The Standard

Volume 2 No 1 Internet Protocol version 6 Newsletter June 2011

The Standard is the journal of IPv6 Now Pty Ltd, the only dedicated IPv6 company in Australia. IPv6Now has specialised in IPv6 training, consulting and services since 2007.

ipv6now.com.au
ipv6now.co.nz

ipv6.org.au

Learning From Experience

Melbourne
17-19 Oct 2011
ipv6.org.au/summit

What is IPv6?

IPv6 is the replacement protocol for IPv4, the current, almost-exhausted addressing system for Internet devices. Modern IPv6 offers massive numbers of addresses, and better networking, mobility and security.

Contents

IPv6 Roundup

It's been a while between editions of *The Standard*. It's not that nothing's been happening in IPv6, it's that too much is happening! We all watched, fascinated, as the last few blocks of IPv4 were allocated in February 2011, months earlier than predicted. The sky stayed firmly in place, but the world finally paid attention (for 5 minutes) to a real Internet milestone.

World IPv6 Day is on Wednesday 8 June 2011 – a day of IPv6 implementation and testing, to motivate service providers and vendors across the globe to get moving with IPv6 – see articles inside.

IPv6Now has launched its first book, a pocket-sized guide to technical IPv6 definitions. Called ***Easy IPv6: The Lookup Book***, see *ipv6now.com.au* for sample pages, contents and order details.

Sometimes we hear mutterings around the place that IPv6 tunnels aren't a 'real' solution. But large-scale, reliable, fast tunnel mechanisms exist for when native IPv6 is not available (or not guaranteed) – which covers a very large chunk of the Australian Internet! See more about **carrier-grade tunnels** and **IPv6 performance** inside.

Launch of *Easy IPv6: The Lookup Book*, October 2010. IPv6Now personnel Mike Biber, Kevin Karp, Tony Hill and Kate Lance, with Latif Ladid, President of the IPv6 Forum.

Please feel free to contact us to discuss anything that interests you in this edition of *The Standard*. Dr Kate Lance – Communications Manager, IPv6Now

World IPv6 Day, 8 June 2011

The 3rd of February 2011 marked the beginning of the end for the IPv4 Internet when the final remaining free blocks of IPv4 were allocated by IANA, one to each Regional Internet Registry. Even just a few months before, most observers had not expected the end to come so soon in 2011. It was something of a wake-up call.

Shockingly, just over two months later, Asia Pacific Network Information Centre (APNIC) announced on 15 April 2011 that it has reached the last significant block of IPv4 addresses in its available pool – making the Asia Pacific region the first to reach the point of being unable to meet regular IPv4 demand. Again, exhaustion arrived months before it had been expected.

Building on rising awareness that a major turning point had been reached, the Internet Society proposed a global trial of IPv6, during which participants would enable IPv6 on their services for 24 hours, called World IPv6 Day.

World IPv6 Day will start at midnight UTC (10.00 am AEST) and run for 24 hours. As of 1 June 2011, around 300 providers and 600 websites have agreed to participate, including major networks and sites such as Google, YouTube, Facebook, Yahoo and Akamai.

In Australia, World IPv6 Day is being coordinated by the Internet Society of Australia. Check out *ipv6.org.au* for more information.

By providing an opportunity for the Internet industry to collaborate to test IPv6 readiness it is hoped to lay the groundwork for large-scale IPv6 adoption. One of the goals of World IPv6 Day is to expose any issues under controlled conditions and address them as soon as possible.

The vast majority of users should be able to access services as usual, but occasionally misconfigured network equipment, particularly in home networks, may impair access to participating websites during the trial. Estimates are that 0.05% of users may experience such problems, but organisations will be working with operating system manufacturers, home router vendors and ISPs to minimise the number of users affected.

Participants will also be providing tools to detect any problems and offer suggested fixes in advance of the trial. For instance, the Regional Internet Registries APNIC and RIPE-NCC have prepared a new tool to measure the speed of connections to websites via IPv4, IPv6, and dual-stack. A simple script shows statistics on the experience Web users would have if connecting via IPv4, IPv6, or dual stack networks, an invaluable tool for improving IPv6 performance. See *labs.apnic.net* for details, sample code, and existing reports.

Encouragingly, preparations for World IPv6 Day seem to have already had an effect on IPv6 performance. Figure 5.1 on the monitoring site *ipv6monitor.comcast.net:8088/monitor/* shows that for the top 100 websites in the world that run dual-stack servers, IPv6 has usually been slower than IPv4 by about 50 ms. But late in May 2011, that difference suddenly fell to zero, making IPv4 and IPv6 equally as fast. A great achievement, but just the beginning.

