if you are an ISP serving say, for instance, a small island nation, a slash 32 will be plenty for you for more than a decade, probably. Of course, there is no shortage of address space and very large ISPs can get much more than that, so there is no problem.

ICANN, in performing the IANA function, allocated slash 12 blocks to the RIRs. And a slash 12 is enough for basically a million slash 32s. So what we're going from here is 18 trillion addresses in a LAN segment, 65,000 LAN segments in a home connection, possibly, 65,000 customers for an ISP, at the very smallest, and there is loads of room for ISPs to have millions of customers if they need. And room for a million small ISPs in the initial allocations to the RIRs.

So, plenty of room and there is loads more to come. Because I haven't put it in the slide, but there are 512 slash 12s in the current one eighth of the IPv6 address space that we're drawing from. And what that means is, if the Internet grew at a fantastic rate, and each RIR needed a whole slash 12 every year, which is very difficult to believe.

It would be 100 years before we run out of the current one eighth of the IPv6 address space that we're using. So we've got a lot of space to play with. So, moving on, let's talk a little bit about feature equivalents. So it's really important, we have a very highly developed and deployed Internet already.

We must do, because we've used up almost of the IPv4 address space. So, if people are going to adopt IPv6, IPv6 needs to offer those people everything that they have on IPv4, otherwise they're actually not getting as good a deal. So, what have we got? Well, all of the desktop and mobile operating systems have IPv6 support. All of the main server operating systems have IPv6 support.

Most of the infrastructure equipment has IPv6 support. Basically, if it's connected to the Internet, the chances are it has IPv6 support. Now, there are always some things which get updated so very slowly that they will never have IPv6 support. The elevator control system in a tower block might never have IPv6 support. And a nuclear power station built in the 1960's might not have IPv6 support.

But most things that you can buy off the shelf today if you're an ordinary consumer, if you're a business deploying a service, IPv6 support is not a problem. So where is the problem? The problem isn't the device and it isn't the network, the problem is the little boxes in the middle, the customer premises equipment, CPE. The problem with CPE is that the people who sell them, sell them based on price and that means the cheaper they can make them, the more profit they make.

And they don't really sell on features because consumers, on the whole, don't differentiate based on features, they only go on price. So the shoestring budget means that a device will never get an update after the first six months that it goes on sale. It will be deployed or sit in the dark probably, it will never have any maintenance, and it will only be replaced when it breaks.

And most of these devices do not have IPv6 support. The new generation do have IPv6 support, but we can expect minimum of five years before there is a significant amount of IPv6 support in customer premises equipment. The kind of ASL or cable modem that people use to access the Internet.

But is this a major problem? And maybe it's not. What I've got here is a screenshot from my personal mobile phone. It's using T-Mobile's network. T-Mobile's network is a native IPv6 network. You can go and see on that picture, I've put a little red oval around the IPv6 address that the RIPE NCC says I'm connecting from. And increasingly, people are connecting to the Internet on a mobile phone.

And for the next billion people, a mobile phone might be the primary device that they connect to the Internet with. So, maybe the CPE problem isn't a huge issue, at least initially, because we've got a few years to resolve it. But it can be resolved in other ways. And one of the ways it can be resolved is using large scale, or carrier grade, network address translation. So, I mention network address translation earlier. It was developed in the mid-1990s.

And the way it has been used traditionally is that each ISP has given a subscriber a single unique address, and the subscriber has shared that address between multiple devices. So there is a private address space on the inside of the network, and as packets parse out to the Internet and come back from the Internet, their IP address is rewritten by the customer's modem.

And that has, on the whole, worked okay. Lots of people hate that. They go and say that NAT, network address translation, is evil. And they have a point, but we are where we are, and it seems to have worked okay. So, carrier grade NAT, or large scale NAT, does the same sort of thing, but it does it for the ISP. So what happens is, instead of a subscriber getting an unique address, a subscriber gets use of an address that is shared between multiple other subscribers.

