APNIC IPv6 Deployment

Tue, Feb 28, 2023 4:29PM • 2:10:37

Kenny Huang
00:59
Okay, good afternoon, everyone. And good morning, good evening, to anyone joined online remotely. Welcome to the session. My name is Kenny Huang. Thank you for joining the session. The purpose of the session is try to encourage all operators to share their experience in terms of IPv6 deployment. And so, how do they do… what kinds of problems they experience and what kinds of solution they deploy? What kind of business operation they do? And even, they can share about a customer acquisition strategy as well. So thanks again for joining the session. Can I invite, to all speakers who are on site, can join the stage with me. So, I'd like to invite our first speaker, Ondrej Caletka from RIPE NCC. He's going to introduce deploying IPv6-mostly access networks. Let us welcome Ondrej. Is he online?

Ondrej Caletka
02:01
Yes, I'm here. I hope you can hear me, and I will just share my presentation.

02:10
Yes. Okay, so I'm ready. Thank you very much for having me here in this conference. And I would like to talk about deploying IPv6-mostly access networks, which is sort of a new term, that is there since let's say last year or maybe two. And it's basically, in a nutshell, IPv6-only network and dual-stack network in one network together. How exactly does it work? I will cover it in a moment. But first, let's talk about transition mechanisms because that's what we do if we want to deploy IPv6. And we used to say that the best transition mechanism is dual-stack, which basically means that you have IPv4, just like you had, and then you add IPv6. And in this case, that means basically that IPv6 works completely separately, there is no common infrastructure. If IPv6 is not available, the resources are lowered over IPv4. If there are some problems with IPv6, there is this algorithm called Happy Eyeballs, just that is just masking these problems by switching, firstly, into IPv4. It works very well, with one minor issue, and the issue is that it does not address IPv4 address scarcity at all. By deploying IPv6 with dual-stack, you will not save any IPv4 addresses. So, this will just not help with the fundamental issues of IPv4.
03:49
Therefore, we have another solution that sounds much better. It's called NAT64, which basically means that you can have IPv6-only access networks, and provide IP for only as a service over IPv6 network. In this case, the service is a translation. So basically, you take part of IPv6 address space and translate it into IPv4. And whenever somebody wants to reach IPv4-only resource, they can do it via NAT64. And together with DNS64, which is like a trick with DNS, you can… the end device can have a feeling that everything is accessible over IPv6, so it doesn't have to care about IPv4 at all. It works well, except for some very well-known cases when it doesn't. For instance, if you type IPv4 literal in somewhere, then it of course doesn't use DNS and therefore DNS64 will not work. There is also legacy software which does

not do anything else and in IPv4 it will not work. And there is also a problem with… is everything working? Yeah…….with dual-stack servers, which are broken.

05:03
Anyway this solution is already very well used in the mobile world. So mobile operating system vendors are trying really hard to make it work on their own IPv6-only because the mobile carriers tend to deploy IPv6-only networks with NAT64. So in case of Apple, they did this approach that they basically forced everybody since 2016, to have all apps working properly on NAT64. In the case of Android, it's… they implemented a thing called 464XLAT, which means there is a second translator, called CLAT, which is inside the device. And this second translator basically makes it available, that even end user apps with IPv4-only can work in such a device, even though the network is IPv6-only.

06:03
The problem is that desktop computers don't have such features. So if you have IPv6-only access network and you connect some windows machine for instance, it will not… it will work mostly, but there will be some small programs, that are very well known, most people will not notice them because they are quite small. But for some people this can be like a showstopper. And as I said, like the legacy applications that are using IPv4-only or IPv4 literals, or just…yeah, the problem is with Happy Eyeballs – that algorithm that allows us to quickly switch from IPv6 to IPv4 in case IPv6 is broken – this of course cannot work on a desktop that has only single stack IPv6 connectivity, because there is nowhere to switch to.

06:53
So the question that this talk is about to solve is, can we somehow do IPv6-only at least for some devices, that means devices that are supporting it, like mobile devices and not for the others that would suffer on it? And the solution that already exists – and it's very well-implemented – is an RFC 8925. And it basically uses DHCP, the old protocol that is used for IPv4 address configuration – so this is IPv4 DHCP, not IPv6 DHCP – and there is a new option for this legacy IPv4 DHCP that is called 'IPv6-only preferred' and this option is numbered with number 108. I'm going to use number 108 because the name of the option is very complicated and it gets confusing very quickly.

