
John Crain: Okay, good morning everybody. Welcome everybody this morning. My name is John Crain. In a galaxy a long time ago before I learned how to spell domain names, I used to be able to count, and I did IP addresses, numbers, etc. And today I have the pleasure of opening up this panel.

We're going to hopefully talk a little bit about the need for the support for rather than new names which everybody's been talking about all week, actually a not so new numbering protocol that's been around a while called IPv6.

I see a lot of faces that I recognize, so I think at least 99 if not 100% of the people here actually know how to spell IPv6 and routing, or routing, depending where you're from, and such technical terms. So without further ado, I want to thank Keith Davidson for agreeing to chair this panel and moderate it actually, and pass over to our organizers.

I'd really like to thank APNIC and the NRO and all the guys from the numbering industry for helping pull this together and, Miwa, if you want to kick off, that would be great. Okay, my IPv4 microphone is working fine; the v6 one over there seems to be broken, so...

Miwa Fujii: Good morning everybody. ICANN staff, in working with this room, if you can look after this microphone, that would be very helpful. Thank you.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

So good morning everybody. My name is Miwa. I am from APNIC. I'd just like to introduce our moderator, Keith Davidson. Keith is Internet NZ Director of International Relations. Internet NZ is the dedicated monitor of the .nz domain name. Keith also is the Chairman of APTLD, the Association of ccTLD Operators for the Asia/Pacific Region. Within ICANN he chairs the same work of Interpretation Working Group which is investigating aspect of delegations and redelegations of ccTLDs and also is a member of the ccNSO Council within ISOC. Keith is one of three Co-Chairs of the ISOC Advisory Council. He strives for an open and uncapturable internet. So, Keith, thank you.

Keith Davidson:

Thank you and good morning everybody. We have a panel discussion today and I think the main view or the main desire today is have a look at what's happening with IPv6 deployment primarily in the domain name business or in the DNS.

So we have a very useful panel, and I think today we'll have each presenter do around a 20 minute presentation. After each presentation, if we can allow just one or two questions that are pertinent to the presentation, work our way through the presentations in an orderly fashion to build the total picture, and then open up for broader discussion on the floor after everyone's had a chance to present their aspects of this important topic.

So we do have also remote participation in the Adobe Connect room, so if people in the room are also able to connect and help



anyone who's remotely participating and you'll answer their questions along the way, that would be much appreciated.

And it is my pleasure to introduce our first speaker who has already introduced herself. But Miwa Fujii from APNIC is probably the chief enthusiast for IPv6, certainly in our part of the Asia/Pacific region and probably in this part as well. Miwa is the APNIC Senior IPv6 Program Specialist. Her role is to assist the Asia/Pacific internet community in managing issues related to IPv4, address exhaustion and the transition to IPv6. She is responsible for various projects that support the smooth transition to IPv6, including outreach activities to multi-stakeholders.

Miwa regularly represents APNIC's view on the transition to IPv6 at Global, IPv6 summits and forums, inter-governmental organizations such as APIC, APT, SPC and other conferences around the Asia/Pacific region. Miwa, it's very important that you set the stage for us today. Thank you.

Miwa Fujii:

Thank you for your kind introduction. The microphone in front of me is not working - just resetting myself. Just give me a sec. Alright, Keith, thanks a lot again for your kind introduction.

My name is Miwa. I am from APNIC. APNIC is delighted to organize this conference, this IPv6 session for ICANN participants by corroborating with ICANN and IANA. We successfully located great speakers today, this morning besides me, we have four more speakers. They are all outstanding speakers in IPv6 field.



I would like to give as much as time possible to those speakers so I'd like to make my presentation very quickly, just touching the fundamental issues of IPv6 transition.

So the IP – on the internet you are nothing but IP addresses. Most of you are domain field people, I suppose, in this room, so please do not confuse domain name to IPv6 internet protocol. IP has three functionalities – 1) defining packets; addressing interfaces and routing those packets through interconnected interfaces.

We have two types of IP now – IPv4 and IPv6 and every device directly connected to the internet needs to have a unique IP address and IP address is finite. This is especially the case for IPv4. So when you connect to the internet, you are assigned with an IP address and all other domain names also are assigned with an IP address. Obviously, the internet will function without domain names, but internet cannot function without IP addresses.

The internet is a big network of interconnected networks and this interconnectivity is provided through the electronic communication protocol standard called IP, Internet Protocol. We have v4 and v6. We call these internet resources and these internet resources are managed by regional internet registry. APNIC is one of them. We are all membership-based industry bodies, not for profit, neutral and independent.

And internet resources are we... our mission is internet resource allocation, registration and other services such as training, supporting infrastructure and facilitating community cooperation.



This is a simple chart to compare IPv4 and v6. Basically both are IP. The fundamental functionalities I mentioned earlier – the three functionalities are exactly the same between v4 and v6. There is no difference.

Only the major difference is address space. IPv6 has 128-bit address space which is much, much, much bigger than IPv4 at 32-bit address space. So reflecting this huge address space the projected lifetime for IPv6 is indefinite versus IPv4 is virtually exhausted in APNIC's region. We reached to the final unit in IPv4 in April, this April, 2011.

So the immediate impact of IPv4 address exhaustion, obviously, when ISPs come back to APNIC to get more IPv4 addresses, APNIC cannot provide ISPs IPv4 addresses anymore. They need to have IPv6 transition plan by now. If not, their business continuity would be at stake.

And if they are resisting still moving to IPv6, obviously they need to find alternative source of IPv4 addresses. However, there is no sustainable alternative options available at this point and prolonged the impact as well. As mentioned earlier, internet is a big network of various size of networks, interconnected networks. Unless these interconnected networks move to IPv6 collectively, the internet itself will face the difficulties to maintain sustainable growth.

And because of lack of IPv4 address space will create difficulty for new entries coming into the market. And also it becomes an impediment of farther technological development. So obviously



IPv6 is the only viable option we have now. And please don't forget – more non-computing devices and mobile devices – Smart Phones and so on – are coming into the market and they are very IP hungry devices.

So what does this mean to all of us? Internet industry is facing a biggest challenge and we as a multi-stakeholder internet need to deal with this problem together, collectively, including domain name registries and registrars, ISP's content providers and government need to be aware about the climate and importance to move to IPv6 and take actions now.

So what TLDs need to do. TLDs need to enable AAAA records for their domain name service and insert glue records to the root servers and name resolution over IPv6 by adopting the DNS software that allows domain name resolution with IPv6 and accepting incoming connections with UDP, TCP and over IPv6. And there are TLDs infrastructures including routers, firewalls and so on need to move to the IPv6 as well.

And there are tools, operational tools such as web, WHOIS, EPP need to be operable over IPv6. And registrars are not exceptions either. Registrars need to be aware the people and organizations connecting to the internet after mid-2011 may face difficulty to obtain IPv4 addresses, may be only given IPv6 addresses. IDN and gTLDs will push more IPv4 addresses demand.