IPv6Now, Studentnet and World IPv6 Day

IPv6Now is coordinating with Studentnet® (*studentnet.net*) to support educational participation in World IPv6 Day, the first coordinated global industry effort to implement the next-generation protocol. The IPv6 Internet is not backwards-compatible with the existing IPv4 Internet, so techniques like 'dual-stacking' help them inter-operate. World IPv6 Day will allow users and Internet content providers to internationally test dual-stack and IPv6 capabilities.

Studentnet already uses IPv6Now infrastructure to supply email and collaboration services seamlessly to 11,000 Australian school students.

World IPv6 Day provides the opportunity for students to experience IPv6 directly for themselves, as it is the communications foundation that these young people will be using over their lifetimes. The experience will offer staff and students a foretaste of the time when classes collaborate internationally via IPv6 as a matter of routine.

Two groups of students from Waverley College and Wollondilly Anglican College will be given the opportunity to 'break' the Internet, by testing IPv6 and dual-stack access to major content providers as robustly as possible! Will they succeed? Perhaps more importantly, both success and failure will help build a stronger IPv6 Internet.

Studentnet has set up a facility to collate the students' experiences, but in fact anybody in the world is free to use it, to put on record what they experience on World IPv6 Day – have a look at:
studentnet.net/wipv6day.php

While overseas there are many examples of participation in IPv6 Day, IPv6Now and Studentnet are so far the only companies in Australia organising a major demonstration of IPv6 capability. We have optimised IPv6 access for the participating schools with minimal disruption to existing technical and commercial Internet arrangements, even though school networks can be extremely large and complex.

This is possible because the IPv6 infrastructure established by IPv6Now provides immediate access to the IPv6 Internet without changes being required to current hardware, ISP or hosting services.

To make certain that we continue to deliver the highest levels of services, training and consulting, we are using IPv6 Day as an opportunity to test our networks. It will be a welcome challenge for IPv6Now – and one that Studentnet's innovative and resourceful young people will probably take to the extreme!

Australian IPv6 Summit 2011

Learning from Experience

Melbourne 17-19 October 2011
ipv6.org.au/summit

A vast new IPv6 world awaits - but which path will get you there?

Pioneers have forged ahead to find a home on the (address) range. They've hacked through the RFC undergrowth, dual-stacked the hostile servers, and even transitioned the mighty DNS... And they're here to tell you how they did it: the pitfalls, the breakthroughs, & the successes of moving to IPv6.

One-day IPv6Now training workshops will take place before and after the Summit itself, tutored by IPv6 Forum Gold Trainers. Get a practical introduction, and experience hands-on IPv6 for yourself.

Demystifying Tunnels

Q: Aren't tunnels flakey, slow, difficult, subversive and Not A Real Solution?
A: Yes and no! Let's untangle the mythology surrounding tunnels.

What does Dual-Stacking Mean?

Dual-stacking is the process of making every device capable of using both IPv4 and IPv6. The capacity to use IPv6 must be implemented in hardware, software, routing and peering, data services, load management systems, operational support, user access, authentication and accounting, and customer premises equipment. This includes:

- Routing and switching hardware
- Internal and external routing protocols
- Peering and transit services
- Network management subsystems and information bases (MIBs)
- Network measurement and reporting systems
- Firewalls, incident, spam and content filters
- Virtual private networks and other types of tunnels
- Data centre clusters, load balancers
- Cloud servers, database servers, virtualisation servers
- Web servers, email servers, domain name servers, Voice over IP servers
- Authentication, authorisation and accounting (AAA) systems
- Network address translation (NAT) devices, customer modems and routers
- Any other provider specialist services

Dual-stacking requires a complete audit and overhaul of every element in a provider's logical and physical network. It's hardly surprising that only a few have begun the arduous and expensive process.

Is There an Alternative to Dual-Stacking?

Yes. Providing IPv6 via controlled conduits over the existing IPv4 infrastructure. *That is, using tunnels*.

Tunneling is simply the technique of wrapping data with one type of protocol header inside a second header, to be able to traverse a network that uses the second protocol. This 'encapsulation' is used in all sorts of ways on the Internet and other carrier telecommunications networks, not just for IPv6.