07:49
So first, how does it work when a client joins the network and regarding DHCP? This is pretty common. So first, the transaction starts with discover message where the client lists what kind of parameters it's interested in. And in case of client that is supporting this option, the client will also request option number 108. The DHCP server will respond with an offer of parameters, and in this case, the server doesn't know anything about this option 108 so it doesn't offer such option to the client. And in this case, the transaction continues normally, so clients will request the parameters and the server will allocate this request and acknowledge it. So in this case, the client asked for running IPv6-only, but the server ignored it, so the client got assigned IPv4 address like normal and everything works like normal on dual-stack or even IPv4-only.

09:01

Then, what if the DHCP server is configured to respond with this option 108? In this case, before the discover message, there is an offer message. And in this offer message, the option 108 is offered together with, sometimes, in this example 30 minutes. And that basically means that the server is aware that the network, where the device is connected, is capable of running IPv6-only. And therefore, it announces to the client that IPv4 does not have to be used for 30 minutes in this case, and the reaction of the client is that, in this case, the client will stop the transaction here. So, the transaction is not finished because there are no four-message exchanged. And it is like there was no DHCP transaction at all, because it was interrupted in the middle of the of the transaction. And in this case, the device will just work on IPv6-only, and the DHCP pool will be one address, one address bigger let's say.

10:08
The question here is, does this protocol get at least some implementation? And the answer here is, yes. Quite a lot actually. It's in all new recent Androids, iOS, and pretty surprisingly, also MAC OS desktop operating system from Apple. Actually, one year ago in Berlin, during RIPE 84, I did a measurement on DHCP service where I basically looked up how many clients are requesting this option 108. And I found out that two-thirds of devices in the network already requested the option 108. Back then, the network was not ready for it, so the devices operated dual-stack normally. But yeah, you can see that quite a decent number of devices are already using it. So let's say, devices are ready; the networks are a little bit lagging behind.

11:07
One more thing is about the macOS, because this is interesting in a way that macOS allows you to run any software. You are not tied to a particular app store, like with iOS or Android. So you can run any software and you can run even software that is, like I said, is broken on IPv6-only networks. And it seems that this has been solved by Apple by implementing CLAT in macOS as well. So that means, if you connect to IPv6-only network with macOS, there is a CLAT engine inside that will get activated. It used to be only activated when the option 108 was present, and also when there was PREF64, prefix option. But this is not the case anymore, starting with macOS 13. It gets activated every time the network has NAT64. And, at the same time, it has this disadvantage, that if the network is pure IPv6-only, that means there is no NAT64. The macOS will refuse to use it. It will basically disconnect and says that this network doesn't work. So there is like now a hard requirement that Apple devices have to be connected to a network that supports IPv4, either natively or via a NAT64.

12:35
Yeah, so we list on the right-hand side, you see this is how it looks when the CLAT is activated. I thought I mentioned briefly the router advertisement option, PREF64, which is also another new thing and basically in case of setting up this CLAT, this translator inside the host machine, the CLAT has to be set up with the NAT64 and somehow it has to discover which IPv6 prefix is used for NAT64. And it seems that the most …like, the best option how to do it is with the new option of router advertisement that is called PREF64. And it's a very simple option that just adds to router advertisement. So whenever the net device receives the network configuration, it will also receive its NAT64 prefix information. And this option used to be required for the previous version of Apple operating systems, in order to activate the CLAT. It's not the case anymore, but chances are that it's going to happen, because this is a little bit more trusted than other ways of discovering a NAT64 prefix.

13:44
Okay, let's now talk about if you perhaps want to give this network a try, how exactly you could run such a network. First of all, you need DHCP server that supports the option 108. And surprisingly this is very easy, because first of all, there is native support in Kea from ISC. So in this case, you can just configure it, it's easy. But even if you use a different DHCP server, as long as the DHCP server supports custom options, you can create option 108 by defining a custom option, because there is no requirement for any special handling of this option. So, any like 20-year-old DHCP server that supports custom option can do it. There is an example of a DNS mask. Yeah, the only thing is that you still have to have addresses in pool of IPv4, because as soon as the pool gets depleted, the service will stop responding with the offer. So even though the address is not needed at the end, it will not work without at least one free address in the DHCP pool.

14:00
Regarding the option PREF64, it's a little bit harder, because in general, routers don't support custom RA options. Every time there is a new option standardized for router advertisement, it requires firmware update for that particular router in order to support it. So this is really painful. We already had this issue with recursive DNS server option, which was already put there later. And now we have the same problem again, and we probably are going to have it again. So, I would really like to see vendors picking up with some standard of customer options for router advertisements, so a new option can be used without requiring complicated software upgrade.

