OLIVER CREPIN-LEBLOND: Thank you very much. So welcome everyone to the IPV6 webinar for ALAC, ALAC and At Large actually, for our community. We have the chance today of having a number of people join us with guest speakers Leo Vegoda, and also Jason Schiller, who will be speaking to us about a number of things today. But first we have Leo who will be speaking to us about a short introduction about IPV6 – what it is, what does it enable.

I think that must of our community know, but it's always good to have a good reminder of what it is. Leo Vegoda is the operational excellence manager of ICANN. Next we'll have, taking part in policy development in the five regional Internet registry regions.

As you know, the ASO does not conduct its policy development within ICANN, it's conducted outside. And there is somehow, I think, a lot of… There are questions in our member's mind as to how does one get involved and, in fact is one able to get involved with policy development as far as the numbers are concerned.

Ricardo Patara and Jason Schiller, who are the ASO representatives, will be able to help us on this. And then after that we will have a little update on the IPV4 address recovery policies. They are in some cases different in some regions, and some are not. In fact, I seem to recall that there was some question about being able to have an uniform policy across the world. We'll know all about this, whether there is or whether there isn't, again, from Ricardo and from Jason.

And finally, IPV6 allocation policies in the GSO five regional Internet registry regions. Policies are sometimes different, sometimes the same. It would be good to have a good update on that. Finally, questions and answers. We might be running a little bit over time because I understand that the presentations are going to be quite extensive.

So without any further ado, let's go to our agenda [?] the short introduction about IPV6. Leo, you have the floor.

LEO VEGODA: Thank you very much Olivier. So, I'm scrolling down slides, I don't know if everybody can see that. But for the overview of this, I'm starting with what is Internet protocol and I wanted to do that just to make sure that everyone has a good basic understanding of where we're coming from.

Going through IPV4, the design and development of IPV6, and then where we are with current deployment, and some things that At Large might want to consider for getting involved. So firstly, Internet protocol is basically about sending messages; you break them up into small pieces and then you reassemble them at the other end.

One analogy [Vince] sometimes uses is it is a bit like writing on postcards, but sometimes you send one postcard and you'll get two at the other end because the way the network works is that sometimes it would prefer to deliver two copies if it doesn't get an acknowledgement that the first copy was successfully delivered rather than not deliver anything.

So the other similarity between Internet protocol and postcards is that as Internet protocol is, as a default, there is low encryption or security built into it, at least in IPV4, and that was something that was added at a later date. Of course, the topography… Public key topography was only published as a concept in 1977 while Internet protocol was being developed, and it wasn't computationally possible to deliver that kind of IP sec security in the initial IPV4, which, as you can see, only really started in 1983.

There have been several sort of little revisions to IPV4 as a protocol. So initially, there were just going to be 256 networks, and each of those networks would be about 16 million addresses. And that was considered fine for an experiment because, remember, the Internet protocol and the Internet as we now know it, was started as an experiment.

It wasn't expected that, it wasn't expected that the experiment was to have the tremendous success it has, and that's one of the reasons that there are just over four billion IPV4 addresses. It wasn't considered important to have an experiment that would provide enough addresses for everyone in the world.

Now, a couple of years after the experiment started, they said, "We've got a little bit of a success problem." And they've introduced some additional technology, the class four concept, class A, class B, class C, and different portions of the address space were set aside as having different sizes of networks. So there were 127 of these class A's which have 16 million addresses in them.

And then another portion of the address space was set aside for class B networks, which would have 65,000 addresses. And another portion was set aside for class C networks that would have 256 addresses. Basically, small, medium, and large. And that worked well for a while, but by 1993 it was recognized that even this enhancement wasn't going to allow IPV4 of the every growing Internet.

And so some classless inter domain routine was introduced and that allows any size of network prefix, which is based on a power of two. So you can have one address or two, or four, or 16, or 32, and so on, in your network. So that's IPV4 and we've come up to the mid 90's.

And if I scroll forward into deploying IPV6, you can see at the same time that cider was being deployed, there was a protocol design competition for the next generation of Internet protocol because people recognized that IPV4 wasn't big enough. And the selection was made and standardized by 1995.

And six bone, which was the experimental test bed network for IPV6 began operation in 1996. ICANN IANA department actually made the first IPV6 allocations in 1999, and the six bone ended on the sixth of June 2006, so that was 6/6/6. And that was sort of like the overview of the design and deployment history up until the point where we have production.

So we're now at the point where the vast majority of the IPV4 address space has been fully allocated. Everyone who wants IPV6 address space can get their IPV6 address space and start deploying it on production networks. And about a year ago, the Internet Society led the world IPV6

deployment, and if you go and look at the picture on this slide, you'll go and see that some of networks here have well move half of their traffic coming out over IPV6 and some significantly more than that up towards 90%.

Now, what this tell us, and admittedly these are the most highly deployed IPV6 networks, these results tell us that IPV6 is now a deployable production protocol and you can get access to all of your cool content with that. There are maps of the deployment from Cisco, from Google, from APNIC, and they all use slightly different methodologies, but they all go and tell you basically the same kind of thing.

Now, it's important that if you do buy and deploy network equipment that you don't turn off IPV6, apart from anything else, you might damage the security of your network. Christian [?] from Microsoft made this statement last week saying that Microsoft doesn't test Windows without IPV6 turned on.

And if you go and disable IPV6, well who knows what results you might get. You'll also have seen that the Microsoft's new gaming consul relies on IPV6 with IP sec to work, and so if you don't have an IPV6 capable network, it will go and use a transition technology, a tunneling technology, which isn't quite as good as native.

This is a consul which is designed for the long term and IPV6 is the future, so it relies on IPV6. Now, if you as At Large are thinking, "What should I be doing? What can I do?" You might want to consider engaging with Internet service providers, with regulators, and with over

the top service providers, and asking will the features that you use on the Internet, gaming, video, voice, maps, and so on work?

Will audit trails work? So for instance, if you are involved in any kind of financial transaction, and you have to keep an audit trail of where stuff comes from, will that functionality work with IPV6 as well as it does with IPV4? So on. I know that on the right meeting this week they've been testing functionality on an IPV6 only network and found a significant amount of stuff that works seamlessly, and a few things that don't work on an IPV6 network.

And especially when you are a service that runs over a network, it's important that you have IPV6 functionality built into your application, because without that functionality, your customers might no longer be able to be your customers, you won't have control over what network your customers are to use.