When they're starting operations in earlier stage of their bootstrap period, they need to have IPv4 v6 for address to enable their infrastructure dual stack.

So registrars need to be aware due to changes coming to their business. So for example, just randomly I picked up one major registrar. Their website it seems like running on IPv4. So if users with v4 address only and all v4 and v6 addresses, these address users access to this registrar site to buy a domain name, they can have access successfully. So they are happy, they can do domain name shopping.

But once we start having users with IPv6 address only, then obviously these IPv6 address only users cannot use content or cannot do online shopping which is available only *via* IPv4. So obviously this added to their business model too. How can they continue their business model if they don't move to IPv6? So and also more and more devices and Smart Phones are coming in and that would only push the demand of IP.

So with reflecting such kind of circumstances, we successfully organized the multi-stakeholder speakers in here from government, ISP, TLD and also the transit IPv6 transit service provider. We can have great content from these speakers.

So one last comment from me is IPv6 is a prerequisite to a sustainable long-term development of the ubiquitous and open internet. We need to move IPv6 now. So thank you.



Keith Davidson:

Thank you very much, Miwa, for setting the scene for today. Does anyone have any questions or comments? If not, I'll introduce our next speaker, Mr. Mun-Yuen Leong who is from Info-Com Development Authority of Singapore. Mr. Leong is the Chief Technology Officer and Director of Technology and Planning at the Info-Com Development Authority IDA of Singapore. He is responsible to provide technology leadership and technology planning at IDA as well as IDA's vision for Info-Com directions.

Before joining IDA, Mr. Leong was the Global Chief Technology Officer at Avaya. At Avaya he was responsible for the technology strategy, architecture standards and R&D investments. Mr. Leong was also Vice President and General Manager of the Communications Appliance Divisions of Avaya.

Prior to Avaya, Mr. Leong held senior positions at Lucent Technologies, AT&T and Hewlett Packard. Thank you, Mun-Yuen. The floor is yours.

Mu-Yuen Leong:

Thank you. Good morning. Maybe some background about IDA. IDA is kind of a combination of organizations for Singapore. It is the government CIO Office. We do CIO functions for our government agencies. We also have policy for the country in terms of Info-Com and we do industry development for info coms in Singapore.

The team that I have works a bit more on the technology side – the CTO team and IPv6, the task as IPv6 unfolds within our



organization. As Miwa so eloquently put it, we have run out of IPv4 addresses, and with a variety of new applications coming in, we do need to get on to IPv6.

So this smooth transition from v4 to v6 is kind of where we have a program going and you would see that we have two parts of the program – one is how we encourage more IPv6 adoption, and we'll talk a little bit more about that in a second. The second part is kind of how those IPv4 addresses start to exhaust. There are various techniques also to extend it a little bit but not much, so the key part is how we encourage IPv6 adoption.

So if you look at our transition program, it's kind of in two parts. If you look at the left it's a whole category of stakeholders we'll have to work with – end users, enterprise, etc., and also the network providers and hardware and software vendors.

And then you really have the two programs – how we increase the adoption and how we extend v6 and extend v4. In the increasing of the v6 adoption, it's really kind of a pull factor and then kind of a push factor, right? I think the encouraging corporations, encouraging SMEs, etc., to adopt v6 starts off with raising awareness. So I'll talk a little bit about some of our awareness programs.

And then in terms of the kind of pool factors, we've got to kind of create demand so part of it is how we work with various organizations to create demand for IPv6, as well as to create supply



of technology better than the consultants, etc., help people move along.

Then as we get v6 kind of networks or v6 running more in Singapore, then of course, one thing that we really want to make sure we kind of constantly monitor is the IPv4, v6 performance equity. V4 is pretty well-tuned today and v6 is new and kind of would be better if v6 is even better than v4, but at least let's make sure it's not worse. Then of course, there's the v4 exhaustion programs.

So in terms of raising awareness, we've kind of been doing two kind of sets of briefings in the last few years. We do kind of half yearly executive briefings out to the industry – kind of encourage various SMEs, large companies, suppliers, etc., to kind of talk about IPv6 raising kind of those companies do it or supplies bring them together, share experience so you can see we've done a variety and will do more of this next year.

We also are conducting technical conferences. We did one in April and actually worked with APNIC to do some technical training together with that as well. And we'll be planning one for next year as well.

But besides just kind of the executive, most senior business level briefing and technical conferences, we also are kind of looking at a way we can put some of the results. So we have an IPv6 microsite where we put program details, latest updates, etc., if you're interested in taking a look.



But we also have put up a marketplace and the marketplace is really a listing of where there's technology supplies, vendors, system integrators, solutions suppliers who can help companies, SMEs move to IPv6. So kind of help them understand it's an issue with all our awareness and then help them resources or they can go search out resources and we, of course, work with various companies to list them on the microsite.

We also conducted a v6 readiness adoption guide so we got kind of a consulting firm to start talking to data centers, end users; the kind of category I had in the previous chart and kind of understand where they are. You can actually again look at a website to see it, but the companies – I would say the suppliers, of course, definitely are more aware of it. A lot of the kind of enterprises, end users are much less aware, so the education process still needs to continue significantly.

The other issue has always been when we see IPv6 there are quite a few elements in the IPv6 standard so we give (inaudible) in Singapore, our local standards organization; and we have worked an IPv6 profile to define more clearly the various parameter items in the v6 standards that we believe are key for adoption in Singapore.

And from a policy point of view, we've also released some information papers on why v6, what we're doing, what the program is, the rationale behind it. This is all available on the website.



In terms of creating demand so kind of more about the pool, we've looked at our government organizations to enable v6 government access from September 2012 and also our procurement of government IT systems will be guided by this and the v6 profile.

We also basically leveraging NGN, BN or next generation broadband networks internet exchange to also enable v6, and other development issues to insert v6 as well as our public-facing websites.

We also recently announced so called seeding program to work with some of our local websites to enable them for v6 as well. I mean, part of it is how we start getting a few people up running and then getting the success stories and bringing it up to the rest of the public. So then we do that, get it to the public, executive briefings and go into detail discussions of the net. And of course, we talked a little bit about marketplace to deliver some of the solutions.

In terms of competency development, in Singapore we have a framework called National Info-Com Competency Framework and as part of this framework is kind of a broad Info-Com set of expertise, core competencies that we try to insert in overall development program. And we've incorporated the IPv6 requirements in there and organizations who want to train can also get grants, etc., as they participate in some of these programs to train personnel in this.



And we also are working on developing kind of a competence work force basically appointing training providers, getting the competency in Singapore to train IPv6 and certification of people of IPv6 as well. And as we do that we're also going to include IPv6 requirements into some of the institutes of higher learning, university, polytechnics, etc. in the course, as well as encourage some IPv6 competitions to encourage more knowledge or IPv6, awareness of IPv6 at the technical level.

In terms of the performance equity, we are going to conduct a pilot and measure the performance across the v4 and v6. We've been measuring the v4 network across a variety of ISPs in Singapore and we'll now look at measuring the v4 and v6 performance equity as well.