In relation to IPv6 there's also a lot of nonsense flying around about tunnels, and much of it is based upon confusion about two major functionalities, very different in practice. They are:

1. DIY tunnels, initiated by end-users ('roll-your-own')
2. Carrier-grade tunnels, operated by service providers

1. DIY Tunnels

These include 6to4, Teredo and manual systems, set up by users behind IPv4-only networks:

6to4 Tunnelling

This depends on a special address format and relaying servers to get v4 traffic to and from v6 destinations. Addresses of 6to4 devices use the special prefix 2002::/16 followed by the IPv4 address. An outgoing IPv6 packet is wrapped inside an IPv4 header with the protocol number 41. An IPv4-to-IPv6 relay strips off the v4 wrapper and sends the v6 packet to the v6 destination. The IPv6 return traffic goes via an IPv6-to-IPv4 relay to get back to the originating host.

Teredo Tunnelling

This also uses a special address format, 2001:0000::/32, followed by the IPv4 address. It is for hosts located behind NAT devices, which usually pass only traffic with TCP and UDP headers, not the protocol 41 of 6to4 tunnels. IPv6 packets are wrapped inside IPv4 and UDP headers, then

Teredo servers provide routing information and Teredo relays pass the encapsulated traffic to the IPv6 destination. To negotiate NAT devices, Teredo depends upon ICMPv6, but ICMP packets are often blocked for security reasons.

Manual Tunnelling

This requires that users are able to configure their tunnel endpoints, so they need control over their routing infrastructure. IPv6 packets are wrapped inside v4 headers with protocol 41, the v4 header is stripped at the tunnel endpoint and the packet sent to the v6 destination.

All three DIY techniques have major limitations: reliance on uncontracted third parties, reliance on often-blocked protocols, and need for high levels of technical expertise and infrastructure control.

A: Yes – DIY tunnels can be flakey, slow, difficult, and (if transiting systems with explicit exclusion policies) they are subversive.

2. Carrier-Grade Tunnels

These are tunnel systems owned and operated by a service provider, spanning the provider's network infrastructure, from the incoming client systems to the provider's border relays. These include 6RD, MPLS-PE and hardware tunnel devices[1]:

6RD Tunnelling

6 Rapid Deployment is similar to 6to4, but with a major improvement – it doesn't use the 6to4 prefix, but the service provider's own prefix, so outgoing and incoming relays are under the service provider's routing control. This removes the unreliability of third-party relays or protocol blocking, but still requires IPv6 implementations for security, monitoring, access and accounting devices.

MPLS-PE Tunnelling

MPLS (Multiprotocol Label Switching) encapsulates packets of any network protocol, creating end-to-end circuits across any type of transport medium. For IPv6 traffic, an IPv4 header is added then the packet is directed to a PE (Provider Edge) router, where the encapsulation is stripped and the original packet sent to the IPv6 destination. This scenario also demands IPv6 for security, monitoring, access and accounting systems.

Hardware Tunnel Server

These are stand-alone devices that not only provide several carrier-grade tunnel protocols, but also *integrate* the functions needed for handling IPv6 in a provider network, such as security, DNS operations, access, authentication and accounting, service monitoring, NAT functions, and IPv4-over-IPv6 services. They operate at near wire speeds.

In all of these cases, service providers do not need to immediately dual-stack their entire infrastructure, they can continue to run IPv4 internally and provide conduits for their IPv6 customers, as they transition carefully and incrementally to IPv6.

A: Carrier-grade tunnels are emphatically not flakey, slow, difficult or subversive.

But are Tunnels a *Real* Solution?

Given that separating out functions onto different, connected devices is a well-proven provider strategy, it's clear that IPv6 carrier-grade tunnels, whether at a protocol level or on dedicated hardware, are as real a solution to the problem of IPv6 transition as full dual-stacking.

In fact, depending on network size and staff expertise, carrier-grade tunnels may well be the most economic, simple, and well-staged solutions to IPv6 adoption.

Over time, as IPv4 gradually becomes the minority protocol, tunnels will also carry legacy IPv4 across the new IPv6 Internet. And beyond transition to IPv6, tunnels still have untapped potential – for instance, IPv6 in IPv6 (6in6) tunnels will be able to provide roaming geographic independence.

1 My thanks to Geoff Huston for his March and April 2011 articles on *potaroo.net*, which were very useful in clarifying the differences between end-user and service provider tunnels.

Carrier-Grade Tunnels in Practice

Large-scale implementations of tunnels are currently fairly rare. In Australia, AARNet and Internode offer tunnel device services along with their native IPv6 product, but their IPv6 addresses are not permanent, reverse DNS is not available, and services are not guaranteed.