So I'll try not to take too much time. I hope my introduction has been useful, and I'm happy to take questions at any time, but I should hand back now. So thank you very much.

OLIVIER CREPIN-LEBLOND: Thank you very much Leo, that's very helpful and very concise, and thanks so much for sharing with us. Will this presentation be available to our members to be able to download?

LEO VEGODA: Yes. I believe it will be available both from the At Large site and also on the IANA site.

HEIDI ULLRICH: Oliver, this is Heidi. He can also send it out to our list so all of your presentations on today's call, we'll send the presentations and the recordings.

OLIVIER CREPIN-LEBLOND: The link to the presentation recording, yeah, that's great, a perfect thing. Olivier speaking. So the floor is open for questions. Well, I think the first question was actually a question from Glenn McKnight, and that was exactly the questions which I addressed quite independently.

So that's fine. With regards to questions, I see Yaovi has put his hand up. Yaovi Atohoun, you have the floor.

YAOVI ATOHOUN: Thank you Olivier. Can you hear me?

OLIVIER CREPIN-LEBLOND: Yes we can hear you very well, go ahead.

YAOVI ATOHOUN: I just have one question. Because [?] talking about IPV4 translation to IPV6, so my key question is to know if [?] IPV4 network will stop working because why? We're talking about, people are talking about transitioning to IPV6 from [?] and I think both networks will continue working talking to each of them. Thank you.

OLIVIER CREPIN-LEBLOND: Leo?

LEO VEGODA: I think I should punt that question to someone from the ASO who actually runs the network like Jason or Ricardo or Louie Lee because I don't actually run a network, and I think it would be unfair for me to answer the question that they're going to be authoritative on.

JASON SCHILLER: Jason Shiller. I'm with Google, I'll go ahead and try to take a stab at this. I also was previously with U-Net and Verizon as a transit provider. Basically for networks that continue to grow, and their need for addressing continues to grow, they're going to eventually find difficulty in getting additional IPV4 space.

So their options are going to be either stop growing, take IPV4 address away from low revenue generating customers to make it available to high revenue generating customers, or to simply do IPV6. Initially what we're hoping in terms of the transition is to get all the Internet to dual stack so that means all of the servers and all of the clients speaking both IPV4 and IPV6.

And that would allow, if there are IPV6 only networks to showing up in the near future, that they can get to all of the content and all of the peer to peer networks that they need. Eventually what you'll start to see is new content will show up on the network and that might be IPV6

only because as people turn up new servers, they might not have the four addresses readily available.

So eventually, if you're IPV4 only, you'll start to no longer be able to reach certain sites, certain new sites. But ultimately, we don't think people are going to turn off IPV4 in any short time period. There is not going to be some legacy invented system that still will only speak IPV4, think about your HP direct jet card for example.

So likely what's going to happen is people will just transition all of the new growth to IPV6. At some point, when a small, small, small amount of traffic is before, then providers might consider the cost of continuing to support V4 in their router configurations and whatnot, and considering turning it off, but that's a long, long time away.

OLIVIER CREPIN-LEBLOND: Thank you very much Jason. Some people say 50 years, is this real? Considering the Internet itself, and the [?] IP is no more than 35 years.

JASON SCHILLER: Yeah. I mean, I've heard estimations between 30 and 50 years before people just completely turn off IPV4 and I don't find that unreasonable.

OLIVIER CREPIN-LEBLOND: Wow.

JASON SCHILLER: In terms of… The more important question is, when will all of the Internet be able to speak IPV6? And that could be as soon as a two years or five or 10 years.

OLIVIER CREPIN-LEBLOND: It's Oliver speaking. So the question, I guess, if it's going to take a few more years, two or five years or so, and I understand that this has been on for a while, why has it taken so much time to deploy IPV6?

JASON SCHILLER: Well, the issue is there is real cost involved in deploying IPV6. Oftentimes old hardware has to be replaced, there is training issues. You have to get all of your staff trained up on IPV6. And what you get from deploying IPV6 is the ability to continue doing business as usual when IPV4 addresses run out.

So there is added cost in terms of hardware upgrade, there is added complexity in terms of support, there is added security concerns that you have to address, there is – you're doubling the amount of testing that needs to be done to certify equipment. So there are significant costs involved in doing V6, and you don't get any new revenue, you don't get any new products and new services.

You just get to keep doing business as usual. So a lot of providers have tried to do the minimum in terms of testing and certification to make sure that they're equipment supports IPV6 and turning it on anywhere that they can for free, but they've been trying to defer all the

equipment upgrades until the last possible moment so that they will be able to do IPV6 everywhere just as IPV4 addresses are running out.

OLIVIER CREPIN-LEBLOND: Okay. Thank you Jason. Yaovi Atohoun, you have the floor.

YAOVI ATOHOUN: Thank you. Yaovi speaking. Just a comment. I have a question about the term transition because I [?] …I see [?] meeting, when people are talking about the issue of digital broadcasting, you know? There have been deadline, 215, to move from analog to digital TV broadcasting.

So when people also in the meeting were talking about transitioning, I was thinking that we can lead to conclusion because for instance, as you said, before I think IPV4 and IPV6 will continue working. The point is that IPV6 is providing more space, so can you tell the difference from digital from [?] …they are going to use a frequency called [?].

But for IPV4, I think that it will be working, so I'm glad that you said that. Maybe transition will mean preparing [?]. So IPV4 AND IPV6 can be able to talk to each other and not [?] to stop IPV4, [?] user. And the second question I want to [?] …care about this because for user, the user doesn't mind, he just wants the Internet, so what is the place for the end user?

Or why is the introduction not about this technical asset? Thank you.

OLIVIER CREPIN-LEBLOND: Thank you Yaovi. I'm not sure who wishes to answer those questions. If we can get a volunteer.

LEO VEGODA: Could you repeat the question? I didn't hear it clearly.

JASON SCHILLER: Leo, I think the short of the question, the comment was we're not really transitioning away from IPV4, we really need to get to dual stack, and I think that that's spot on. But the question was, why do users care about IPV4 or IPV6, they're just trying to get on the Internet.

LEO VEGODA: I completely agree with that. Users shouldn't need to know anything about how the Internet works. Users should just need to turn it on and there it is and it just works out of the box and they don't need to do anything special.

LOUIE LEE: If I may say, this is Louie.

JASON SCHILLER: Go ahead Louie.