15:50
There is some adoption. There are some patches for software routers or pull requests. There is also better MikroTik RouterOS added support to this option. So it is slowly happening, but unfortunately, there are still many, many routers that don't support it.

16:14
Also, if you run it on macOS, there are some interesting, let's say, features or maybe even bad. First of all, there is this thing that if your network has more than one network prefix, the CLAT engine will choose one basically randomly. So in my case, there is a ULA prefix and there's a global unicast prefix, and the CLAT engine picked up the ULA prefix to be used for the CLAT. So, this probably is an expected behavior and something, yeah, that this probably should be… there should be probably the global unicast for CLAT to use, but in the end, most of the time it doesn't really matter because probably most of the networks use only one prefix there.

17:08
Another interesting thing is that the CLAT engine in macOS does not work for all applications. For instance, it doesn't work for DNS resolver. So, if you have custom settings of DNS server and you set up IPv4 addresses for DNS server, it will not work. Despite that you can ping that address. You can even use commands like 'host' that will resolve using that address. But if you try to connect to some domain name, it will tell you that 'unknown host', that DNS doesn't work, because the DNS resolver itself does not work over the CLAT engine, so cannot work with IPv4 literals.

17:55
Yeah, so that's basically it. If I wrap it up, this approach of having both IPv6-only and a dual-stack in one network, the main advantage is that there are no two networks, only one. So people are not confused which network to join. And there will be no… you know, with the growing number of mobile devices, there will be no growing number of consumption of your DHCP pools, or you have to expand your address ranges because of that. Because most of the new devices will just work without IPv4. This is especially good if you don't use NAT for IPv4 and you are using public addresses, then you have a very big problem with more and more devices. Also, because the dual-stack uses DNS64, it means that even in dual-stack, the use of IPv4 would be minimal. Most of the traffic will go translated by NAT64. So this is like on the positive side.

19:02
On the negative, this network is the most complex network because it has everything from dual-stack network, plus everything from IPv6-only network. So, you still need to deploy IPv4 everywhere. So there is no like direct saving of IPv4 addresses because you still have to have it deployed for the devices that don't support IPv6-only. You also need NAT64, so in case you don't have any NAT solution, you still have to deploy NAT64. And there are some problems with interoperability. If you have one device that is running dual-stack and others running IPv6-only, chances are that they just cannot talk to each other. In my case, it was if I tried to set up a Chromecast from an Android phone. So Android can work on IPv6-only, Chromecast works on dual-stack and they just were not able to see each other.

19:56
So, to wrap it up. When you should consider IPv6-mostly? I came up with these options. So basically if you don't use NAT and your DHCP pool is filling up, or if you do use NAT but you are even running out of private addresses because you are a very big network, if you are a network that has mostly Apple devices or mobile devices and if you already have NAT64 and would like to be producing more, and if you would maybe even consider two un-deploy IPv4. So these are the most options when this makes sense in my opinion. And because with the right meeting network that we do for our RIPE meetings, we basically check most of these buttons here. We have.. we don't use NAT and our DHCP pool is filling up. There are mostly mobile devices there. and we already have NAT64. We tried and we deployed it during our last meeting 85. And we actually did also deploy other networks and the numbers were pretty good. So in the end, 74% of devices were using IPv6-only. The biggest issue we found out was that MAC users have set up custom DNS servers. Therefore they had no DNS and therefore nothing worked for them. Or even, we found some devices that had this disabled IPv6, and of course if the network has disabled IPv6 and the IPv4 is disabled by DHCP option, then there is nothing to connect to.

21:34
We saw still some Apple users that connected to legacy network instead of the IPv6-mostly network. And the only thing that we observed as an issue, was that Cisco AnyConnect or even OpenConnect to open [unclear], didn't want to connect. But on the positive side, something that I was really worried about was our printer. That is a network printer, which supports Apple AirPrint over dual-stack, and it indeed worked very well even with IPv6-only devices. It printed without any issues.

22:09

So I think this is like very interesting standard and I would really… I would really like you to consider this maybe first as an experiment. Or maybe it could be the way forward on how to finally get rid of dependency on IPv4. And with that, I'm finished and if you have any questions, feel free to ask.

Kenny Huang
22:34
Okay, thank you, Ondrej. And any other question from the floor? Okay, any question from online? Okay, Hearing none. Thank you, Ondrej. And actually, we do prepare some valuable gift to speaker but unfortunately you are not here. So I'll probably just keep it for you and you can collect it from me next time, if you can find me. *(laughs)*

Ondrej
23:00
Okay. Thank you.