IPv6Now provides gogoSERVER hardware tunnels, which include permanent IPv6 addresses, reverse DNS and commercial Service Level Agreements.

The company has made a practice of running its mission-critical internal communications infrastructure over tunnels to test and demonstrate their reliability, without any problems since establishment in 2007.

Studentnet®

On a wider implementation scale, Studentnet has used tunnels for several years to support 11,000 Australian school student email and collaboration services, with outstanding success – it is is the third-largest IPv6-based school network in the world.

Studentnet also offers services that can only be supplied by IPv6. For instance, the massive abundance of addresses means that school alumni can retain them permanently for their own future use.

Another IPv6-only capability is that every device connected to Studentnet has a unique address – there is no IPv4-style network-address translation (NAT) aggregation.

In a school environment, with its strict obligations to protect student well-being, privacy and exposure to age-appropriate material, IPv6 provides excellent functionality for tracking down instances of cyber-bullying, illegal behaviour or security breaches.

Bechtel

Internationally, tunnels were used by the giant engineering company Bechtel, when it decided to implement IPv6 from 2005 on its complex networks in a carefully-designed process[2].

'Frustrating' difficulties with getting native IPv6 from service providers meant that Bechtel had to implement tunnels over virtual private networks for IPv6 wide-area network connections. Bechtel's IPv6 transition has been vast and systematic, and clearly, carrier-grade tunnels were an integral and useful aspect of their strategy.

2 Grossetete, P., Popoviciu, C., and Wettling, F., *Global IPv6 Strategies*, 2008, Cisco Press

Summary

It's clear that once the fear, uncertainty and doubt is dispersed, IPv6 tunnels are practical solutions, with capabilities far beyond anything previously imagined in the IPv4 world.

They have inherent value simply because they work, and work across the Internet without requiring full dual-stacked IPv6 adoption.

Carrier-grade tunnels do not depend on IPv6 implementations by anyone else. Those administered by service providers are:

- *reliable*
- *innovative*
- *available now*
- *cheaper than full dual-stacking*
- *independent of the IPv6 progress of others*

Any provider of network services – whether enterprise, government department or small-medium ISP – owes it to themselves to recognise what carrier-grade tunnels can offer their organisation, especially if they require a fast, easy IPv6 implementation.

IPv6 Performance and Tunnel Brokers

It's an Internet cliche – YMMV – Your Mileage May Vary, and it's true. Every instant on the Internet, traffic conditions are changing, so no study can ever replicate what happened a moment before. Despite this, some excellent long-term work is being done by APNIC and RIPE-NCC on the performance of browsers accessing their sites through IPv4, dual-stack and IPv6 servers (more about this on page 2.)

But just to get a snapshot of the situation, here's a simple, end-user test of IPv6 reachability, done using a carrier-grade IPv6 tunnel broker. On 30 May 2011 I did a series of ping reachability tests to dual-stack servers, including some of the target sites for World IPv6 Day. My computer is located in Melbourne, my IPv4 traffic goes via Internode, and my IPv6 traffic via a gogoSERVER tunnel broker at PPS Internet in Sydney (i.e. not a DIY tunnel like 6to4 or Teredo).

Here are traceroutes showing my path under both protocols – in this case to the site *xkcd.com* in the USA. Both protocols take 16 hops to reach the site, and their performance is nicely comparable, both inside Australian networks and overseas.