LOUIE LEE: Oh, yeah, the complexity itself is more than what users are accustomed to. But what they will find is that at some point they'll find that there are certain websites they can't reach anymore, if there only on V4 and

those new sites are only on V6, unless there is some kind of [?] or some sort of tunneling technology in place that has to be consciously installed.

It's not just going to be automatic for the most part. There will be certain sites that they cannot reach and that would be the kind of pressures that service providers will start hearing, are they hearing email directly to the website? Or they will start complaining to their own ISPs.

OLIVIER CREPIN-LEBLOND: So Louie, it's Olivier speaking. So effectively you're only going to start seeing users pushing for IPV6 when things start getting broken. Is that correct?

LOUIE LEE: That's how I see end users asking for V6. Unless you have the folks like you and me who are, yeah, yeah. But yeah, that's the chicken and egg problem.

OLIVIER CREPIN-LEBLOND: And I guess, yeah, that's where we're going. The chicken and egg problem is one of the major problems that we have with this because it's the technology that end users are not really aware of. What they want is to have their Facebooks and their Gmail and their services running, they don't really care how they're run.

And I guess that's one of the things the industry itself, the ISPs need to plan for their future in order to keep their customers happy, and that's probably something which goes better in some ISPs than in others. I'm mindful of the time, and I think we're on the subject of addressing, and I wonder if we could skip part three for the time being, maybe moving it to further down in the call, and just continuing speaking about addresses and in particular the IPV4 address recovery policy.

One of the ways to push back the day when an end user will try and access a website and will not be able to access that website and we'll get a blank screen or get an error instead. Is to try and recycle an IPV4 addresses and for this we have Ricardo Patara and Jason Schiller who will be able to provide us with more info on that.

And specifically, what's happening in the region Internet registries.

JASON SCHILLER: Can we have this line? Okay, so I'm going to do quick introductions and then we'll do the RAR policy development. My name is Jason Schiller, I'm an SOAC member from the [Aaron] region. I'm an arch engineer for Google and I also do the numbers administration for Google. Ricardo, do you want to introduce yourself?

RICARDO PATARA: Yes, thank you Jason. My name is Ricardo. I'm on the [?] integration. I currently work at [?] it's [?] for the domain name dot DR, but we also, we are also [?] so we allocate IP addresses and [?] through ISPs and other organizations [?].

JASON SCHILLER: Okay, we've also got Louie Lee who is the chair of the AOAC on the call. Louie, can you introduce yourself? Louie?

LOUIE LEE: Sure, as soon as I unmute myself. Yes. Can you hear me okay? Okay. Chair of the SO address counsel, I work at… Okay. I am currently working at Equinix. Prior to Equinix I was at Netcom, one of the first commercial Internet service providers.

My responsibility at Equinix does include working with all of our customers. I try to build products that are relevant to their needs. Our customers include all the very large carriers, all the very large content providers like Google, Facebook, and Yahoo, so on. We also have financial customers that would like to get together in the same data centers. We have data centers all around the world.

JASON SCHILLER: Okay, so we're going to go ahead and jump right in and first we'll talk about the status of IPV4 address space. The number resource organization, the NRO, quarterly puts out statistics showing how much address space is left. You can go to the website and see every quarter where we're at.

This is the last time they did it, as of June, and it basically shows where all of the slash eights are allocated. We've got these 91 central registry slash eights, these are the pre-RAR allocations, sometimes called legacy addresses. We've got another chunk of addresses that are unusable, 35

of them, those are allocated to things like [?] or class E experimental usage.

We've also got some private use stuff in there as well. And then the remaining 130 are broken out among the various regional Internet registries. APNIC has the most 45, followed by [Aaron] and Ripe which have 36 and 35, and then AfriNIC and LACNIC have five and nine.

On this next slide, this shows the allocations or assignments that the RIRs make to their customers. Their customers are either ISPs that turn around and give them to other downstream ISPs or customers, or they also make direct assignments to customers as well. And it's broken up by the five regions.

This next slide is looking at what is the available inventory of each of the five RIRs. You'll see that AfriNIC has the most with almost four slash eights, followed by LACNIC which has 2.32 and [Aaron] with 2.32 and then APNIC and Ripe have less than one. Both of those regions have created a soft landing policy where they've decided to designate the last slash eight for special usage, and that's only for new organizations that don't have any addresses.

So those two regions are essentially out at this point already even though you do see some addresses in their inventory. This next slide is generated by Jeff Houston who works for APNIC. He's been doing a lot of research on the draw down on the IP addresses, and this is showing when Ripe and APNIC ran out, as well as the projections for [Aaron], LACNIC, and AfriNIC.

And he keeps this up on his website [Pottery] you can check it daily, he'll tell you when we're expected to run out. So we want to talk about IPV4 address recovery. So the easiest way to figure out if there is a policy for something in the different regions is to go to the NRO, RAR comparative overview, the URL is there at the top of the slide.

And when you do that, you can quickly compare what are the policies for recovering unused resources, for example. In APNIC and LACNIC, a resource is valid as long as the original criteria continues to be met, and they actively recover unused resources. Resources of organizations that cease to exist, companies that go bankrupt for example, are returned to the free pool for that particular region.

Ripe is also very similar. They are valid as long as the original criteria continues to be met. And if an organization ceases to exist they return to the free pool. [Aaron] is a little bit more different. They look at organizations that are out of compliance with the current [Aaron] policy, and those organizations are requested to return address space, and they can be compelled to return the address space.

They also revoke resources for lack of payment and for fraud. We can actually dig down into the [Aaron] number research policy manual, and look at the various places where these policies come from. And it's fairly complicated to do because they're not all consolidated in one place.

They have a section 12 for resource review, and that basically says that [Aaron] can go back to an organization and review how they're using the addresses to make sure that they're efficiently utilized. And you'll

find here in bold a number of paragraphs down, organizations found by [Aaron] to be materially out of compliance with current [Aaron] policy shall be requested or required to return resources as needed to bring them into, or reasonable close to, compliance.

And if an organization does not voluntarily return the resources, as requested, [Aaron] may revoke any resources issued by [Aaron] as required to bring the organization into overall compliance. There is a lot of policy surrounding the review [Aaron] can do a review at any time it deems, as long as it hasn't done one in the last two years.

They can do one when you come back for additional resources, or they can do if they have a suspicion of fraud. [Aaron] also has an annual renewal. They require everyone holding resources to pay [Aaron] yearly. In some cases, it's quite a nominal fee for end users, for example.