Kenny Huang
23:02
Welcome the next speaker, Kams from Akamai, a CDM prospect provider prospective IPv6 in 2023. Let's welcome Kams. Thank you.

Kams Yeung
23:16
Hello, can you hear me? Okay. Thank you. So, I'm Kams, one of the network architects in Akamai. So today I'm going to talk about from a CDN perspective, what we observe on the IPv6. So this is more about the chance data that we observe. Do I have a clicker or…?

Kenny Huang
23:41
Yes, sorry. Sorry.

Kams
23:47
Yeah, so I just wanted to talk about a brief history about Akamai. IPv6, so we have been doing IPv6 since 2011, so the first time that we implemented this. And now we have been deploying IPv6, not only on the server [unclear]. We also do a lot of IPv6 like internally for our security, for our compute, and also of our other products, DNS. So we also support a full stack of IPv6 as far as possible. And we have been always encouraging our customers to do dual-stack since a long time ago, and nowadays we do dual-stack by default. However, there are also some customers that remain on IPv4-only for their own reason. So some high-level statistics is, we are doing Peak IPv6 traffic now, for the whole network about 47T, versus total dual-stack IPv4 plus IPv6, is about 250 Terabit per second. So this is like 1/5 of our traffic is now about IPv6. And we have IPv6 deployment in like over 120 countries and 650 cities.

25:02

And this is proof that you can see the trend over the last 10 years. So these are statistics we collected last year. So you see that in 2012, this is the IPv6 launch. Of course, we are one of the members where [unclear]. So the peak IPv6 is barely just 1Gig, because we only have like maybe a handful of 1000 hostname that v6 enables. And at that time, 10 years ago, of course there's a lot… not a lot of IPv6 devices at the time, as you can imagine. But still, we got mobile feedback on that period of requests per day. And IPv6, just observe I mean, the end user address that we observe is about 19 million. So that is like the IPv6 population 10 years ago. Then you fast forward like looking after 10 years later, that is last year. We are doing a peak traffic of 41T. And the daily requests nowadays coming from our CDN servers have become 4 trillion per day. And the IPv6 addresses observed, that means kind of like the IPv6 population, is like 7.5 billion. So you can see that across this 7.5 billion, of course there's a lot of different [unclear]. So we've looked into 64 [unclear], we observed like 2.2 billion – which is huge terms, exponential growth I mean, if you compare the last 10 years.

26:30
There's a lot of driving factors on that. So we're going to look into why we see this traffic? Number one is IPv6, I think everyone knows, it's definitely more scalable in terms of the address space. So a lot of ISPs and carriers, when they have to grow their network, especially like mobile operators, definitely they have to shift to IPv6. Of course there's some mobile where nowadays they do IPv6-only. And nowadays, you can see them more and more, IT computer applications are moving to cloud. And if you want to speed up a [unclear], and when you get up so many machines on the cloud, then you need IP addressed for them, right? So, you can imagine that cloud is getting more and more popular. An initial [unclear] the machine, a virtual machine, on the cloud needs an IP address. So that's… if you don't have a [unclear] enabled, you can't get into a very large scale deployment. So that's also one reason.

27:27
And IPv6 definitely simplifies the address management. Even if you do perfect IP, you have to do NAT, or NAT and NAT… carrier NAT. So there's a lot of management that you have to do if you do [unclear] your space.

27:44
If you do see what's your security, the IPv6 space is also easier compared to before. And if you go direct to v6, you'll find that it's easier to manage because you don't need to manage dual-stacks before v6. You just manage v6. Of course, externally, you still need some translation. I think the previous speaker talked about that. And IPv6 also enables you to have other new surfaces that you can create, in terms of the [unclear] DB for restriction. Of course, v6 actually also have a better performance, so we have a lot of statistics showing that, because IPv6 will allow you to bypass the carrier [unclear] or translation device. So that actually have a better performance, when you compare to before v6.

28:34
Of course, there are also a lot of driving factors, because a lot of governments are trying to have policies that enable, or at least from the government department, they have to enable v6. So we see a lot of governments have these kind of mandatory requirements. And some of course, even push down to even become a national goal, if you read those government news.

28:59
So he has some changes that we observed for the IPv6 adoption. We only selected the top countries here, that if they have an IPv6 adoption over 40%. So IPv6 adoption in Akamai means the number of IPv6 requests that we observe, versus the total. So if you see from a country perspective, you will see the requests coming for v6. We see that there are 15 countries that they have IPv6 adoption more than 40%. So this is the list of countries here, like Belgium, Brazil, France, etc. And from a network perspective, they are half of… it's about 50 networks that have a IPv6 adoption of over 40%. And you can see the growth in trend here. So I've just listed the top ten countries here, that they have an IPv6 adoption over… have a high growth of that.