IPv4 Route	**IPv6 Route**
traceroute to xkcd.com (72.26.203.99), 30 hops max, 60 byte packets	traceroute to xkcd.com (2001:48c8:1:d:0:23:5482:d026) from 2406:a000::6:d, port 33434, from port 61616, 30 hops max, 60 byte packets
1 192.168.0.1 (192.168.0.1) 3.905 ms 2 lns20.mel4.internode.on.net (150.101.212.44) 21.511 ms 3 te2-2.cor3.mel4.internode.on.net (150.101.208.65) 22.048 ms 4 gi7-1-0.bdr1.mel4.internode.on.net (150.101.210.102) 194.104 ms 5 te6-0-0.bdr1.mel6.internode.on.net (150.101.210.50) 194.411 ms 6 te2-0-0.bdr1.cbr1.internode.on.net (150.101.160.173) 195.207 ms 7 te6-0-0.bdr1.syd4.internode.on.net (150.101.160.170) 196.903 ms 8 te0-0-0.bdr1.syd7.internode.on.net (150.101.197.5) 198.597 ms 9 pos2-0.bdr1.sjc2.internode.on.net (203.16.213.50) 199.302 ms 10 equinix-ix.sjc1.us.voxel.net (206.223.116.4) 200.776 ms 11 ve16.tsr2.iad1.us.voxel.net (208.122.63.233) 276.235 ms 12 0.te1-1.tsr1.phl1.us.voxel.net (173.231.161.37) 280.813 ms 13 0.te6-2.tsr1.ewr1.us.voxel.net (173.231.161.33) 283.752 ms 14 * 15 0.ae57.csr2.lga6.us.voxel.net (208.122.44.210) 288.851 ms 16 72.26.203.99 (72.26.203.99) 310.487 ms	1 2406:a000::6:c (2406:a000::6:c) 30.594 ms 2 2406:a000::5:1 (2406:a000::5:1) 32.391 ms 3 2406:a000:ffff:ffff::1 (2406:a000:ffff:ffff::1) 36.867 ms 4 as4826.ipv6.sydney.pipenetworks.com (2001:7fa:b::8) 33.904 ms 5 ge-0-1-0-136.cor02.syd03.nsw.VOCUS.net.au (2402:7800:0:1::35) 31.667 ms 6 ten-1-3-0.cor01.syd03.nsw.VOCUS.net.au (2402:7800:0:1::59) 32.948 ms 7 2402:7800:0:1::82 (2402:7800:0:1::82) 188.675 ms 8 pos-0-1-1.bdr01.pa001.ca.VOCUS.net.au (2402:7800:100:1::2a) 188.473 ms 9 paix.ipv6.he.net (2001:504:d::10) 195.898 ms 10 paix-ix.pao1.us.voxel.net (2001:504:d::3b) 193.506 ms 11 ve16.tsr2.iad1.us.voxel.net (2001:48c8::801) 265.252 ms 12 0.te1-1.tsr1.phl1.us.voxel.net (2001:48c8::84d) 267.439 ms 13 0.te6-2.tsr1.ewr1.us.voxel.net (2001:48c8::849) 265.469 ms 14 * 15 0.ae2.csr2.lga6.us.voxel.net (2001:48c8::82e) 264.365 ms 16 2001:48c8:1:d:0:23:5482:d026 (2001:48c8:1:d:0:23:5482:d026) 263.734 ms

Interestingly, the IPv6 route was faster than the IPv4 one. How could that be? Everyone 'knows' that IPv6 has performance problems! Clearly, the mileage varied. In any case, it is obvious that the IPv6 traceroute did not suffer because of the use of the integrated carrier-grade tunnel broker.

Here are the results of the ping tests to dual-stack sites around the world (based on 10 pings per site, all with zero packet loss). Many show IPv4 with a tiny advantage over IPv6, but some show the reverse. No statistical significance is claimed – these are just a snapshot of conditions at one moment and prove nothing, except perhaps that IPv6 packets can actually get around in today's IPv4 Internet.

Sites pinged	**IPv4 ping round trip time**				**IPv6 ping6 round trip time**			
	Min	**Average**	**Max**	**MDev**	**Min**	**Average**	**Max**	**MDev**
(Australia) www.apnic.net	43.96	44.50	45.04	0.37	45.10	46.47	48.38	0.98
(Australia) ipv6now.com.au	29.65	30.83	31.59	0.65	32.72	33.57	35.25	0.79
(USA) www.isoc.org	326.42	327.18	327.84	0.55	333.46	336.13	341.68	2.84
(USA) xkcd.com	289.18	289.95	290.91	0.56	263.81	265.72	269.70	1.79
(UK) www.gigatux.com	323.00	323.74	324.74	0.96	326.37	328.50	331.92	1.59
(UK) www.campaya.co.uk	343.15	343.92	344.51	0.58	350.77	353.18	356.72	2.34
(Netherlands) www.ripe.net	328.96	329.55	330.99	0.78	333.86	336.26	340.87	2.42
(Netherlands) python.org	332.53	333.49	334.75	0.88	333.08	335.26	338.02	1.62
(Germany) www.heise.de	336.78	337.34	337.99	0.63	337.53	339.25	340.65	1.24
(Germany) picpaste.de	337.41	338.12	339.05	0.82	335.98	337.73	338.64	0.99
(Slovenia) go6.si	376.49	378.01	379.52	1.26	359.40	360.83	362.43	1.08
(Slovenia) www.arnes.si	359.58	360.80	362.78	1.03	356.49	359.08	361.86	1.63
(Brazil) lacnic.net	375.63	376.45	376.79	0.69	384.62	386.33	389.39	1.78
(Brazil) registro.br	374.74	375.37	376.05	0.51	383.54	385.52	389.67	1.77
(Canada) www.tatacommunications.com	344.21	344.67	345.34	0.57	263.08	267.30	276.26	4.19
Column averages	**301.44**	**302.26**	**303.19**	**0.72**	**295.99**	**298.07**	**301.43**	**1.80**

Again, surprisingly, the IPv6 averages seem slightly better than the IPv4 – but that's mainly because one overseas site has spectacularly fast IPv6 performance.