But the importance of having the feel there every year is so that [Aaron] can maintain the relationship with the resource holder. And if that relationship goes stale, if they don't pay their bill then [Aaron] can assume that the addresses are no longer in use and revoke them. The other point that's worth making is the IPV4 transfers, and there is two important points here.

The first is, transfers is another way to try and liberate unused addresses and get them to people who need them. So that way, rather than revoking your returning addresses, they can be transferred. But the other thing that's worth pointing out is, when a transfer happens

the new organization that is receiving the addresses has to justify the usage of those.

And you'll see here at the bottom in bold, if a company goes out of business, regardless of any reason, the point of contact listed for the number resource does not have the authority to sell, transfer, aside or give the number resource to any other person or organization. The point of contact must notify [Aaron] if their business fails so that the assigned number resources can be returned to the available pool of numbered resources if a transfer is not requested and justified.

And additionally, in the event that the number resources of the combined organizations are no longer justified under [Aaron] policy, at the time, [Aaron] becomes aware of the transaction, through a transfer request or otherwise, [Aaron] will work with the resource holders to return, aggregate, transfer, or reclaim resources as needed to restore compliance by the process outlined in [Aaron] policy.

So addresses are actually moving around. [Aaron] has recovered quite a bit of address space, and usually when they get a large block, they return them to IANA for redistribution. We now have a global policy that allows IANA to split what they're holding five ways and allocate them to the five RIRs.

Also have some stats on LACNIC, they've done some recovery as well, nearly a slash 14 equivalent. Additionally, I was talking about transfer, there are transfers that are going on, so organizations that have addresses that are underutilized, they can sell those addresses to another organization and transfer the recourse. This is just showing the

number of transfers that have occurred in the APNIC region, as well as the number of transfers that have occurred in the [Aaron] region.

The interesting thing about APNIC and [Aaron] is, they have a bilateral arrangement in which you can transfer resources from one region to another. And you'll note that about half of the [Aaron] transfers were out of region to APNIC. There are also transfers going on in the ripe region as well.

These three regions are the most interesting because, as I said before, both ripe and APNIC are essentially out because they are down to their last slash eight, which is reserved only for new organizations. And [Aaron] is interesting because it is nearly out and it has a lot of legacy address space and it has the ability to transfer addresses from the [Aaron] region to the APNIC region.

So those are the three regions that typically report on transfers. So shifting gears and looking at the V6 address space, you'll notice that these graphs look very different from IPV4. We're using a very small amount of V6 address space in the global routing system, and of that, each RAR has one slash 12, and then there is an additional slash 12 which has some various registry space out of it.

That was where the initial allocations of the RIRs got that were much smaller than 12. And you can also see the allocations that the RIRs are doing to their downstream customers and downstream ISPs. And this…

LOUIE LEE: If I may, if I may just note, a couple of slides back, where it shows, yes, very small slices. So this gives us a way to make policies without worrying too much that we're going to run into the same issues that people say we did with IPV4. We can make policies, there is some flexibility, and we can learn from them so that the next tiny slice can be better used. So go ahead.

JASON SCHILLER: So this graph here is showing the percentage of members that have both IPV4 and IPV6 broken down by region. The interesting thing here to note is that the ripe region has the most widespread IPV6 adoption contrary to the belief that everyone thinks that Asia is the continent that is running out and needs to embrace IPV6 strongly.

So IPV6 allocation and assignment policy. And please do stop me if I'm saying strange words. We do have some fairly technical jargon, for example, allocation is a term that is typically used when NRAR makes a number resource available to an ISP that is going to then in turn make it available to a downstream organization.

Assignment is when a number resource is given to an end site that will use the addresses themselves, and they don't have the ability to make further downstream reassignments or reallocations. So again, the best place to go and look for what is the policy for V6 initial allocations is the [?] comparative RAR report at that same website.

And if you look at it, the one thing that you'll find is that all of the RIRs provide a slash 32 by default to ISPs. [Aaron] is a little bit different in that they're default is a slash 32, but upon request you can ask for a

slash 36 instead. The reason for this is the [Aaron] building is based on different sized buckets, and extra small pays a certain amount a year, a small pays a certain amount a year.

And a slash 32 because the address space is so large, and because it covers a large majority of the ISPs that come to [Aaron] for address space, that's all been to the small bucket. So currently an extra small ISP in IPV4, as soon as they do IPV6, they're upgraded to a small and have to pay more. So in order to avoid that, they've offered a smaller slash 36 which is available upon request.

And site assignments typically start at slash 48. So digging in a little bit deeper, we can look at the AfriNIC policy. In terms of AfriNIC ISP allocations, the organization must be an ISP, not an end site. They have to have plans to deploy IPV6 in Africa. They have to show plans to have downstream customers within a year, and announce the IPV6 space to the Internet within a year.

By default, they get a slash 32, but they can get larger than a slash 32 if they can justify it with their existing customer base number structure. So you can basically count up all of your IPV4 customers and IPV4 infrastructure and assume that it will be doing IPV6 one day and get the appropriate addressing for all of that network and all of those customers.

You come back and get additional space when your HD ration slash 48 is greater than zero point nine four, and what that means is basically they take the log of the slash 48s that are in use, divided by the log of the total number of slash 48s that the organization is holding.

In terms of end sites, you must be an end site and not, what AfriNIC calls a LNR, which is an acronym for ISP basically. You must be able to qualify for AfriNIC provider independent IPV4 space.

You actually don't have to hold it, but you have to be able to qualify for it under the current policies. And you must show a need for IPV6 provider independent space. You must put it into service within one year, and you'll get a slash 48 by default but you can get a larger block if it's justified.

There's really no good guidelines of what is acceptable justification though. For APNIC, also a default of slash 32, if you already have IPV4 ISP address space from APNIC, you qualify for a slash 32 of IPV6 space. If you don't already have IPV4 ISP address space, then you have to show all of these three things. You have to be an ISP and not an end site. You have to have a plan for downstream customers, and have a plan for at least 200 downstream customers in two years, or provide IPV6 transit to customers within two years.

You can larger than a slash 32 based on your current customer base and infrastructure, assuming that it will do, that [?] V6, how much would you need? Or, if your current IPV4 current customer base and infrastructure plans to support IPV6 within two years. You can also get additional space with a HT ratio of 0.94, but in this case they look at slash 56s instead of 48s.

And when you come back to get additional space, they allow for a two year supply. For end sites, it's a slash 48. You have to have… You need to have a provider or independent IPV4 address from APNIC or be multi-

honed in IPV6, or demonstrate that an ISP block from an upstream ISP will not work.