In terms of the exhaustion management, it's not much you can do. ISPs have some addresses left and you can, of course, kind of reuse it, but really it's to get people, the companies especially, to be aware of all the issues as Miwa put so clearly just now about if we do run out of addresses, what happens and plan for the adoption of v6 essentially.

And these last two slides just talk about the v4 allocation that's been going up, but v6 really has gone up significantly in Singapore and part of the program is to continue to encourage more of this and to really get the awareness and the demand up, but also some of the push factors as well. And that's really all I have. And I probably did it in 10 minutes, 12 minutes, maybe a little bit more. Any questions?



Keith Davidson: A very efficient and useful presentation. Thank you very much, Mun-Yuen. Before opening up for questions from the floor, I have two questions. Firstly, you mentioned in training like in universities and so on, that there's courses. Is that of the core curriculum for computer science students and so on?

Mun-Yuen Leong: So we actually not creating IPv6 courses in there yet, although some of them may. What we're actually doing is kind of more conducting sessions within their courses to help them understand IPv6, implications of IPv6 a little bit more mechanical aspects. We are looking at how we can get some of the instructors to go through some of the training or how we can get trained, certified IPv6 trainers to do some of these sessions at the universities and polytechnics.

Keith Davidson: Okay, it's just a thought I have that I've seen a tendency in universities for many to not have any IPv6 training as part of the core computer science curriculum and I'm thinking by the time those students are in the market, they'll actually need to know this, so perhaps there's a message to us all that universities having v4 and v6 simultaneous training for those people who need it.



Mun-Yuen Leong: Right and then that's part of the incentive. You start off with this and then maybe next year it gets more integrated more and more. I guess you don't really want to say it's v4 or v6 training; it's IP training and it should be agnostic v4, v6.

Keith Davidson: Excellent. Thank you for that clarification and also your graphs at the end showed the allocations of IPv4 and IPv6. Is the traffic actually shaping in v6 as well?

Mun-Yuen Leong: We have not really measured v6 ourselves cause it's more with ISPs, but I personally haven't looked at that. But I would imagine there will be some, but I have not actually monitored that. We have monitored more enabling rather than actual volume here at the moment.

Keith Davidson: Okay. Thank you. Are there any questions from the floor? There are online participants so we need to have it on the record if we can.

Paul Foody: Hi, Paul Foody. There was an awful lot there but in one of the slides I noticed that there was... you had insure IPv6 performance is equal to IPv4. I was just wondering what way is IPv6, its performance less than IPv4.



Mun-Yuen Leong: Well, we don't know yet. Part of it is when... Theoretically, from a client when you access a service at the back end, whether it's v4 or v6 enabled, depends on how peering, etc. is done between here to a local site to an international site. Could be going through a variety of parts. So part of the objective is to make sure it doesn't get penalized. This part of the measurement is the same. Technology-wise it shouldn't make a difference.

Paul Foody: So there's not any difference in terms of reliability or speed?

Mun-Yuen Leong: No, as far as we know but maybe later on there could be some discretions about that. I think we have person who is talking of IPv6 backwards.

Miwa Fujii: Any other questions?

Paul Wilson: It's Paul Wilson here from APNIC. I've got to head off to the GIC so I hope I'm not stealing anyone's thunder here, but on the question of v6 performance and readiness, as Miwa said, thanks very much. The provision of IPv6 services by registrars and registries is pretty important as it is for anyone who's an online service provider of any kind because as we move through the



transition, there are going to be users changing their configuration of access to the network.

It's actually not as simple as a case of on or off, but it's actually a case of degrading performance for users, for instance, who are using IPv6 plus private IPv4 addresses in a dual stack configuration. And also for users who may be using some of the very handy but pesky automatic tunneling systems which are automatically enabled on many operating systems.

So I hope I'm not stealing thunder about the IPv6 day but the recent IPv6 day gave a huge chance to test in the field what the performance was and what the prevalence was of user connections that are out there actually existing on the network. And so we did actually find and APNIC was involved in a measurement project which is still accessible on our labs site, that's labs.apnic.net which allows a website holder to monitor and to count the different types of connections that come from their users to their website and it will give you a full summary of how many v6 connections we have, how many v4, how many tunnel connections of different kinds, what the performance of those connections is and the failure rates.

And one of the performance issues is definitely that if you're a website and you're kind of relying on the internet to solve the transition problem for you rather than deploying IPv6 on your own network, then your users can have a very hard time because they'll be given private addresses with network address translation, they'll be given automatic tunnel services which often don't work or have



a very compromised performance on the order of many seconds of additional waiting in order to reach a website. And I think as any online service provider knows, the performance of your website to your users is absolutely critical. A user will only wait for a few seconds before they decide, “Well, I’ll go off and try another registrar or registry.”

And so if you’re not deploying IPv6 on your site, then as Miwa said, you can have a much degraded performance. As I said, it’s not just a case of on or off, it’s actually a case of degrading performance and that’s something that’s going to be happening on the internet in increasing numbers now that the IPv6 deployment is having to accelerate in lieu of IPv4. So that’s just the comment I wanted to make. Thanks.

Keith Davidson:

Thank you, Paul. Any further questions? I thought there was a hand up over here. No further questions? Well, thank you very much, Mun-Yuen. And can I just make a final comment? Governments have, I guess, options that vary from being able to legislate or doing nothing about this issue. And I think what we’re seeing is a very interesting multi-stakeholder approach of the government playing its part and showing real leadership in a methodology by which you’re stimulating the interest and the uptake. So, congratulations to you.



Mun-Yuen Leong:

Thank you very much and you probably could see if you go to our website, we just recently issued a public consultation as well in a process called No Islanding and our Minister announced it, I guess, on Monday. And it's really not having v4, v4, v6, v6 and get islands what happens.

So we actually have a public consultation on some questions out in that and you could read that as well on the website. It's something that we are actually asking for comments from the public.

Keith Davidson:

Excellent. Can we show you our appreciation? Thank you. Our next speaker is Hideo Ishii, if that's roughly right. My apologies. She is Vice President of IP and Cloud Services at Pacnet Global Singapore. Prior to her current position, he served as CEO of Pacnet Japan for two years and head of IP Engineering for eight years. He has been with the company since 2000 and is active in the area of network operation communities and NPLS - I guess that's NPLS and IETF terms rather than ITU terms? That might be another debate for another day- and transport technologies areas.

He is a program committee member of NPLS Japan Technical Conference since 2001 and was a founder member of JANOG, the Japan Network Operator's Group and Co-Chair of [APOPS] until 2008. Ishii, welcome and the floor is yours.



Hideo Ishii:

Thank you very much for introduction. Okay, first of all I come from Pacnet which is a service provider as well as a cable provider in Asia/Pacific. So I'm on today as a service provider. We have two service provider phases.