In general, some IPv6 pings are a few milliseconds slower than IPv4. More importantly, the IPv6 standard deviation is much greater, showing far more variability in speed of response. Although simple ping tests are not directly applicable to the much more difficult task of moving large chunks of data across the Internet, it's heartening, from two points of view, to see how little difference there is.

First, when sites and their networks are carefully dual-stacked their IPv6 response can be excellent, and second, the use of specialist IPv6 hardware (carrier-grade tunnel broker), imposes no detrimental effect on performance.

In fact, types of traffic identified as 'unicast' and 'brokered' seem to have almost identical behaviour when fetching an HTTP object, according to RIPE-NCC statistics[3]. Their final graph shows the 'Average extra delay when using IPv6', broken down into 6to4, Teredo, unicast and brokered. The 6to4 and Teredo traffic have distinctly large delays, but the unicast and brokered are lower, both hovering at around the same average level of 35 milliseconds delay.

This suggests there is not much difference in unicast and brokered IPv6 purely as delivery mechanisms, and the measurement simply reflects the accumulated tiny delays of the many IPv4 paths that IPv6 packets must traverse to and from their destinations.

Over the last decades, Internet software and hardware have been systematically tuned to get the best out of IPv4. It will be wonderful to one day see that kind of effort directed towards optimising IPv6 networks.

Psst! Wanna Buy a Used Address?

A market for buying and selling IPv4 addresses may be emerging. In March the bankrupt company Nortel was offered US$7.5 million by Microsoft for 666,624 IPv4 addresses.

Nortel had realised that a block of legacy addresses it possessed were a potential asset (legacy addresses are those handed out before the Regional Internet Registries (RIRs) were set up). After expressions of interest, Microsoft emerges as the highest bidder. The deal was finalised on 26 April, and worked out to US$11.25 per address.

With IPv4 address exhaustion looming, the RIRs had changed their policies to allow transfer of addresses between local registries within their own regions, but transfer across regions is currently not supported, although at least one broker is offering it as a 'grey-zone' service. A few websites calling themselves IPv4 brokers have popped up in the US[4], but they are fairly light on detail about what they actually do.

TradeIPv4[5] is slightly more informative. It appears to be the only one with any actual activity, but the gap between potential vendors and buyers is enormous: a seller of ARIN addresses wants $200 per address, but a buyer is offering only $3.50. Blocks must be a minimum of /24, i.e. 256 addresses, so the seller wants $51,200 and the buyer is willing to pay $700. Not exactly a harmonious meeting of expectations!

An APNIC seller wants a significantly lower $10 per address which is odd, considering APNIC is the only regional registry currently to have exhausted its IPv4 allocations, so according to market mythology it should have the highest scarcity value.

Stephan Lagerholm[6] points out that before the IANA pool was depleted, the average allocation rate of IPv4 addresses at APNIC went from slightly below 0.4 million per day, to almost 1.2 million per day immediately after IANA exhaustion. Such a leap in demand suggests a land-grab mentality, and probably includes people hoping to sell the addresses for profit. It would be nice if their hopes were dashed.

And as someone from a company that's been using IPv6 painlessly for years, I can only shrug and say "why bother?" It will be much cheaper in the long run to just leave IPv4 behind.

---

3 albatross.ripe.net/v6-clientresolver/site_ncc/
4 e.g. denuo.com, depository.net, addrex.net
5 tradeipv4.com
6 www.ipv4depletion.com

Go-Go-Gadget SERVER!

Professionally managed IPv6 tunnels are a safe, secure, standards-compliant, cost-efficient means of implementing full IPv6 connectivity.

The gold standard in carrier-grade tunnel hardware is the gogoSERVER from *gogo6.com*.

IPv6Now is the sole supplier and support for gogoSERVER in Australia, and has been using the device as an integral part of its own network services for over three years.

gogoSERVER provides carrier-grade functionality, including high capacity throughput, redundancy, high reliability, scalability and in-service upgradeability.