Ripe is also very similar. For ISP allocation, you must be an ISP, you must have plans for V6 customers in two years, and you get a slash 32 by default. There are a little bit more liberal than the other three in that you can request up to a 29 with no additional justification. Beyond the 29, it's based on your customer base and your infrastructure.

And like the other two that we have seen, it's a HD ration of point nine four of slash 56s. For end sites, it's a slash 48 and it can be larger if it's justified by the number of subnets or the number of discrete routing sites. By they specifically state that you cannot reassign or reallocate these addresses. These are for use within your organization, you cannot reassign them to downstream customers.

LACNIC is very traditional. And the V6 initial allocation or assignment policy that we saw throughout the globe, but the others have drifted away from it over time. You must be an ISP, you must have a LACNIC IPV4 ISP address space. Or, if you have a plan for downstream customers using slash 48 and you announce the space within 12 months, and you have downstream customers within the LACNIC region within two years, then you qualify.

By default, this has 32. You can get more based on a four year time horizon of what your projection is to need for V6 addresses the next four years. And that's based on your current customer and infrastructure, as well as your projected growth of your current customers and infrastructures.

You can also assign aggregates to parts of your network, but those parts of your network have to be at least 30% utilized. You also get additional address space with a hd ratio of point nine four on slash 48s. You can get a two year supply, and if you got your addresses and you come back within six months and you decide that you asked for the wrong amount, you can return what you've got and reapply as an initial allocation again.

The [Aaron] policy is one of the most complicated, but also one of the most liberal. [Aaron] has decided that they only want to give out ISP block on nibble boundaries, which is basically counting by forms so you can do a 36, 32, a 28, a 24, etcetera. By default, it's a 32 by as I said previously, you can request a 36 to reduce your billing.

And the way they decide how much you can get is they basically look at the amount of address space you need to basically look at your largest serving site. So if you imagine that you're aggregating customers on edge routers, you can take your edge router with the largest number of customers, you could say that you're giving all of your customers a slash 48 each, count up how many 48s you need for that edge router, round that up to the nearest site, and then you can use that block size to give to every edge router in your network.

And then when you add up all of those blocks from all of your edge routers, you round up that up to the nearest nibble. And if you're very close to the nearest nibble, you round it up to the nibble above that. So basically the idea here is to give a large oversizing to allow people to easily manage and grow their networks.

You must have downstream customers, and either be an [Aaran] IPV4 ISP, be an IPV6 multi-honed or have a plan for 50 IPV6 customers within five years. And you can get additional address space when any of the three occur. If you're 75% total utilization, if you're at 90% utilization at any single serving site, or if 90% of your serving site blocks have been allocated and are still justified.

Generally what they'll do is that they will bump you up to the next nibble, until you reach a slash 12, and at that point they'll just give you a slash 12 at a time. For end site, just like everyone else, it's a slash 48. You must either be an IPV4 end site holder provider and dependent based on [Aaron], or be multi-honed, or use 2,000 IPV6 addresses within one year, or have 200 slash 64 subnets in one year.

And you can get larger if you can justify a better number of subnets or the number of discrete routing locations. So I think you all wanted to know how can you get involved in making RAR policy. So, what is a RAR? A RAR is a regional Internet registry. They're established in IP ICP two, which you can find on the ICANN website.

And they're basically defined as operating a continental sized international geography, have broad support of their community, have a bottom up self-governance, they're transparent, they're neutral and impartial, and they have technical expertise and adhere to the global policies of conservation, aggregation, and registration.

So they've basically gone and divided up the globe into five continental sized regions. [Aaron] has North America and the English speaking Caribbean. Ripe has Europe. APNIC has Asia and the Pacific. AfriNIC

has Africa, and LACNIC has Latin America and the non-English Caribbean.

Each RAR establishes its own regional policy. Each RAR has its own policy development process. And this policy is developed from the community in a bottom up manner, it's consensus based. They all have open policy mailing list and open public policy meetings, and it's all transparent. They're policy development process is documented and their number policy is documented.

So there are some differences in flavor, but generally each RAR has its own published policy manual. They each have their own policy development process. And then the publish draft policies that are under discussion, they're discussed on the mailing list, and in meetings which happen twice a year, they're open and they also allow remote participation.

So generally the way it works is any individual in the community, or a group of individuals, can submit a proposal. The community discusses that proposal on the mailing list. Staff provides their assessment in terms of how changes – how it changes to how they operate. They also provide a legal assessment if there is a legal impact. The community will discuss these proposals in an open public policy meeting.

They will judge their consensus, and if there is they will move to last call. The board will view a fiduciary risk assessment. After that it gets adopted or ratified and implemented. Global policy is a policy that requires the IANA to do something towards the RIRs. In this case, you have to pass five regional policies that all have the same language.

They would follow each of the normal regional policy development processes in all of the five regions. And then it gets passed to the NRO for review and up to the ICANN Board. So once the same policy is passed in all five regions, the NRO NC validates that the proposal went through the correct development process in each region.

That it was discussed for the right amount of time. That it was posted on the mailing list for the right amount of time. That there had been adequate consideration of viewpoints. They also have to judge if the same policy has passed in all five regions because sometimes the same text doesn't make it through all regions, but it's essentially the same policy.

Sometimes the NRO NC will rewrite the policy and it's a substantive change, in which case the NRO NC will send it back to the communities because it's not the policy that was passed. If the NRO judges that it followed the right process, and if the same policy passed in all five regions, then it will recommend to the ICANN Board to ratify it.

And the ICANN Board can ratify it. They can sit on it which it then becomes ratified by default. They can ask for clarification or they can send it back to specifically reconsider certain concerns. So I have a bunch of slides for the five RIRs that basically gives you the pointers to where all of these documents are.

So if you were particularly interested in say the [Aaron] region, you could go find a number of resource policy manual, read it, see what the current policy is, pull down their draft policies that are currently under discussion, read them. You can join the mailing list. You could certainly

comment on those draft policies and say whether you're in support or against them, or how the language could be changed to be better.

And you can certainly attend any of the RAR meetings. They're open. So I've listed where the meetings are, and when they are, and websites where you can find those meetings. And with that, I guess we'll hand back the microphone and open it up for questions.

OLIVIER CREPIN-LEBLOND: Thank you very much. It's Olivier Crepin-Leblond speaking. That was, I would say, quite involved and some of it might have been a little technical as we've seen in the chat. Thank you very much for going through this. It's quite fascinating, it's a role that our community doesn't know very much about.