One is a regional transit service provider as well as a domestic Singapore service provider which is a former Pacnet Internet Singapore. Pacnet Internet Singapore is the oldest ISP in Singapore at a commercial point of view. So I'm covering Pacnet Global side because Pacnet domestic is still around IPv4s and so on.

What is Pacnet is that, as I mentioned, we are a service provider over the internet as well as a cable provider in Asia. We have three cable systems in Asia/Pacific. Within Asia/Pacific, two cable systems, ESE and CSE as sea to sea between multiple Asian countries, as well as we have transfer capacity from L.A. to Japan. We are utilizing this capacity for IP and the other services right now.

We are now in the top-something in IPv6 service provider in Asia. This is an Asian region top 25 transit upstream AS number. So we are now, it's around, something a 10 or seven around. So we have 51 transit AS number right now in Asia/Pacific.

Next I would like to explain our IPv6 transit story since 2001. First I joined the global crossing Japan 2000 and I was working with a group of crossing guys to deploy IPv6 globally. So and then we finally got the IPv6 address from APNIC and then we tried to deploy the v6 network using global crossing platform.



Then at the time that we tried to deploy IPv6 on the IPv4 network was simply GRE tunnel between routers. Because at the time it was very, very tricky to tunnel IPv6 in IGP in routing in the domain, as well as eBGP. So then we have been working since then for IPv6 deployment since Asia/Pacific region due to some reason at global crossing and Asia crossing.

And we have our own completely independent from global crossing network in Asia Pacific and we tried to adopt best technology into IPv6 network into this region.

For five years we had been working on that using multiple vendors, smart view in OS and so on. And also at the same time we providing IPv6 transit in Asia/Pacific as a trial base so customer want to connect to the IPv6 with Pacnet, we give them IPv6 route using GRE tunnel or the other way.

Then finally, we go the in operational team approval to deploy IPv6 routing protocols into the network, then finally 2008 we deploy native IPv4 IPv6 throughout that network group globally, especially almost rather in Asia/Pacific region anyway.

And then at the same time we could provide IPv6, IPv4 dual stacking for a solution to the ISPs as well as a corporate customer to enjoy IPv6. Then finally, last year we launched the IPv6, IPv4 native IP transit in this version to the customer as product. So it's very, very long story, but we are now providing IPv4, IPv6 dual stacking service in this region.



Okay, as I mentioned that we never generate names fast because this is no impact to the IPv4 network. It's completely isolated from IPv6, IPv4 has routing perspective, but this is of some issue as some BGP router reflectors and an (inaudible) mesh and a (inaudible) JIRA on the back and so on. So service provider point of view, we have to deploy IPv6 routing protocols into network in order to optimize net performance as well as routing visibility.

So then second step was deploying IPv6 IGP, interior gateway protocols, as a BGP. But we couldn't deploy these routing protocols globally one time because we haven't seen any potential issues and so on at the time. We should be carefully deploying those protocols into the network because this is a production network anyway. And finally we turned on every single router IPv6.

This is a very simple steps of the service provider, prospective anyway. So then in terms of the operation and the provisioning perspective, this also very crucial for the service provider and then we did a long, long time to providing trial IPv6 transit to service provider. This is an operation support, to be honest. It is just a best effort type of solution to the customer because at the time no match in high speed want to use IPv6 too much critical. It's just a trial, just doing ping, just doing aware that IPv6 are around and growth and interface also emailing, inquiry, trouble report and so on. It just goes through the email, not phone call anyway.

At the same time we have been doing IPv6 technical training for operations, programming and customer support and so on. This is



too much a headache for us because new protocols, new concept, new routing, new OS and so on is very much complicated at that time. So we also need some hands-on training environment for those guys because they want to touch a router, IPv6 router that is not only on the paper. So we deployed IPv6 in Japan to do some tryouts using those environments at the time.

And also some commercial documents - should we support IPv6; and cautioning database and cautioning some documents and so on we completed.

So then before we tried to provide IPv6 status as service provider, we have to complete every single objectives and some activities. First of all we need IPv6 address anyway. We obtain IPv6 address from APNIC and also providing IPv6 transit to the customer, we need a global route from someone because independent IPv6 network is not internet like BPN anyway.

So also need to set up IPv6 peering router, to exchange routes to optimize the latency and so on and also IPv6 numbering plan, the interface, routers, and so on. It's too much processes in the design. But it's similar to IPv4 process anyway. So then no service impact. This is the key, it is obvious anyway.

So what is happening during IPv6 deployment is very important as operation and sales and so on. We had considered some way how do we deploy IPv6 without service impact. Some of technology we have is 6PE using MPLS or pure like IPv4, IPv6 in IGP. And



also we enable IPv6 on existing IGP on OSPF or IS-IS or completed in isolated routing protocols within domain.

Let's say I am now using IS-IS as IPv4 top of OSPFv3. But finally we enable IPv6 in the IS-IS using in multi-topology way. And also we have to change config., let's say Cisco anyway. Cisco support IPv6 but configuration is not the same; needs to support in address-family IPv6, IPv4 and so on. It changed the configuration temporary as well.

And we did those processes and we had some observations as a service provider point of view. First of all, traffic volume is still very low even though while IPv6 Day. So as Pacnet point of view, it's not cover whole router to be honest but I resolve some points, 0.3% traffic compared with IPv4 right now.

I think Marty has a more big volume I think but in Pacnet in Asia-Pacific is doing this. Also lack of customer demand in this region. As my understanding, as a regional service provider point of view, New Zealand, Australia and Japan is a huge demand. But the other countries it's still struggling, but it's IPv4 anyway.

And routing optimization as some have covered this point as well, is still not completed. I mean, during IPv6 Day I tried to do some trace route from my server to somewhere. I was surprised because Singapore to U.S., let's say, like I see in [ICMP] traffic go to U.S., Europe, back to the U.S. it creates latency because this is the wrong way because this is a trans-route. I'm not sure return traffic from there still traffic go to U.S., Europe, back to U.S. It's



amazing. So we have to optimize IPv6 routing as well in order to match IPv4 currently anyway.

And then monitoring tools is still service provider trying to achieve because simple SMB means cannot get the IPv6 traffic data. Some tricky way we have to implement. Let's say we try some filter into the INET6 to see IPv6 traffic or create a (inaudible) map – this is one solution, I think. But the (inaudible) IP map to see IPv6 traffic to get IPv6 traffic data. We are not doing such in a tricky way, so but maybe something, a new more simple way in the market, I hope.

Netflow data – we have to count on IPv6 netflow correcting using version 9 and try to get the IPv6 netflow data. But needs to upgrade IOS, so let's say an IOS12S to IOSXI anyway. It's also big impact to service provider platform.

Also this is a... I'm not sure what is the best way. I think Marty has a good idea. Should BGP related policy be similar to IPv4? Let's say BGP local preference value and community strings and so on. So we are now using IPv4 a similar policy, but I'm not sure if this is right or not.