The device is compliant with IETF standards: RFCs 1981, 2136, 2460, 2464, 2473, 2766, 3053, 3177, 3587, 4213, 4291, 4294, 4301, 4302, 4303, 4443, 4861, 4862, 5095, 5572. It is also covered by U.S. Patents No. 7,305,481, 7,321,598 and 7,388,865.

gogoSERVER Capabilities

IPv6 in IPv4 Tunnels

TSP (Tunnel Service Protocol) provides IPv6 connectivity over IPv4 infrastructure without changes to infrastructure or customer devices. Because the TSP standard is available on all IPv6-enabled hosts and routers on the market, gogoSERVER provides IPv6 connectivity to any existing device.

6RD (6 Rapid Deployment) - provides stateless IPv6 in IPv4 tunneling and offers another means of deploying IPv6 in existing IPv4 networks. It provides direct communication between tunnelled hosts, which increases performance without adding cost.

IPv4 in IPv6 Tunnels

DS-lite (Dual-Stack Lite) provides IPv4 connectivity over an IPv6 infrastructure with shared IPv4 addresses. With gogoSERVER's integrated large scale NATs, providers can share increasingly scarce IPv4 addresses among their users, so ISPs can ensure business continuity regardless of IPv4 supply.

DSTM (Dual Stack Transition Mechanism) allows the deployment of IPv4 services to dual-stack hosts behind an IPv6-only network. IPv4 addresses allocated are allocated only for the period of time that dual-stack clients need an IPv4 connection, reducing the consumption rate of IPv4 addresses.

Integrated Networking Services

Reliable NAT traversal – gogoSERVER supports IPv6 in UDP IPv4 tunnels, so it can establish IPv6 connectivity even with IPv4 NAT devices in the path. Works through all NATs, including symmetric and cone devices, as well as nested NATs. The NAT is discovered automatically, requires no user intervention.

DNS registration and reverse delegation – Addresses assigned by the server are automatically registered in a DNS server. Static (permanent) and dynamic address support.

Reliable delegation of IPv6 or IPv4 prefixes – The server delegates prefixes of any length (for example, IPv6 /48 or /64, IPv4 /24 or /26) to devices acting as a router. The gogoCLIENT automatically configures router advertisements or DHCP, providing connectivity to all hosts within a network.

Monitoring – gogoSERVER offers extensive monitoring capabilities through SNMP. Administrators can easily keep track of the connected users, bandwidth usage, number of available tunnels, etc. gogoSERVER supports industry-standard Management Information Bases, as well as Hexago's MIB for specific data.

High availability and performance – redundancy, failure reconnection, solid state memory, tens of thousands concurrent sessions possible, near line speed throughput, highly efficient tunnel bring-up, tear-down.

Security – Access control lists, encrypted user authentication. IPSec for management and static tunnels.

AAA – User authentication, authorization and accounting control. Full support for Radius, even address assignment can be done using Radius.

Redirection – modular, can be deployed over multiple sites for availability and scalability – the client can detect connectivity failures and reconnect to alternateservers.

HA (HomeAccess) Platform – provides seamless connectivity to private home networks.

Want your own gogoSERVER – or maybe just a slice of one?

Talk to IPv6Business Manager Kevin Karp, phone 1800 222 085 or email *services@ipv6now.com.au*

Find out more about IPv6Now Access tunnels, gogoSERVER hardware and more at *ipv6now.com.au/services.php*

Front-End IPv6 Packages

User packages for those who don't want to change their current services

Your IPv4 Internet access and website hosting are currently supplied by ISPs or hosting services. But they don't do IPv6, and you may be losing out by being invisible to the IPv6 Internet.

***Don't* change your ISP. *Don't* change your web hosting provider. *Don't* change your hardware. Just get effortless IPv6 access and visibility for existing IPv4 services.**

Mini6

For small offices and home offices e.g. architects, vets, accountants, B&Bs, ebay shops, sole traders, galleries, tradies...

You get Internet access via IPv6 for a single network with 256 subnets, plus a permanent IPv6 address for one website.

Your Internet presence is now effortlessly futureproofed.

$160 per year (ex-GST)

Midi6

For smaller networks, such as school sites, family retail and wholesale businesses, real-estate agencies, health providers...

You get IPv6 Internet access for up to 3 networks, 256 subnets, and IPv6 addresses for up to 3 websites.