So there are multiple layers of network address translation. And this is what has happened on mobile phone networks for smart phones and things like that for years. And for telecos who have run mobile phone networks, they go and say, "We do this already. What's the problem?"

And the problem is that as phones get faster, and bigger screens, and they can do more, or people connect to the Internet using a laptop or an iPad or something like that that has more capability, they want to do more stuff. And if you use services like iTunes, or Google Maps, or things like that, you can open dozens of connections at a time, and it's the number of connections per subscriber that counts, because when you're sharing an address what you're doing is you're getting a portion of the port space.

And a way to think about ports is they're a bit like channels on a radio, and there are 65,000 ports, and normally if you are sharing 65,000 ports between half a dozen people in a home, or some people in an office, it's not a big problem to use network address translation. But if you only have 200 ports, and you want to download music, and watch a YouTube video, and look up your route to a customer tomorrow on Google Maps, you can have a problem.

And it's not just a problem for you as the user. One of the problems is that in a lot of countries, not everywhere, but in a lot of countries, there are regulatory requirements for logging IP address usage. And if you have to log usage when you have carrier grade NAT, then you have to log a lot more data, because you have to log for each session flow. And that can be very, very expensive, because it's not just when did the customer connect, when did they disconnect, and what address is assigned, it's a lot more information.

It requires very accurate time keeping, and it takes a lot of disk space. The amount of money that you can spend on network attached storage, or network accessible storage, is huge because if you're running a big ISP with millions of customers, and those customers are doing tens of thousands of session flows each day, per customer, it can cost you a lot of money.

So, this can cause problems to some extent, and it could be a business advantage to some ISPs. They could try and use this as a way of getting content providers to buy transit from them. Basically get the content provider to pay them money to access their customer's faster. That could also create regulatory issues. It also makes the end to end model, which is one of the design goals of the Internet protocol, much more difficult.

And it makes it very difficult for startup, over the top service providers like [Viber, What's App], various other things that maybe haven't been invented yet to get going, because if they have to arrange a separate agreement with every single access provider, than that's a lot of agreements before they can actually start their service. And if they're a startup and they only have a small amount of capital to work with, that might not be practical.

So, let's go and look at where we are with IPv6. And what we've got here is a graph showing the Alexa top 50, which is averaged over all the countries which Alexa does reporting for. And we've got almost 9% of the Alexa top 50 accessible over IPv6. And you can go and see a really big spike there, and the first spike was IPv6 global launch, and then the more recent spike was Cloudflare switching on IPv6 support for all of its customers by default.

So this is content becoming available over IPv6. Another way to look at it is, how is that content accessed? Google publishes information about its properties, and they have almost three and a half percent of their connectivity going over IPv6. So three and a half percent of their customers, or three and a half percent of their traffic maybe, is IPv6.

So that's quite good. Coming up to 10% of the Alexis top 50, and about three and a half percent, and quite a steep curve. I wouldn't want to cycle up that hill for IPv6 deployment. Cisco does another analysis, they go and look at what has been rooted on the Internet. Basically whether an ISP has announced an IPv6 address block.

Goes and looks at the website content availability. It goes and looks at several other factors, and the more green a country is, the better its IPv6 deployment. So you can go and see some countries in Europe that are particularly green. And you can go and see some countries that's almost completely white and with no IPv6 deployment.

It's not unusual in the early stages for something to go into the – something that is so distributed. But what we do have, and I'll just go back a slide, we've got some real sharp growth there with some vertical lines and some steep curves, and then another very steep curve in the growth for demand of content over IPv6. And so, this map that Cisco shows, also is taking account of whether that can be used, and we've got quite a lot of green and it's spreading, if you've been watching this map over the last couple of years, it's very good.

So, the initial plan was up to 256 networks, each of which could have 16 million addresses. And we now have a good toehold in there for IPv6, which gives 18 trillion address to every network segment, and it gives plenty of network segments to every single Internet subscriber. And there is no shortage of address space for ISPs.