Also in Asia/Pacific region, some countries are good; some countries not demand; some countries, let's say, in the need of v6. It's too much variety requirement in the Asia/Pacific region so we have to make sure the Singapore IP was kind of a requirement they have why we should deploy IPv6 quickly or spend the next six



months or one year. I'm not sure. But it's very much a very requirement around this region.

And next question is who is IPv6 Tier 1 ISP in globally? So I think it's very easy to predict that in (inaudible) is one Tier 1 ISP globally, but I'm not sure the other. IPv for Tier 1 ISP is not IPv6 Tier 1 in some cases, so we have to make sure that who is it IPv6 Tier 1 run globally. Otherwise, routing is go to Europe, back to U.S., go to Asia, but it is very much complicated.

Also we need to deep dive into IPv6 requirements of broadband customers. We have broadband in IP service in Hong Kong, Singapore and Australia so, you know, 6 RD or whatever is very complicated but we have to looking at what is the best way those countries.

Lastly data center and hosting we are now providing a data center space and hosting and so on. So we have to correcting what is the customer demand in the Asia/Pacific region and to prepare the right involvement, right solution to them. That's all. Very quickly go through, but I think a light time.

Keith Davidson:

Thank you very much, Ishii. I think I've just got one question. You were talking about the long routes to Europe, USA and back and so on. What's the timeframe for resolving some better, shorter sort of hopping process?



Hideo Ishii: Normally from Singapore to U.S. is around 200 milliseconds, but it is 300 something milliseconds in addition of 150-180 milliseconds in latency, even if site in the U.S.

Keith Davidson: So there's ongoing discussions to try and...

Hideo Ishii: Yeah, yeah. Resolve the situation. I'm not sure what is the best solution. We peer more or we buy (inaudible) from someone like an IPv4. I'm not sure right now.

Keith Davidson: Okay, thank you. Are there any questions from the floor?

Female: Hi. Good morning. One of the objections of the ISP in Malaysia is because of the cost of converting IPv6 infrastructure to IPv4. So out of curiosity I would like to know what is the figure to setting up to IPv6 infrastructure for ISP.

Hideo Ishii: Very good question. Thank you. So I thought that I should add, but it's too much commercial point about it. So some cases, it depends on the infrastructure, how old those equipment are anyway. Let's say, some routers can support IPv6 just software upgrade only and infrastructure-wise, just routing, so not much,



CAPEX and OPEX required. However, some of the related servers and internal process change, data base change and so on is required at some OPEX and CAPEX, as my observation.

Keith Davidson:

Could I tease that out a little further. Given that you're potentially replacing equipment all the time, I know some service providers and so on will go as far with IPv6 as seeing that there's a tick on the box, on the kit, to say it's IPv6-capable and that's all they want to do. I think the test is to actually have an IPv6 deployment plan and when you're next upgrading that part of your network actually doing it at that point. Is that a fair assessment of what you're seeing with your connecting customers?

Hideo Ishii:

Almost, to be honest, our customer almost service provider, they came to provide IPv6. They're already their IPv6 platform in their domain, so not much crazy come from within. I think the next two, three years, maybe we will see some such changes, I think. But, again, in Asia/Pacific region, it's too much gap in the situation. It's not like a simple comment to them.

Keith Davidson:

Thank you. Any further questions?



Sivasubramanian Muthusamy: I'm Sivasubramanian from ISOC India Chennai. For new entrants to the ISP business, it should be easier to get started on ISP in the first structure. What would it cost for a new entrant who would like to set up a regional internet service provider to start entirely on IPv6 infrastructure? Is there a very approximate figure? Particularly the choice of IPv6 in prospect or IPv4 in prospect if everything is new.

Hideo Ishii: I'm not from [Bender] so I cannot say something but if platform must support IPv4, such router scale must flow in IPv4 demand in the performance. Because I think IPv6 a very small router can support because traffic is still small and routing table is also very small.

But if you want to support IPv6, it's completely different game. So if you enter those market, I think IPv4 needs a big size router, so you have to buy a big router to support IPv4. So big server can support IPv6 anyway, so it's more easier step.

Keith Davidson: A question I've noted that I think should be asked more generically in other environments and perhaps would be a useful little case study to find the answer. There's a further question over here.

Modou Sall: Good morning. Modou Sall from France Telecom Group. I have first a contribution. According to the ISP who wants to move to



IPv6, I think that something that is very important is to have a plan to move to IPv6, a strategy. And one thing which is really important is to be able to integrate in the requests for proposal for the occasion of new (inaudible) to be able to have a dedicated paragraph related to IPv6, comprehensive, for this new (inaudible). This is something which is really important.

I have a question. You talk about your program customers have some issues. My question is do you perform some IPv6 tests with if you have existing DSL customer entities and (inaudible) customers? Other question – why 6PE and not 6VPE? You tell that in your presentation you have 6PE. Okay, why not 6VPE, which according to me is more simple if you want to set up some MPLs, VPMs without no change in your core network. You will remain in IPv4 and only enable the DS stack in your P routers. Thank you.

Hideo Ishii:

Okay. First question, unfortunately, I don't have any test results because we are still looking at. Second question is why we haven't used 6PE. So it's a very interesting question. So at that time, to be honest, we did some tests with the vendors, but one difficult issue we had is that, at that time, set up the TE tunnel between two PE routers, IPv6 address is required and how do we distribute IPv6 routes to the other PE router? This is the challenge for us because why we have to deprive v4 with v6 within the domain to set up that TE tunnel in between. It doesn't make sense to us.



So of course I set up that IPv4 TE tunnel then to enable 6PE to exchange BGP if possible. But BGP's next stop is at IPv6. So this is too much complicated for us. BGP is IBGP's neighbor but 6PE, if set up TE tunnel using IPv4. We cannot set up IPv6 neighbor BGP. If so we have to turn on IPv6 from TE tunnel but IPv6 T tunnel needs IPv6 IGP. So it's something very tricky, so that's why we changed.

Keith Davidson:

Thank you very much, Ishii. There's no further questions so please can we thank Ishii. Oh, there is? No, probably not. Sorry. Hopefully we'll have a bit of time at the end. Thank you very much, Ishii. Can we show our appreciation?

Our next panelist is Joe Waldron, Vice President Core Name and Services at Verisign. Joe manages the domain name registries including .com, .net and the root zone. This includes responsibility for all related services, most recently the deployment of DNSSEC.

Joe has graduated with a BS in Electrical Engineering from the U.S. Military Academy and graduated with an MS in Systems Engineering from the Colorado School of Mines. Before joining Verisign, Joe was Field Artillery Officer with assignments including canon, missile and raid air systems and was an instructor at the U.S. Military Academy's System Engineering Department managing the Operation Research Center and Combat Simulation Laboratory. That must have involved some fantastic Star Wars type games. Thank you, Joe. The floor is yours.



Joe Waldron:

Thank you. I'm going to talk a little bit more of a high level description of IPv6 from a registry point of view so I won't be going into as much technical detail.