Package includes one two-hour briefing per year for 3 people on IPv6 basics and strategies.

$1,790 per year (ex-GST)

Maxi6

For medium to large business networks with multiple branches, department stores, franchise operations ...

You get IPv6 Internet access for up to 10 networks, 256 subnets, and IPv6 addresses for up to 10 websites.

Package includes one day per year for 3 people of an information or training IPv6 workshop.

$4,120 per year (ex-GST)

Back-End IPv6 Packages

Service Provider packages, for those who supply network services

You are a distributed enterprise, government department or Internet Service Provider of any size, with control over your own routing. Your networks don't currently run IPv6. The cost of resources and training for full dual-stack adoption is not in the budget, but you need to provide both IPv4 and IPv6 services.

Give your clients *immediate IPv6 access and websites* while you phase in IPv6. Packages include Silver or Gold technical training, taught by IPv6 Forum Certified Gold Trainers.

SmallSP

For smaller enterprises, ISPs, and government departments.

You get IPv6 Internet access for up to 10 networks, 65,536 subnets, and IPv6 addresses for up to 10 websites.

PLUS 3 days per year of IPv6 Technical Training for 3 people, IPv6 Forum Silver Certified.

$9,520 per year (ex-GST)

MediumSP

For mid-sized enterprises, ISPs, and government departments.

You get IPv6 Internet access for up to 100 networks, 65,536 subnets, and IPv6 addresses for up to 100 websites.

PLUS 5 days per year of IPv6 Technical Training for 3 people, IPv6 Forum Gold Certified.

$21,520 per year (ex-GST)

LargeSP

For larger enterprises, ISPs, and government departments.

You get your own hardware carrier-grade gogoSERVER, providing IPv6 access for up to 50,000 users.

PLUS 5 days of IPv6 Technical Training for 3 people, IPv6 Forum Gold Certified.

$40,500 (ex-GST)

ipv6now.com.au/packages.php

For all packages (except LargeSP): IPv6Now traffic is protected via IPv6-specific firewalls. IPv6 DNS has both forward and reverse domain name resolution. All IPv6 addresses are static (permanent) at no extra charge. IPv6 traffic transits via Vocus – fast, large-capacity native IPv6, with failover to NTT. Our servers are in Australia, providing low-latency round-trip, local support and local IPv6 peering. Services are monitored for alerts, performance, resources and 24x7 availability. Commercial-grade Service Level Agreements apply.

IPv6 Advice, Workshops and Training

IPv6Now is the *only dedicated IPv6 company* in Australia

Our education services are available for all levels of business, government and service provider knowledge. IPv6Now is an Approved Supplier of IPv6 training to the Australian Government and major telecommunications companies.

Our trainers and courses are Certified Gold by the IPv6 Forum

"The IPv6 Forum has selected and certified IPv6Now for its pioneering work ... and approved Michael and Karl as among the very first IPv6 trainers in the world to qualify for the new certification", wrote Latif Ladid, President of the IPv6 Forum.

We offer real-world educational advantages

Our courses are based upon decades of commercial and technical Internet experience, understanding of strategic directions in IPv6, development of our own independent IPv6 network, and real world understanding of IPv6 hardware and applications.

Consulting and Training Programs

- **Strategic Analysis Reports** – tailored for specific organisations and circumstances
- **Business Briefings (2 Hours)** – global, political and commercial impact of IPv6
- **Information Workshops (1 Day)** – interacting with IPv6-using suppliers & customers
- **Technical IPv6 (3 Days)** – details of IPv6 formats, techniques, hands-on training
- **Advanced IPv6 (5 Days)** – Technical Course, plus two days for provider specialists

Courses are held in capital cities once per month. They may also be arranged at non-scheduled dates or locations if there is sufficient demand. Discounts for multiple attendees are available. Training is held either on-site or at specialised service premises. Bookings are *essential*, courses must be paid for in advance, and we reserve the right to cancel any course.

Want to know more?

ipv6now.com.au/consulting.php ***ipv6now.com.au/training.php***

Talk to IPv6Now Professional Services Manager, Michael Biber
Mobile 0412 058 808 or email training@ipv6now.com.au

Writer: Dr Kate Lance (editor@ipv6now.com.au)
See earlier editions of *The Standard* at *ipv6now.com.au/standard.php*

For advertising space and rates in *The Standard*, please contact Kevin Karp, Business Manager
Phone 1800 222 085 or email *services@ipv6now.com.au*

ipv6now.com.au ipv6now.co.nz