30:09
Next, [*warped audio,* unclear]. This is a network view of IPv6 adoption, but with [*warped audio,* unclear] networks. So you can see that about approximately 50 networks, they will have – about 80% of them – have v6 enabled. If you look into 800 networks, the top 800 networks in the world, about 20% of those have 40% adoption. So, you can see that we see a lot of large networks, they tend to have a higher v6 adoption, because the last level usually are those large mobile operators. And mobile operators have more need for IPv6 addresses. So that's why the adoption is here. You can see this here. And, if you zoom into Asia, so these are the Asian countries, of the percentage of v6 that we observe. They have an IPv6 adoption rate of over 5%. So, Asia varies country by country. But you can see that some are very high, like India, or like Malaysia. So they are like 59, 58%. There are also some countries that are catching up, so like Pakistan, and Singapore and Mongolia. So they are catching up on that. So it's really varied country by country that.

31:48
And now if you look into… we zoom now into the Philippines, since we are here. So this is the, what we call SOTI. SOTI is like the State Of The Internet report. You can check it on our website as well. So we measure the rates of adoption in a map format, that you can click into. Like, oh this country have a higher provider… a higher v6 adoption. So we found that Philippine – this data I think we just captured this morning – is 15.7%. But this is talking about the whole Philippine user. And then, if you look into the trend for the Philippines, they have been starting quite a big jump since like mid-2020 [unclear]. And they wish to… now, it's about 15.7%. And in terms of country ranking, they are 61st in terms of the ranking. It's not that bad, but they could be better. And the way that we observe the IPv6's adoption is by IPv6 request versus total requests. There are other measurements, right, so we can benchmark with Google, and APNIC. They also do their own IPv6 measurements. So by Google, it's measured by… for the Philippines, they see 16.91%. And APNIC, we see 13.28%. Of course, this is gonna be measured differently by different measuring mechanism, but you can have a reference about the Philippine IPv6.

33:21
And the major v6 in the Philippines that we observe is Dito Telecommunity. That's because, you know, they are an emerging new mobile operator. So, I mentioned that mobile operators, they have a more urgent need for v6 because they're new, so they can't get much IPv4. So that's why when they deploy, they have to be more reliant on IPv6. And next after Dito, would be PLDT. PLDT, of course they also have a mobile smart. So that's why they also have a very high v6. So you can see the provider data again

on our website, so I don't list them all here. But this just gives you a bird's eye view on how v6 is doing in the Philippines.

34:11
And coming back to Akamai, we have been doing dual-stack content since 2011. So you can see that the number of hostnames… the number of hostnames I mean, for the top 100 customers, so the top customers, they have 1 million requests per day. So we are looking into how many user content that we are serving for the top customers. So you can see the trend is like, on 2015, it's not that much of our content, customer content, on dual-stack. But now you see the trend, because on-and-off, we have new customers, right? And by default, new customers will be defaulted on dual-stack. If it's a legacy customer, they are on v4-only. So the new customer has to agree for us to migrate them to dual-stack. So this trend is, number one, it's because of new customers coming up with default v6… default with dual-stack. And the other too, is that we are also converting some of the existing legacy customers, because before, only to dual-stack. So we expect that this will keep going up. So will allow more and more Akamai content to be on dual-stack.

35:32
So, [*warped audio,* unclear]. All Akamai content to date, are dual-stack, because of different reasons. So number one is, when we start to do with dual-stack default, there are some legacy customer that they don't move and we can't force them to move. And so they can have an option to, "I only want v4-only". It's up to the customer's choice. We provide that flexibility. But then some customers have an impression that if they turn on v6, then there's a performance issue, because they think v6 adoption is still not that high and they think that we have a performance issue. But that's actually not the case. So that part, I think we need more education on that.

36:18
And of course, there are also some customers, their back end system can only process before. Like the lock system, onboard the application. So they can't support v6 trend. This means that we have to keep their content only on v4. So this also relates to how customer integrate in their back end.

36:37
So, I just want to repeat again why we need to… why we are moving to v6. Number one is v6 is much more scaled, much more address space. It's enabled the carrier network and the cloud provider to grow. It's simplified all the address management, so we get before for the NAD, carrier NAD. It also provides you more flexibility to create a new surface. And some government regulation policies were also mandatory to enable v6. So these are the driving forces of why we, at Akamai, have been pushing so hard to v6.