First we've all heard that the IPv4 space was exhausted and I think that if you look at the parallel between concern over the Y2K bug and implementation of IPv6, the only difference is that for Y2K everybody knew when the deadline was. Everybody knew that on January 1, 2000 something might happen, so everybody prepared for that.

With the exhaustion of IPv4, it's been something that has been rumored and talked about and people have worried about for a long time, but there was no concrete driver. Well, back in January we had the first hint of that and the first real big announcement that the last block of IPv4 space is gone. Now what that really means if you look at the way IP addresses are allocated, that's the allocation from IANA and the last of those allocations were made.

And then to each of the four regional internet registries, they still have IP space, although APNIC did exhaust theirs in April. And then from there all of the allocations are made from those regional internet registries to the ISPs and ultimately to the end user who has a website that they're trying to get on the internet.

So if you look across the globe, this is not an issue that's unique to any one; it's going to be everyone's problem to be prepared for a world where there is no more available IPv4 space to be allocated.



So if you look at each of the regional internet registries – this is a chart that is available online – you can see when the exhaustion timelines are predicted. Then they can accelerate very quickly as you get toward the end line and everyone wants the allocations.

So IPv4 space is about – the end is in sight in terms of available IPv4 space and there is a lot of discussion around what that market looks like once the market is completely exhausted.

This brings us to what does it mean and that's the question I always ask myself is what does it mean for Verisign as a registry operator to implement IPv6? So as I look at the infrastructure and the ecosystem that we operate in, we have a direct relationship with registrars – I think that's a very well known and documented relationship – but we also have to be able to interoperate with ISPs, with root operators, so there is a very loose connected system of relationships that are both formal and informal.

And all of these connections up until when people started implementing IPv6 which has been going on for a while, but, the IPv6 deployment needs to be addressed across all of these. And even within the registry – and I'll speak a little bit now about what it means to implement IPv6 as a registry.

The one Miwa had on her slide is that you have to be able to accept Quad A records. So I need to be able to accept a Quad A record from a registrar that wants to associate that with a name server. In addition to that though, I also need to be able to generate a zone file that has that Quad A record in it and then make sure that when



someone does a look up on a DNS server that it's available in the response. So that sounds fairly straightforward and Verisign has been doing that for a number of years.

But the next component that I would look at being able to also make that DNS response available over an IPv6 network and not just an IPv4 network. So I put my name servers into a different network space, I allocate or assign those addresses and now I've got to put the Quad A record for those name servers into the root zone. So the root now comes into play in the ecosystem and the root zone, I think it was in 2004 or 2005, started allowing adding TLD name servers, could start adding Quad A records.

So that's again another component of the registry view. And the other piece, though, is what happens when I have registrars, all the registrars that connect to Verisign today operate and connect to us over a v4 network. That's just the legacy system that's always been in place. So what happens when I have a registrar that wants to connect over v6?

So all of the other services that I offer connection to the shared registration system, connection to our web tools, access to servers that they use for reporting – so all of those other services also need to have IPv6 connectivity.

I think we've had some discussion about how to do the... what the choices are. We've really identified four different choices or approaches for migrating to IPv6 and I almost think migrating is a little bit of a misnomer. But you can build parallel structures, you



can have dual stacked or use of tunneling and network address translation – we talked about that before.

What's important about this is that any of those solutions for existing customers trying to get into an IPv6-enabled environment – all of these different techniques require that you have both IPv4 space and IPv6 space, so that's an important consideration, especially as you think back on the depletion chart.

I feel compelled to also make a reference to IPv6 Day. Verisign participated in this. There is a link here that has some information about what Verisign did during IPv6 Day. One was to enable our corporate website to make sure it was accessible over IPv6.

More importantly, I think, from my perspective as a registry operator, I look at the traffic that we saw. On a typical day we see about 60 billion DNS queries across our infrastructure and on World IPv6 Day the IPv6 traffic represented 607 million queries, so it's a very small percentage and I'll show you another chart on that. We also saw similar rates on the A and J routes but a much smaller scale.

This is a chart that is generated and I'll describe it very quickly, from a web crawling service that we offer called Internet Profile Service. On the left side are the monthly statistics in the blue line that show the number of web servers that we've identified within .com that were accessible over IPv6. So that they were a web server that had an IPv6 address associated with it and was accessible.



And you can see from November of 2010 through May, there's a pretty reasonable growth rate in the numbers, but we're talking here in the range between 100,000 and 200,000 web servers across all of .com. Because we haven't completed the data for June yet, I wanted to look at some impact that we would have seen during World IPv6 Day, so if I just looked at the subset of what was on the blue line, I'll plot this on the right to make it hopefully clear – we see that a similar growth over those same months.

But then when we get to June of 2011, the accessible websites out of the top one million .com names goes from 1,000 to over 5,000 so that's a significant increase. So what that tells me is that we're definitely seeing more interest in World IPv6 Day than the about 400 participants that were there when we see this many increased web servers that were available on that day.

Within the data of the .com registry itself, we have name servers that have Quad A records associated with those and you can see over time a very small percentage. If I go back a year and a half ago, it was only .06% of the name servers had a Quad A record. And now we're all the way up to .17 which is still a very small number of name servers that have Quad A records.

But if you look at the graph on the bottom though, how many domain names are actually hosted on those name servers, that has increased from less than ½% in January of 2010 to over 2½%. So we're seeing the name servers that are hosting a large number of domain names are being... when they're enabled you get a



significant improvement in the percentage of domain names that are ultimately accessible over IPv6.

And I think the other important point is looking at the number of name servers that have both IPv4 and IPv6 which is an indicator that most people are implementing a dual stack approach for their adoption.

Just as I said earlier, Verisign has been operating the registry and we've implemented the acceptance of Quad A records. We did that nearly 10 years ago. We also have implemented and provided information for our registrars to help them with the integration and connectivity with us and some of that is covered in our software development kit and other services.

The IPv6 transport I referred to earlier in terms of the connectivity for the name servers has been in place since that was essentially allowed to be added to the root where we could add the Quad A records for our name servers into the root zone and on our IPv6 website we have a graph that you can see. It's probably a little bit hard to read, but it shows monthly percent of traffic of our overall DNS traffic to our infrastructure.

So if I go back to January of this year, it was only 0.6, so less than one percent and increasing. And I expect that when the data is complete for the month of June, we would see some increase in the amount of IPv6 as a percentage of our overall traffic.

Finally, just in the interest of time, I'll just make a reference and there's a link here for a white paper that we published, again trying



to help prepare the community for the challenges and solutions for implementing IPv6. So I'll leave that slide up and ask if there are any questions.

Keith Davidson: I think we're chancing the room unless there's one burning question that relates to and we'll stop after one question.

Male: The last few years two guys at Google have been doing work on providing Google services, particularly search of IPv6, [Gus Lees and Lorenzo]. They were finding latency and EPV failures that were unacceptable in Google's eyes. Their unusual case. They can value thousands of dollars a minute of what this stuff costs.