I have a few questions, but I'll first open the floor for questions from others, questions and comments from others here, and then we'll take a few more minutes for the questions that I do have. Anyone wish to start? I don't see anyone putting their hand up, so maybe I can throw in a few questions.

The first thing is, we've seen all of this IANA and the RIRs and all of the address policy making and so on. And this is really something that would advocate, in a way, how important are the RIRs and how important is the mission of the RIR? Couldn't we just get rid of them and have numbers directly allocated from IANA and the NSO within ICANN doing the policy?

LEO VEGODA: Olivier, may I answer?

OLIVIER CREPIN-LEBLOND: Yes please, if you introduce yourself.

LEO VEGODA: Sure. I'm Leo Vegoda from ICANN, I spoke earlier. I work in the IANA department. Just on the issue of do we need the RIRs or could it all be done by IANA. The RIRs have significant staffs and significant memberships. I believe, for instance, the ripe NCC has in excess of 9,000 members.

And we have just a couple of people who provide service to the RIRs, and we're not, at the moment, scaled up to provide service to literally tens of thousands of customers around the world. I think that the hierarchal distribution system is an advantage and it means that people get good customer good service in their own time zone, and there is a language support where that's required, and all sorts of things that are very difficult to build into single organization.

So I think you could argue that RIR, for that specific reason, are absolutely essential. I think that things would be significantly worse without them.

JASON SCHILLER: Jason Schiller. The other thing too is that having the five different regions allows for regional variations, which there is quite a bit of. And in addition to the needs of the region being different, there is also

cultural differences. If you attend an [Aaron] meeting, you'll find that the Americans and the Canadians tend to be very vocal, they like to argue, and they're very anal in terms of making sure the policy language is just right.

And they kind of view it as, if I'm in a court of law and I need to get the addresses that I think I deserve, and I'm held to the letter of this policy, would that be a good thing for my company and for the Internet as a whole? When you go to another region, say like the Asia Pacific region, it's much less confrontational. A lot of the discussions have been in one on one, and a lot of times you'll hear them say things like, "The actual text of this policy doesn't really say what we all know what we want it to be. But we're all comfortable that the principle is sound and we'll go ahead and we'll pass it and we'll all know what it means."

So I think that there is also a lot of value in terms of allowing for the regional differences. And I think because of the cultural sensitivities and because of the process is bottom up and it involves with an open community, I think it would be much harder for that to be successful at an international body, whether that be ICANN, or IANA, or somewhere else.

OLIVIER CREPIN-LEBLOND: Very, very interesting. It's Olivier speaking. Very interesting indeed. As you know, the At Large community is divided in its five regions, and we find this to be particularly important for the balance of interest that we get from the different parts of the world, and I think that you'll hear

many echo exactly what you just explained there with regards to how culture plays a very big part in the way things are done.

There is a question from Evan Lebovitch, he's not able to speak at the moment, but he texted this over to me. And so I'll read it over to the record and then either of you can decide on responding to it. So here's the question, given that this is an At Large audience, how would you describe efforts to advance the IPV4 declaration, as an issue that matters to end users?

Most of them don't see a problem, and map techniques, therefore the [?] seems to get more advanced, so most don't even know what's happening underneath. How are end users enacted, and how can they be motivated to take this issue as their own?

JASON SCHILLER: I don't think end users should actually know what their IP address is, or whether they're using IPV4 or IPV6, I don't think they should care. They should be able to get to the Internet, and it should work. But their upstream providers need to be very concerned.

Their upstream providers are the ones that need to continued IP addresses for growth of their customer base, or they need to deploy a carrier great net type solution. And a lot of people have concerns about how the [?] net would work because it basically has to be an application of a gateway, it has to understand all the applications that it's letting through. It has to be SIP aware, it has to be FTP aware and so on and so forth.

So there are issues with that. But beyond that, people are also concerned about the scalability of it and the performance of it. So it depends on where you deploy your carrier great net inside your network topology. If it's very close to where you terminate your customers, then it's going to be a very expensive solution, but there is not a lot of latency and there is not a lot of bottle necking if you funnel all the traffic through the CGN device.

If you deploy it more centrally, then that brings costs down quite a bit, but it increases the latency because you have to funnel all of your traffic and haul it back to whatever that CGN point is. And it can also potentially be a bottleneck. There is also privacy concerns. When you're running a CGN, you basically have to keep your TCP level transaction information.

So if law enforcement shows up and says, "We need to know who posted this child pornography," they might have the IP address and the port of who hit the webserver, but that IP address might be shared by 1,000 people. So the ISP will have to go back through their logs and figure out exactly what inside address is using what outbound port in order to map it to a particular unique individual.

And that basically means they're keeping TCP level transactions. They're knowing every website you visited and they're recording that and they're holding that for five or seven years. So there are certainly privacy concerns with regard to that. The other issue is Konami which is a game manufacturer in Japan. They're going around and talking about the implications of having CGN net in terms of performance for cutting

edge gamers who want the best throughput and the lowest latency to their servers for their gaming performance.

So some of the power users may be concerned about going through a CGN, but I think all and all, it's either going to work well and people aren't going to notice and care, or some things are going to break. And in that case, people are going to choose, do I go with my current provider that's forcing me to go through CGN that impacts what I can do? Or do I switch to their competitor that still has IPV4 addresses and won't make me go through a CGN?

Or their competitor that's deployed IPV6 because it just so happens, I'm a Xbox one customer and it does IPV6 and IPV6 works perfectly well for everything that I want to do.

OLIVIER CREPIN-LEBLOND: Thank you Jason. Louie?

LOUIE LEE: I would add that for the users that do hear about V6, they can help the effort by asking their ISPs if IPV6 is supported on their connection. With the idea that not too long, if you're IPV4 only, you're not on the whole Internet anymore, you're only part of the Internet. So ask to be on the whole Internet is what I would say.

For those users that are, that know enough about V4, V6, and if they want to be involved in the effort, but otherwise we are pushing for enterprises to have their, have them ask their ISPs. They spend more money on a per user basis than your home customer. So the

enterprises, for their business continuity, for them to be able to reach all of the Internet will need to be also on IPV6. For that to happen, they have to ask their ISPs. Thank you.

OLIVIER CREPIN-LEBLOND: Thank you very much Louie. It's Olivier here. I mean obviously that would be valid if there was some IPV6 only content, but is there at the moment IPV6 content? I thought most of it is just dual stack.