37:18
And that's all my presentation. If you want to learn more about the Akamai IPv6, you can go to this website. And this slide actually, I got most of the material from my colleague, Erik Nygren. So he kept a blog about … a bit much more detailed blog about what we have been observing over the last 10 years. Not only about specifically Akamai, but just also like what has been happening in the IPv6 world. So you

can look into that link for more information. And if you have any questions, comments, you can ask me now or you can send me an email. So my email is there. I think that's all my presentation.

Kenny Huang
37:58
Okay, thank you Kams. Any question from the floor? And say your name, organization Thank you.

Jan Zorz
38:09
Yes, it works. Hello, Jan Zorz from 6connect. Do you maybe have any recent graphs or measurements or numbers, on how much does the performance of reaching the content over IPv6, increase? Because I know, Akamai publicly shared some numbers years ago, but I didn't see any recent numbers or measurements. Do you have any newer information?

Kams Yeung
38:43
I think I could send some new information to you, but the performance difference between v4 and v6, I just don't have this off my head, but I can follow up with you on that. There is not… the rate is not a very big difference. But this does depend on the network. So let's say if the network, if the carrier have a IPv6 server to talk into a v6 [unclear], the performance will be better than a v4 server going for a translation and talking to a v6 user, for example. Or talking to a v4 user, but via NAT. So because sometimes we see there could be a performance bottleneck on the carrier grade NAT. So it's not like generic, sometimes it could be that different networks have different setups. So we will find that we will try to talk with the provider to see whether we can bypass the carrier NAT, because sometimes you just forgot to enable v6 for example. So we have to go through the NAT. So in general, I think if you could do direct v6-to-v6 server to client, would be always better than any translation or NAT proxy between them. That's the performance trend that we observed.

Jan Zorz
39:55
Yeah, because I remember Facebook saying that they measured even 30% of better or faster reachability of the content over IPv6 for some networks. And I'm wondering if you're seeing some numbers like this

Kams Yeung
40:13
I don't think it's that drastic, I mean, from our measurements, okay. It's just more like you go to less players. So it's more like, for example, latency or potential carrier grade NAT device capacity that causes the problem. So causes, like more like, delay on there. So it's not… I would just say 30%, but I don't know where this metadata has come from. But the data that we have is just a bit more subtle. So that's why… yeah, it's better.

Jan Zorz
40:47
Thank you.

Kenny Huang
40:48
Okay, next?

Jordi Palet
40:49
Jordi Palet. I think I have a point about the question from Jan. I recall. Eric Nygren presented already similar measurements as the ones done by Facebook, and it goes like 30% improvement or 30-something. So probably it's online, we can find it.

Kenny Huang
41:10
Yeah, thank you. Okay, any other question online? Okay, hearing none. Thank you Kams. Before you leave…I give you.

Kams Yeung
41:25
Thank you.

Kenny Huang
41:29
Okay, moving on. Let's see, my next speaker, Adelbert Eslava. He's from Globe Telecom. He is going to introduce Globe Telecom IPv6 deployment. Let's welcome him.

Adelbert Eslava
41:51
All right. Good afternoon, everyone. I'm Adelbert Eslava or Al, you can call me Al, from Globe Telecoms. This will be just a short presentation. It's not necessarily a Globe entirely adoption of IPv6, but it's an operations perspective as I see it, from where I sit in Globe – in the operations department, okay.

42:16
How it started with Globe. Around 2011, I understand there was a global IPv6 connectivity that was being presented and Globe was a part of the world IPv6 day. As soon as July, I think, of 2011, we started to configure IPv6 mainly on our core routers, as it was… where we think is the best position for us to start the IPv6 journey before deploying it anywhere into our customer edges or to customer networks and services.

42:57
The current status right now, as you can see here, we have the snapshot from bgp.he.net. We see that we already have IPv6 prefixes being announcing in to the network and we also have IPv6 peers. We started to grow the network. On the right-side graph, we can see that we have the APNIC statistics on the start of the announcement of IPv6 addresses from our network, AS4775. Okay?

43:31
So on the deployment of IPv6, the question that is always being asked of us networks, is are we there yet? Have we already fully transitioned into an IPv6 network? And that begs the question, when do we really say that we have already reached an IPv6 transition? Is it when there are no more IPv4 addresses to be allocated to customers? Or is it when all services are already serviced by IPv6 and not by IPv4? And for our network, we can say that, no, that's not yet the case. We still have IPv4 in the network, definitely. We are only running dual-stack right now on our networks. The start of the deployment actually for the customers started with our mobile networks, because that's where we feel that we can easily deploy and have the biggest impact into the announcement of IPv6 addresses into the network.