The solution that they adopted was to only serve Quad A records to resolvers residing inside NS networks which had specifically requested IPv6, either access or peering with Google. My question is have any of your .com registrants – I assume larger ones – asked you for the same capability for their NS records that Quad As only get served to designated networks? Or have they taken on all (inaudible)?

Joe Waldron: First of all, all of those questions from the registrants would come from the registrars so we don't have a direct relationship with the registrants...



Male: Okay, so are your registrars asking on their ledger?

Joe Waldron: But no, we operate our DNS servers in compliance with the RFC so I think you're really talking about something that is a bit different from what we would be able to do just as a standard DNS operation, if I'm understanding your question correctly.

Male: What Google's doing doesn't reach any RFC but it's...

Joe Waldron: I think the appropriate place to ask the question is for the people that are building resolvers, should they change the way that they're building their resolvers and ask for questions, but from an operator of authoritative name servers that are responding to queries, I'm going to give the answer based on the question that's asked. And I give the same without regard for where that question's coming from.

Kevin Davidson: I have the question noted and I think it's a very pertinent question and again, it might be one that this panel recommends as a follow up, so I thank you for raising it. Can you all please join me in thanking...



[Applause]

Kevin Davidson:

Now our final presentation is from Martin Levy from Hurricane Electric. Martin has been working for Hurricane Electric and very active in shepherding IPv6 deployment during IPv6's dramatic rise to commercial usefulness. With nearly 30 years of experience in ICPIP protocols, it's been his and Hurricane Electric's goal to see IPv6 as a solid part of the global internet's day to day usage.

During the last few years he's also expanded the understanding of many executives and regulators towards IPv6's importance. Martin was born and educated in the U.K. and now lives in California. Welcome, Mart.

Martin Levy:

Thank you. You're only going to hear the answer yes from me cause that's my mindset towards v6. But those people that know me have heard that before. But let's go run through some stuff. The questions, by the way, that have come up so far have been great but there is just a couple of key points that will come out, I think, and hopefully I'll answer them.

I'm going to run through a couple of things – one really quickly, literally a page, about us, but I won't even talk much about it. I want to talk about why now is the time to talk about v6 and more so even than in the past. The question that was asked by a



gentleman about how much does it cost to build an ISP from scratch. In a way that question doesn't need to get asked anymore because in 2011, it's pretty impossible to buy something or you shouldn't be buying something that isn't v6 enabled.

I'm going to talk about a global v6 report that we've been running for a long time. It has some great numbers that are useful in the TLD and in the DNS environment and then obviously some comments and a summary.

So a quick one page about us. We're a global IP backbone. We run an enormous amount of IPv6 traffic and when I say enormous, as a percentage, it's still small. But it's enormous compared to what people think there is as far as v6 traffic is concerned.

One of the best pluses about running World v6 Day, earlier this month was that a lot of people woke up and said, "Oo, there's more of that than we thought," and we've been saying that for quite some time.

As a global backbone, we exist here in Asia, we exist here in Singapore, throughout the U.S. and Europe as well, all the usual stuff. Go look at our website if you want to hear about us. I'll talk about v6.

Yes, that's just the simple response I always give everybody when they talk about v6. But saying it in 2011 is way easier than it was in the previous years, and that's the honest truth. But as I said I said earlier, if you're going out and building infrastructure today, if you're upgrading something today, if you are reviewing something



today subject to change, whether it be software, firmware, hardware, etc., or a training course for engineers – sorry, or a training course for accountants – the answer should be yes as far as v6 is concerned. It just can't be left off the conversation at the moment.

And in the years that we've been doing this, I've come to a conclusion if they keep asking and you ignore when people say, "You know, haven't had much call for that," when you really dig down, it turns out they have had call for it, they just haven't listened.

So by the time you get to now, there should be no part, registrar and DNS included, where you go and talk about v6. There are some caveats, there are some good caveats, but no different than the caveats that have been seen in other parts of the industry.

So let's look at some numbers. Let's see whether the proof is there. These are some simple numbers that we've actually collected for quite some time, but I just wanted to put in a nice easy graph up there. This is the number of ASNs, the number of Autonomous System Networks on the global internet that run v6.

And that number's been increasing; it's been steadily increasing. And then something interesting happened around the end of January, beginning of February when IANA finally ran out of v4 space and you can honestly say that there was a hell of a lot of press that showed up and said, "Maybe I should write something about this lack of v4 space and this new thing called v6."



Well, okay, v6 isn't new but that press made a difference. If you look at the work the RIRs did, if you looked at the work that many people did to say v6 is important, look what happened in January-February of this year. We saw an honest-to-gosh knee in the curve. When you measure it as to networks that have deployed and running v6, the numbers changed.

This URL that's at the top there has been online I think now for eight or nine years and it has been a collective report on v6 progress. It actually deals with some of the questions that were asked – v6 performance versus v4 performance; number of networks. It has a whole new section on DNS that I'm going to talk about in a moment that's very relevant inside the TLD and inside the registrar environment.

But what's important is that that graph is going up and it's going up good and proper. So keep a look at it and the one interesting take away from this is if you have a service provider today providing you access to the internet and you ask, "Do you have v6," there should only be one answer. If they give you the other answer, the word I try not to use, find another service provider. It's about now that you should be doing it.

Okay, so let's look at what goes on in the DNS world, and let's start looking at the TLDs. So last couple of years, actually about a year and a half, sorry, two and a half years worth of graphs here, we've had a nice steady growth in the number of TLDs – no comment on what happened this week – but the number of v6 glue



records sort of, it increased. It just didn't increase to the point of hitting what I'd like to see which is 100%.

Now maybe 100% is a hard task, but let's go hit a few subtleties about this graph. The green line is the number of glue records in the root name zone that have existed, that have v6 addresses and it has increased. The red line is the number of those name servers that actually do have v6 but aren't yet included in the root glue records. And there are wonderful reasons why that hasn't happened and it is non-technical. It's nearly all paperwork. I don't talk about paperwork so that gets me out of the end of discussing that.

But the bottom line is that both of those are graphs that the closer they get to 100%, the closer you get to the point of having at the root and the TLD level, at having v6 support right at the name servers. Now if you see an explosion in the number of TLDs, then you want to make sure this graph never goes down as a percentage, that the name servers that people deploy have the ability to support that.

So the question is why do you need to support it. Well, again, it's 2011. We shouldn't have to be repeating the many, many slides that talk about the fact that there are going to be users out there that don't have as equal connectivity to v4 as they will have to v6. And yes, I said that the right way around.

The fact is that as you look at the changes, the broadband providers are going to have to implement as they run out of v4 addresses that



it's going to end up that v6 will end up with a better performance. Why? Because they won't be carry a great NATs. There won't be large-scale NATs that have to exist in the v6 world; that will purely something that exists in the v4 world. I'll talk about that in a second.

Alright, so let's look at the next thing – so I picked on .com – nothing personal, by the way – it's just a nice big domain and it's easy to pick on and it's got wonderful legacy in it and it's got a lot of stuff in it.