LEO VEGODA: Yes there is. I just posted an URL for an IPV6 only puzzle game into the chat room.

OLIVIER CREPIN-LEBLOND: Very quick on the trigger, Leo, thank you very much. Olivier speaking. The threat of loss of privacy, and the slow down with net, and all that, is this just scare mongering in a way?

JASON SCHILLER: I don't know if that's well known to be honest. I think map technologies used in the customer premise equipment at the home has been in place for a long, long time. But that's not really what we're talking about here. We're talking about provider based map and potentially two levels of map where they might be doing that inside the provider's network and then map again at the household.

This type of equipment is still very immature, and is not well tested. I don't personally have any experience to know whether the performance is an issue and how good the ALGs are, but I think many of us have experience with the pains of living with firewalls, especially when new protocols come out and those pains certainly are going to be identical in a CGN platform.

OLIVIER CREPIN-LEBLOND: Wow. It's Olivier speaking again. So effectively what you're saying is the CGN technology is less proven than IPV6 performance.

JASON SCHILLER: So I think IPV6 works well in a native deployment, and I think IPV4 works well in a native deployment. And I think most residential users haven't had native IPV4 for a long time. The typical solution is to give a single public IP address to the customer premise equipment and then do map behind that.

IPV6 allows you to reestablish the end to end connectivity between things on the Internet. And they give you enough addresses to continue to grow, and it gives you enough addresses to put things on the Internet as opposed to just people. I mean, people talk about coasters and cars and all kinds of devices on the Internet.

So I think there is real gain within IPV6. The challenge here is, what do we do in the short term where some networks haven't deployed dual stack, IPV4 and IPV6? And we start to see networks running out of IPV4

addresses? Are we in a place where an ISP can just turn up IPV6 only customers?

Imagine if at your house, you were to buy new Internet service, and they just give you an IPV6 address, and they say, "You can get to all of the Internet, but just the things that are on IPV6." I mean, how would your browsing experience be? So the question is, what do we do? Well, yes, we should try to encourage all of the sites on the Internet to get to dual stack and to get it before people start building V6 only networks, but we haven't been very successful with that.

So the alternative is to go through some carrier grade net type solution. The big question is, do people think that carrier grade net is going to be a long term solution, meaning IPV6 is not important? Or will people move to native IPV6 in a real way and to only use carrier grade net for a smaller and smaller number of edge cases where the service hasn't deployed V6?

OLIVIER CREPIN-LEBLOND: Thank you. Yaovi?

YAOVI ATOHOUN: Thank you. Yaovi speaking. My question is to know why somebody can turn on an activity network only now and he knows some people may not be able to reach him. That is my question. Thank you.

JASON SCHILLER: The challenge is, if you're in Asia or you're in Europe, and you are an ISP that already has IPV4 addresses, you're going to continue to use those IPV4 addresses as you add new customers to your network. But at some point, your store of IPV4 is only so large and you run out. The question is, what will that provider do? So this is where you have to make a difficult decision of, can I just turn up only IPV6 customers?

It would be great if I could, because then I would just transition my growth to IPV6 and life is good for everyone. Or, do I have to get CGN equipment, deploy it, support it, figure out what are the latencies and [?] of it, the impact and try to position my service to be as good or priced better than my competitors that haven't been driven to CGN yet.

So the other interesting point here is that everyone is going to run out at different times, and that's going to create a competitive advantage. And then beyond that, the question is, well does it make money to spend money on a CGN solution? How many customers will I lose if I just do V6 only?

Sure, Facebook does V6 and Google does V6, and maybe that's all my customers care about. But beyond that, is the questions too of IPV4 transfers. Do I go out and buy IP addresses on the market? And how much is an IP address worth to me to be able to grow my business? And how much revenue am I extracting from a customer and how much of that do I have to turn around and put into purchasing IPV4 addresses?

And then theoretically at some point, the Internet goes widespread dual stack, IPV6 is supported everywhere, and at that point I can just do my future growth in IPV6 only. And the question is, well when will that

happen? And how does that change my investments? If that happens in the next two years, what do I do with all of the CGN net equipment that I purchased that is now no longer useful for anything else?

So it's actually a very challenging proposition for providers to figure out, what are they going to do as they face IPV4 scarcity.

OLIVIER CREPIN-LEBLOND: Thank you Jason. Olivier speaking. It just sounds like this is all financial [?]. The vector for change will apparently be at all stages be a financial one wouldn't it?

JASON SCHILLER: Yeah. Unfortunately, I mean, 10 years ago, ITF decided to make IPV4 and IPV6 not compatible on the wire, and they said the transition plan is for all of the Internet to go dual stack before the first network is V6 only. And had we done that, everything would be fine.

Unfortunately, people are trying to defer the cost of doing V6, and they're trying to do that because the V6 doesn't generate new revenue, or new services or new capabilities. All it does is allow them to have business continuity. And everyone agrees that they need business continuity and that they need to deploy V6.

The challenge is they want to deploy V6 just before they run out of V4, and the problem that this creates is not everyone is going to run out at the same time. So rather than getting the Internet as a whole to be dual stacked before the first network runs out of V4 addresses, each

individual network is trying to figure out the day which they will run out of V4, and for them to do V6 in a real way.

And the problem is, if you're early on that curve, if you're one of the first organizations to run out of V4, then you're going to be stuck for a while until the rest of them run out.

OLIVIER CREPIN-LEBLOND: That's very good description, thank you very much for that. I cannot help, having visions of the Titanic [?]… [laughs]. It certainly seems that we are drifting slowly, and hopefully the transition will be smooth. I'm mindful of the time, it is already 17 minutes past the hour. Any other questions from anyone? Otherwise, I have one last question I would take us out of the IPV6, the pure IPV6 discussion.

I see no one putting their hand up. So the question I had was with regards to the RIR policy development. How important is it for, I guess, end users, if anyone in our community, what would be the reasons why they would be interested in participating in RIR policy development? And I guess the extension of this, would they be able to do so as well?

JASON SCHILLER: So the question is much easier to answer, absolutely. Individuals can participate in the policy development process. Even when a company or an organization sends someone to an area meeting, it's very clear that everyone participates as individuals, although their opinions are influenced by their perspective and their employment.

Certainly if you work for a transit provider, you see a lot of issues with policies that impact transit providers. But in terms of participation, we all participate as individuals, and it is an open, bottom up, transparent process. Anyone can show up at a RIR meeting and participate. Anyone can join the mailing list and chime in on draft policy proposals.