44:34
After we deploy it into the mobile network, as you all know, we also have the broadband networks. We also have deployed IPv6 already on broadband network, but again, on a dual-stack topology currently right now. Right?

44:53
So this is the current status. We have IPv6 is the network coexisting with the IPv4 network.

45:03
Some early challenges, I might say for the deployment of IPv6 in any networks, is again, for the operations perspective, we do have to look after how we maintain those networks, right? It is important that we are able to support whatever network we put into the whatever technology, or new segments that we put in. So we feel that competency development is one of the challenges we found in our early stages. It's fortunate also that we received a lot of help from our friends from APNIC. They have come here for some workshops and some training materials and training our team for IPv6 adoption from 2011. But they have been working with us most definitely with the local NOGs, the PhNOG, where they conduct workshops for us.

46:06
For me, the next challenge would be compatibility of course. Globe has been a well-established telecommunications company before IPv6, and we feel that to be able to transition, we have to look after all the network elements that are already in the network. Making sure that they are all able to support IPv6 when we do decide to deploy them, right? So there is a little bit of hesitancy from management, as well as from the operations department as well, just making sure that we do not break the network when we do deploy IPv6.

46:52
Of course, there are other equipments that needs to be considered, all the way from the routers to the CPEs, as well as other back end systems that we have in the network. We, like I said previously, we only deployed it on the core networks, but then if we have to go into a full deployment of IPv6 into the network, we also have to consider the back end systems. We have the billing systems or else we break the service flows into our network.

47:26

I wrote here also, number 3 is customer pushback. Customer pushback, is mostly also related to number 1 and number 2. Of course, our customers also have legacy equipment in their networks. That has to be able to support IPv6 if we do push for an IPv6 native connectivity. That is the reason why we still have dual-stack configurations in the network. And still, there are enterprise customers mostly who still want to deploy IPv4. They are still requesting us to provide them with IPv4 upgrades, even though the network is already as a dual-stack. The preference is still at IPv4. Again, on this area, we think our team from APNIC has been also trying to convince or to educate our customers into transitioning fully into IPv6 and to get with that transition with us.

48:30
And the last one, like, what would be the challenges of network monitoring, right? From an operations perspective, again, we [unclear] that we have, to be able to maintain the network at this performance, that we'd be able to support our services, right? So in terms of monitoring, we have to make sure that we have the necessary tools to monitor traffic. Because we understand that the network status would have a great impact in the application performances that are being served into our network. If we are unable to say monitor this traffic and we do not have visibility from an operations or from the NOG department, then we say that we can be blinded with the actual status of the network and we might be providing sub-optimal services to our customers. Also part of the network monitoring that we try to consider when we tried to transition to IPv6, was that we have to be able to have configuration management visibility as well. As you know, when we have dual-stack, we are almost like running two networks at the same time, one IPv4 and one IPv6. And as long as these two coexist in the network, we have to make sure that we have visibility to those two networks and we have the ability to maintain and operate those systems efficiently.

50:05
So, like I said, it will be a short presentation. It's just a perspective from operations, on how we deploy IPv6 on our network from Globe. Thank you.

Kenny Huang
50:14
Thank you, Adelbert. Any question from the floor? Any question? Or any questions online? Okay, hearing none, thank you. Adelbert.

Adelbert
50:26
Thank you.

Kenny Huang
50:27
Before you leave, there is a gift.

Adelbert
50:31
Thank you.

Kenny Huang

50:35

Yeah, you can stay on stage. So let's welcome the next speaker, Mohammad Kimhar Junaide. He's from Dito Telecommunity. He is going to introduce Dito Telecommunity IPv6 deployment. Let's welcome him. Thank you.

Mohammad Kimhar Junaide

51:05

Hello, everyone. So, I'm Kimhar Junaide from Dito Telecommunity. So, I'm here to present the Dito IPv6 deployment.

51:25

So just a brief introduction of who we are. So way back 2018, our government requested to have another mobile operator here in the Philippines. So they have conducted bidding and Dito won that. So by then, we were named the third major telecommunications provider here in the country. However, of course, since we are new, we need to build our network from scratch. We just had commercially launched last March 2021. And we were known as ASN 139831.