And 128 bytes is really going to last us quite some time. So, I have tried to cover the basics there. And hopefully, what I've described is useful to you. I am happy to take any questions that you have. So, if I can hand the floor back to Olivier. Thank you very much.

OLIVIER CRÉPIN-LEBLOND: Thank you very much Leo. And I was going to, just before opening the floor for questions, I was going to give the floor over to Gisella Gruber to let us know of the housekeeping of the how things are done on this call. Gisella, you have the floor.

GISELLA GRUBER: Olivier, thank you. Gisella here. The lines have now been muted, unmuted, apologies, on the Adobe Connect as well as on the audio bridge. And if you'd like to ask a question, please do raise your hand in the Adobe Connect room. If you happen to be only on the audio bridge, please do say your name and we will give you the floor.

If I could also remind you to please state your names when speaking, not only for transcript purposes, but to allow our interpreters in Spanish and French to identify you on the other language channel. And if you happen to be on the audio bridge and the Adobe Connect room at the same time, please do remember to mute your microphone and speakers on your computer. Thank you. Over to you Olivier.

OLIVIER CRÉPIN-LEBLOND: Thank you very much Gisella. This is Olivier speaking. And so now I open the floor for questions. So I don't see anyone putting their hand up at the moment. I was going to ask a question with regards to ICANN's productivity in the IPv4 to IPv6 transition. What can ICANN do to promote this transition? And just as an example, should it, for example, basically tell its customers, so customers being both registrars, registries, and IP customers, and also the end users, that it's going to cut off IPv4 and they have to move to IPv6.

This is meant to done quickly.

LEO VEGODA: Okay. So, if I work backwards on your question, you spoke about cutting off IPv4, and that kind of flag day isn't going to work anymore. Once upon a time, there was a flag day on the Internet, and NCP, which I think was network control protocol, was replaced with IP, and basically everyone was meant to do it on the same day.

But there are now literally tens of thousands of autonomously managed networks across the Internet, and lots of devices, and lots of firewalls, and you can't just cut it off. And certainly, some kind of top down command isn't really what ICANN does. What ICANN can do, and has done, is that ICANN can work with its customers to promote IPv6 in some cases, for instance, with a country code top level domains to encourage them to deploy IPv6 accessibility for their name servers, and support IPv6 blue registration for their customers and so on.

And for the contracted parties, the gTLD registries and the gTLD registrars, there is in the new, or the latest is perhaps the best way to describe it, contract a requirement for IPv6 support. So for the registries, they have to have IPv6 accessible name servers, and they have to support all of the regular registry operations over the IPv6.

And for the registrars, they have to make sure that their WHOIS servers are available over IPv6, and they have to support IPv6 blue registration from requests from their customers. And it's possible that people don't know what blue registration is. And essentially, it's like this. If you have a domain name like ICANN.org, and it has a name server like ns.ICANN.org, you can't go and ask ns.ICANN.org where is ICANN.org, because you need to know what ICANN.org is to ask ns.ICANN.org.

So what happens is the parent has a thing called blue, which is basically the IP address in ns.ICANN.org, and that IP address is a hint to go and ask this name server. And that needs to be registered in the registry, and registrars have supported IPv4 blue forever. And IPv6 blue is now available from any registrar that has signed the latest registrar accreditation agreement.

OLIVIER CRÉPIN-LEBLOND: Hello?

LEO VEGODA: Hello?

OLIVIER CRÉPIN-LEBLOND: Sorry, I'm not sure if I heard the end of your sentence. So IPv6 blue is ready for any registrar that has signed the agreement.

LEO VEGODA: Yes. It's supported by any registrar that has signed the latest accreditation agreement, and it's one of things that the ICANN compliance department looks at.

OLIVIER CRÉPIN-LEBLOND: Okay. Thank you very much Leo. So, let's see, we have a hand up in the Adobe Connect room. So we have Kivuva Mwendwa, Mwendwa. I'm not sure which channel you are on. Whether you are on the French or the English channel. But please, you have the floor.

MWENDWA KIVUVA: Hello?