So every day we grab the domain under contract and we do this massive review of what's in there and we do that basic measure. How many v4, how many v6 addresses are in there in the name servers? Now we also go and look at how many A-Records and Quad A Records exist for www.example.com or any other domain. But that's another point. Go to the website and look at that if you want to.

But what we do find is, we find a little messy set of information and this is the take away on doing v6. And hopefully the rest of the panel has convinced you that v6 is the right thing to do the moment you walk out of this room, get your engineers to think about a few things.

You don't want name server glue records with slightly invalid v6 addresses. Now there are plenty of invalid v4 addresses, but we've sort of already been dealing with that. We have for, well, not for



30 years, but technically for 30 years, definitely in the life of registrars.

In the v6 world you'll see people putting in loop back addresses for name servers, old deprecated [six bone] addresses and my favorite – the 6 to 4 transition address, which all that does is generate a load on some transitional relay of which we seem to run the largest amount of them anyway and then generate a v4 packet in the end, making the whole process of having a v6 Quad A record as a name server somewhat redundant.

If you're going to do anything, look at doing some help, put this in documentation. But the interesting take away from this is put the v6 stuff in the documentation if you're going to do it. There are registrars who have v6 support and they may not have v6 support on their www.example.com or on their registrar webpage, but they definitely have it inside their registry themselves.

And, actually, I'll put credit where credit is due – GoDaddy actually does let you put a Quad A in for a name server record for glue and has had that for a while, even though their website may not be running v6, but at least they allow it.

On the other hand, and again I'm only picking on them cause they're big, Network Solutions didn't. And this is a story from three years ago, so it really doesn't matter anymore, but it's humorous. This is how not to do v6 and then how to get v6 as a high priority.



We had users who wanted to put v6 Quad A records into their glue, like anybody would and they were having a hard time. And finally Network Solutions posted an email address, IPv6request@networksolutions... and said, “Send us an email if you need to do this. So we forwarded that to 50,000 tunnel broker users and they pretty soon fixed the email with a full web interface because there is a demand out there, even if you don’t realize it.

This is a problem that doesn’t exist now, but think about it as you’re doing this. Support v6 like you do v4 and make sure it’s in the documentation so your users know about it.

Anyway, the numbers are interesting. If I zoom in so you can actually read it, we’re looking at nearly 900,000 name servers that are in the .com domains that have local host ::1 as their name server, which is actually clearly wrong. About 71,000 in these mapped v4 addresses, which again are about as redundant as having a transitional address.

Some of these could be teething; some of them could be whatever, but if you go look at this data, if you start accepting this, think about ways of monitoring this. We do it; we’ve been doing these reports for quite some time now. They’re always useful.

So let’s look at why this is all vital. V4 exhaustion – shouldn’t have to repeat it but I’ll repeat it – v4 exhaustion. But the reality is this – that you’re going to start seeing users and you’re going to start seeing infrastructure where v4 is not going to be on par with how we think of v4 today and v6 will be the way out of this.



We do, in fact, actually, I put this data in here before anybody asks the question, but we actually do measure the percentage of reverse DNS name servers. We measure a bunch of different name servers and we look at how close they are on v4 and v6. So we see about 73-74% that run about equal v4 and v6.

And if you look at the full stats, some are running actually faster which is impressive and some obviously aren't. But the whole point about this is that you'll see the stats and you'll see the fact that we're nearly there; we just need to get the rest of these percentages in there.

Alright, last thing that we do. We also look at the number of v6 glue records that are in various domains. So I've just picked the top 15 just cause it fit nicely on the page, it was still somewhat readable, obviously everything's still online. But what's interesting about this is, again, we do this work and if you're wondering why we do this work is we do this work ancillary to our core business of providing internet access. But we do this work so that we can say, "Rah, rah. Good v6 stuff is happening out there."

So these numbers, as a percentage, yeah, we'd like to see them to be larger, but the reality is that these numbers have been increasing. If you look at them historically, the numbers have been increasing nicely over time and we now just need to see that push.

The last thing that I have to talk about is what happened on World v6 Day? Well, very interesting things happened in the DNS world. Nothing, zilch, it stayed the same. The whole point about this was



it was a World v6 Day for content providers. So we saw massive jumps in the amount of web traffic, we did not see and did not expect to see massive jumps in the amount of DNS.

But what's fun is Jeff Houston gave this talk not a week ago in Denver at NANOG, and he showed that there was a big jump in the number of reverse DNS queries that were coming into at least one RIR, APNIC in this case. And what that meant is that people were actually starting to process the log files from seeing these v6 addresses and realizing, "Wait a second. What are they?"

Now maybe it was their software that realized it, but that's okay. The plus here is that there was at least a mechanism that ran through all the way from a website being enable for v6, a website processing a v6 request from a user who is enabled for v6 generating a log record and then finally getting processed because somebody wants to know who used that website.

You know, that's not bad. That's not bad at all. Interesting side note in regards to measuring traffic – as I told you earlier which I didn't talk about – sorry, but it was on there – we did a complete upgrade about five years ago of all our hardware and one of our requirements was to make v6 on par with v4. No ifs, no ands, no buts. And one of the things we needed was monitoring.

So these graphs on there were purely from all of our flow data because unless we can see this, we can't really tell that there really is v6 traffic out there. And we've published some of our numbers. We run about just shy of 10 gigs worth of v6 traffic globally at the



moment which is pathetically small compared to the global internet, or even our part of it. But it's 10 gigs; that's not pings and trace routes anymore; that's real users.

Alright, final thing – just do it. Do it. Do it please. Do it. Thank you.

Keith Davidson:

Thank you very much. I don't know that you can question that. Being instructed to do it is the challenge and besides, I think we've run out of time for any questions. So could I ask the panelists to perhaps be available for anyone from the audience for a one-on-one question after we close the session?

Thank you very much, panelists. This has been a most useful session. I came in just wanting to learn how to spell IPv6 and I think I'm reasonable competent with it now. I think it's really important to note too and remember that the last time in this environment and the ICANN environment that the new gTLD process was held and we had the new gTLDs of .biz and .info, there was a remarkable uptake of domain names outside of those new domain names. The ccTLDs and so on all saw growth.

So I think what I'm taking from this is there is a new gTLD process that's not that far away and registrars, if you're not offering v6 to your customers, you may will be overlooked because we will have a savvy group of potential new registrants coming forward and they will probably want to check that you're v6 compliant.



So you've got plenty of time because the new gTLDs probably won't be happening until about this time next year, but if you start thinking about your v6 deployment, that's not a lot of time. So I think the challenge is to go back and start work today and make sure you have a plan.

So with that, thank you very much, panelists, and I'll hand back now to Miwa.

Miwa Fujii:

Thank you very much for everyone stay back so 10 minutes over actually your morning breakfast time. I already start receiving emails – “Where is the presentation files?” Sorry, we didn't know how to upload the files, so we will learn how to do it and with the panelists' agreement, we will upload the presentation files so it's available for you. Thank you very much.

[End of Transcript]