And if you think a policy needs to be changed, you can certainly draft a proposal or even just submit a suggestion and get others to help you write a proposal. So certainly, getting involved, attending meetings, speaking up on the mailing list, all of the documentation about what current policy is and current proposals that are being discussed is all publically available on the RIR website. So it's very easy to get involved.

In terms of why you should get involved, these policies really inform how IP addresses and AS numbers are given out to providers and in turn to their customers, as well as to direct and [in site] organizations. So if you're a customer of a typical transit provider of say, a Comcast or a level three, it's important that you know that your provider is working in a policy framework that they can get the addresses they need and that there are using them responsibly, which means they can continue to add new customers.

So that when you move to your new home, you can go and buy Internet service and know you can get an address and you'll be able to use it. There are also some policies that directly affect end users. There are, recently there are policies where if you are an end user with a static IP assignment, you are no longer required to publically publish which IP address belongs to which end user.

For corporations, it's still required, and that's a change that happened a few years ago. It was mostly driven by the Canadian law enforcement. There's also recently been a discussion about whether RIRs should give IP addresses directly to individuals. A long time ago, if you were an individual and you talked to John [?] and you said, "Hi. I'd like a solid 24 for my house," you could get one.

Today, in the [Aaron] region, you have to be an organization in order to get IP addresses. So you would have to create a sole proprietorship and use that to get the addresses. Part of the complexity here is the reason why [Aaron] has structured it that way is because they're trying to reduce fraud. So they want to make sure that IP addresses are given to legal entities that they can verify, and it's much more difficult for them to verify that when an individual comes to them and claims to be an individual that they actually are who they say they are.

So that's probably a topic that we're going to be discussing in the next six months or so, in the [Aaron] region.

OLIVIER CREPIN-LEBLOND: Olivier speaking. That's something that actually I guess does affect end users. One last question, I did notice here that there is a question here or comment from Dev Anand Teelucksingh on the chat, whether transcripts and recordings of the meetings will be publically available for those who were not able to attend? Are these all archived?

BARBARA ROSEMAN: This is Barbara. For most of the RIR meetings, they are very good remote participation capabilities. The archiving of them, as usual with these things, may take a little time to make them available after the fact, but all of the RIRs are in fact quite good about enabling remote participation during the meetings.

JASON SCHILLER: Yeah, definitely with the RIR meeting, you can tune into the video stream. They have a Jabber client that you can ask questions. After a policy is discussed, in order to determine if there is consensus, they have a show of hands, you can virtually raise your hand in the Jabber room and be counted.

These videos are also available after the fact to be watched. So you can back and say, in a week or two and, we had discussed this policy and what did we decide? And you can go back and replay it and watch it. And [Aaron] also does publish a fairly good transcript of the meeting, but that does takes a few weeks to come out.

But if you go to YouTube, you could see some of the [Aaron] meeting sessions. They're still putting them up, it's just two weeks back. But yeah, the meetings are all there and you can go back in time and look at older meetings as well.

OLIVIER CREPIN-LEBLOND: It's a [?].

RICARDO PATARA: Yes. [?] if Jason [?] …have a [?] published record of the meetings, and when [?] …in our case, our meetings are translated from Spanish to English, from Spanish to Portuguese, and also the other way around. [?] participation from different regions. And how the recording of the meetings are available on the website.

And also one can check the history of the policy development, as other RIR will have a page where one can find proposals presented in the past, if the proposal reached consensus or not, all this information is available on the website.

OLIVIER CREPIN-LEBLOND: Thank you Ricardo. I realize we are running out of time, it's nearly reaching half the hour. Just one last question, and in fact maybe even a suggestion, are there any beginner's guides to all of what you have just told us today? So both beginner's guide to IPV6 and beginner's guide to RIRs? And to RIR policy development?

LEO VEGODA: Olivier, Leo Vegoda speaking. I just want to let you know that we're finishing up the editing on a new beginner's guide that ICANN should be publishing. We're hoping to have it available in Buenos Aries in at least English and Spanish. It will be eventually translated into all the UN languages.

And I know that the RIRs will sort of publish all sorts of material as well.

JASON SCHILLER: Yeah, certainly all the RIRs publish very good descriptions of the policy development process, the PDP. And that basically walks you through what is the process from an idea to that becoming a RIR policy. So that documentation is very good.

In terms of getting involved, I would say at least for the [Aaron] region, we have an elected body called the [Aaron] advisory council, they do from time to time, do mentor people as part of mentorship program that's actually formalized.

But they're also very good about holding new individuals hand and working with them if they have consideration or a suggestion for how policy should be changed. So it certainly is possible to come in as a new individual completely new to the process find an [Aaron] advisory council member and borrow their ear and say, "I'm really concerned about this potential policy and it really should do that instead."

And they will certainly work with you to write a policy proposal, they'll submit it to the community, and they'll shepherd you and the policy through the whole process.

OLIVIER CREPIN-LEBLOND: Fantastic. Thank you very much. I look forward to seeing the beginner's guides. We are, I think the ALAC is going to be meeting with the GNSO for half an hour in Buenos Aries. I hope that by that time you'll have a beginner's guide with you and you'll be able to bring it over to our members and certainly will shed a lot more light over what the NSO does, what the RIRs do, what the SO address council does, and how does one get involved more in addressing policy around the world.

We keep on hearing about this world, but being a part of those mailing lists, I notice an entirely different world taking place out there with very few hardcore ICANNers [sic] and the rest of ICANN being part of it. I entirely welcome this beginner's guide and I look forward to seeing it, and no doubt many of our members will be very interested in reading it.

With this, I thank all of you for having come on the call. Thank you very much Jason, Ricardo, Barbara Roseman as well who is on the call as well. And of course, Leo Vegoda thank you having arranged, it's certainly helpful for me and for my colleagues. This is recorded, so hopefully we'll be able, I'm sure… Heidi, will we be able to publicize this over to our members and also make up a page to have the PDF presentations too?

HEIDI ULLRICH: Yes, we can add all that additional information to our agenda page, Olivier. This is Heidi.

OLIVIER CREPIN-LEBLOND: That's fantastic thank you very, very much. Thanks to all of you, apologies for the late end to this call but it's a good subject and it's one that we don't often visit. So it's good to have been able to go into it with a little bit more depth than usual. With this, I thank you. Good morning, good afternoon, and good evening, and good bye. This call is now adjourned.

[Various people say good bye and thank you.]

[END OF TRANSCRIPT]