OLIVIER CRÉPIN-LEBLOND: Go ahead Mwendwa.

MWENDWA KIVUVA: Okay. [?]… in Africa is very minimal, apart from [?]… Thank you.

OLIVIER CRÉPIN-LEBLOND: Okay. Thank you very much Kivuva. Have you heard this? Your line, Kivuva, was very faint, but I understood your question. Leo, did you get this?

LEO VEGODA: Yes. I think the question was a statement that East Africa and South Africa have the strongest IPv6 deployment in Africa, and there is very little elsewhere, and what is being done to correct that. And there is quite a lot that is being done from the training perspective. Basically building technical capability in network operators around Africa.

So one of the activities that people who run networks do is they have network operator groups, and within Africa, there is [AfNOG], which is an African wide operated network operator group, and there are also in country and sometimes in region network operator groups. For instance, there is [?] which is the Southern African region network operator group.

I know that Ghana has quite a strong network operator group, and there is more. And these groups have training activities. [AfNOG] makes a habit of once every six months going to a new country, making sure they leave the country with some trained engineers there, and I think they also leave some equipment there.

In addition to that, there is also work for what is done by the NS RC, that's the network startup resource center. And the network startup resource center does similar things with IPv6 capability training, both aimed at network operators and country code top level domain operators, and more. And ICANN has recently brought on a new person to help with this technical training and outreach.

You might be familiar with him, his name is Steven [Contee], he worked at ICANN in the past, he did this sort of thing when he worked at ISOC, and now he is back at ICANN and he is doing more of it. So, if you want, we can put you in touch with Steve [Contee], and he can probably give you a much more detailed answer that speaks directly to the particular place that you live.

OLIVIER CRÉPIN-LEBLOND: Thank you very much Leo. It's Olivier speaking. Thank you for this question Kivuva. Any other questions on this? Certainly as someone who has followed the spread of IPv6 worldwide and also the spread of the Internet in the early 90's, in Africa, an interesting confident in that there appears to be just a few pioneers for several years, and certainly the way it goes across the continent very quickly, certainly I saw that with the spread of the Internet, very quickly, every country was followed by the Internet.

And that happened within just a few months. So it might be with IPv6, one is looking at a similar wave, although this is not guaranteed at this date. I'm not sure if I see any other questions. So no one putting their hand up. We seem to have quite a fast call today. I wanted to thank you very much Leo for this very interesting introduction into the IPv6 and IPv4 world.

And certainly looking forward to seeing more IPv6 across the globe. Certainly your explanation with regards to the network address translation, and some of the costs that appear to be [?] with network address completion, are things that are important for Internet end users since ultimately the costs are paid for by the end user. So, it looks as though IPv6 is the better option that will keep costs low.

So thanks very much again for this clear presentation. The presentation itself will be made available, like all of the presentations in the strategy

– At Large summit capacity building, on the ATLAS II wiki pages. Staff will be providing details of where this is. We've got here, Heidi has already put it in the chat, capacity building webinars ATLAS II.

So you can all download the presentations. And of course, we'll be continuing this discussion when we meet face to face in London. If I could ask what the next week presentation is going to be. I have a feeling I know what the next presentation was going to be. I can see Heidi already putting it in the chat.

The future of Internet governance and it says part one, because both – well, Nigel Hickson and myself will be speaking about this, the future of Internet governance. Nigel, I think, will be speaking the week after, and I shall be speaking on that week. So, that's just after the NetMundial, that will take place in Brazil, which is starting later on this week, on Wednesday.

Several of us are travelling there and that is likely to discuss a lot of different matters of Internet governance. But in my first talk, I think I shall be trying to provide some background on the various aspects of Internet governance and certainly the aspects which effect ICANN directly. Anyway, I look forward to seeing you all next week.

And until then, if you have any further questions for Leo, I'm sure you can send them over to the email, and it will get the…

LEO VEGODA: Okay. Well, thank you very much.

[END OF TRANSCRIPTION]