52:27

Okay. So why Dito pushed for IPv6? So as Kams mentioned a while ago, we are a new mobile operator. So, we were only allocated a small chunk of IPv4. And we aggressively pushed for IPv6 for us to serve the population here in the Philippines, as our commitment is to cover 85% of the different islands in the Philippines. As you know, Philippines has three major islands, Luzon, Visayas, and Mindanao with a population of over 100 million. So a lot of user subscribers to cover. Of course, as IPv4 is getting depleted, so we need to ensure, since we are the new mobile operator, we need to ensure that our network is future proofed. So we need to ready our network for the upcoming challenges of IPv4, and since the technology is fastly emerging, so in the near future, we believe that we will be getting rid of IPv4 and moving to IPv6 – at least by the time our network is ready to support such demand.

54:04

So what are the challenges we have faced during our deployment? Of course, from our perspective, since we are the new mobile operator on highly technical perspective, we already considered IPv6. Yes, because of the wide range of IP address we need. However, one of the challenges is the high investment and the equipment. So not only on our core equipment, but also on the monitoring equipment. This is for us to gauge, what is the common traffic passing through our network? Is it IPv4 or IPv6? And then, another challenge we face is the technical experts. So we need to find experts for us to be able to deploy IPv6 smoothly. And of course, we need technical experts who have already experienced in deploying such protocol, because we need to catch up with the demand of our government. So we need technical experts. And after the implementation also, we need… there's another challenge and that is additional trainings for our operations team. This is to ensure that they are able to support what we have in our network.

55:38

So since we are a mobile operator, our effort is re-configure dual-stack configuration. So that's the [unclear] for us. Of course, since dual-stack, as mentioned also by my co speakers here, there are other contents that are still not available through IPv6. So that's why we consider the dual-stack configuration.

56:18
And also, of course, a lot of consumers users right now, they do not mind what's behind the technology, right? They do not mind if we deploy IPv6 or IPv4. So what's important to them, is they have a fast, reliable connectivity to view those contents.

56:48
So the statistics from APNIC here shows, as you can see, since the beginning of our launching in March 2021, our IPv6 capability is already at 60%... 60 to 70% up to now. So hopefully soon, we'll be able to reach about 90 to 100%. But the question is when? When will that happen? Especially here in the Philippines. You know, there's a slow deployment of IPv6 here in our country.

57:36
So in conclusion, our story showcases our commitment to provide our customers with the possible best… the best possible network experience, while also preparing our network for the future So, despite the challenges that we face, [*warped audio,* unclear] one-time investment, and we see the benefits of IPv6 far more than the investment that we made. So, IPv6, they say it's a future protocol. But no, IPv6 is already a protocol for today. So, of course, it's not 100% perfect protocol, but we have experts, researchers, that can improve a lot of its features, the security, on how we can provide better performance by deploying IPv6, and of course, to accommodate the number of users worldwide. Because as you know, a lot of users, a lot of devices, are very high tech. So a lot of devices are being deployed. So the number of IP addresses, IPv4 addresses, is getting thinner and thinner. So we need a new one. And that's, we believe, IPv6.

59:38
So [*warped audio,* unclear] towards building a better and more secure internet. And we hope that our experience, what we share here, will encourage organizations to enable IPv6 within their network. So that's our short story on how Dito come up with IPv6 deployment.

Kenny Huang
1:00:09
Okay, thank you Mohammad. I believe most of the operators have very similar challenges as you had. And any question from the floor? Or any questions online? Okay, hearing none. I would like to thank you Mohammad again, and there's a gift for you.

Mohammad
1:00:36
Thank you very much.

Kenny Huang
1:00:36

Okay, that's pretty much the end of the session. As I mentioned earlier, basically this session is to encourage the operators sharing… any questions? On the front row. Okay, your name and organization again.

Dave Dawson

1:00:49

Hi, my name is Dave Dawson. I'm with the APNIC Foundation. I just wanted to take this opportunity to let everyone know that our IPv6 deployment grants are open for applications. They're open to pretty much any kind of organization that has some concrete plans that they're working on, to deploy IPv6 for community benefit. So please check out our website ISIF.asia or if you see me around the conference, feel free to come and say hi and I'm happy to chat with you about them. Just wanted to let everyone know while I've got the chance. Thanks.

Kenny Huang

1:01:18

Okay, wonderful. So you're welcome to submit your application for the grant, and specifically for IPv6 deployment. And as I mentioned, the session is specific to encourage the operators to share their experience and challenges they experienced and what kind of technology or what kind of solution to deploy. And also try to find a way to have new customer acquisitions as well. So, thanks again to all operators willing to contribute your experience and your knowledge, to the floor. And before I end this session, can we give a big hand to our speakers here?

1:01:59

Okay, thanks again for your participation. I conclude this session adjourned here. Thank you.

END OF TRANSCRIPT.