Aussie Broadband ipv6 deployment journey

How Aussie Broadband (ABB) has deployed ipv6 from the initial stages to now.

By John Alexander

Network Engineer at ABB

Aussie Broadband

Bloody Good Broadband

Pre-initial alpha stage (< 2018)

- ABB (AS4764) was allocated a /32 back in the mid 2000's
- For some selected customers, we implemented a single router in Melbourne that did ipv6 tunneling over ipv4, allowing ipv6 to work. This was available for many years and worked well if you were in Victoria... Customers got a /64 range and were able to use ipv6 as normally as possible. This was then announced on a public forum on April 23 2018 so that more customers could use it.

Beta program (2018-2021)

- We wanted to change to dhcp instead of tunnels. But we had an issue with our BNGs remarking ipv6 dhcp cos values to 6, which were subsequently dropped by NBN (Australia's National broadband network) thereby stopping it from working. A bug was raised and fixed about a year later.
- https://www.ausnog.net/sites/default/files/ausnog-2018/presentations/2.6_Phil_Britt_AusNOG2018.pdf

20th Nov 2018

- On November 20 2018 we released the beta program.
- The initial rollout was everyone got a /128 wan ip from a /64 range for their dhcp connection which was allocated out of a /48 tied to each BNG.
- They also got a pd (prefix delegation) of a /56 also allocated out of a /40 range from each BNG.
- In order to divide up our allocated /32 range from apnic into chunks for use by customers, we first allocated /36s per state:

```
RP/0/RSP0/CPU0:bng2.vdc03.mel#sh subscriber session filter interface PW-
Ether2005.2005.ip22818

Sun Feb 13 14:48:00.313 AEDT

Codes: IN - Initialize, CN - Connecting, CD - Connected, AC - Activated,

       ID - Idle, DN - Disconnecting, ED - End


Type         Interface                 State     Subscriber IP Addr / Prefix
                                                 LNS Address

-------------------------------------------------------------------------------
IP:DHCP      PE2005.2005.ip22818       AC        100.68.100.100
                                                 2403:5802:ffff:d2:1f2e:2e4:a7c7:e62e
                                                 2403:5802:8888::/56
```

- 2403:5800::/36 and 2403:5800:1000::/36 for network and server use
- 2403:5800:2000::/36 reserved
- 2403:5800:3000::/36 Victoria East
- 2403:5800:4000::/36 Victoria West
- 2403:5800:5000::/36 South Australia
- 2403:5800:6000::/36 South Australia
- 2403:5800:7000::/36 New South Wales
- 2403:5800:8000::/36 New South Wales
- 2403:5800:9000::/36 Queensland
- 2403:5800:a000::/36 VPN/Tunnels (ie the alpha stage)
- 2403:5800:b000::/36 Business
- 2403:5800:c000::/36 Western Australia
- 2403:5800:d000::/36 Western Australia
- 2403:5800:e000::/36 Reserved
- 2403:5800:f000::/36 Testing

So within each /36 we allocated a /48 and /40 for each bng:

- 2403:5800:3000::/48 (bng1.portmel residential ia_na)
- 2403:5800:3001::/48 (bng2.portmel residential ia_na)
- 2403:5800:3002::/48 (bng3.portmel residential ia_na)
- 2403:5800:3003::/48 (bng4.portmel residential ia_na)

- 2403:5800:3100::/40 (bng1.portmel residential ia_pd)
- 2403:5800:3200::/40 (bng2.portmel residential ia_pd)
- 2403:5800:3300::/40 (bng3.portmel residential ia_pd)
- 2403:5800:3400::/40 (bng4.portmel residential ia_pd)

- Allowed for 64,000 /56s per bng (each bng could support a max of 32,000 subscribers, so no issue there.
- We could also support 15 /40's (1 was for the /48s) per /36 which was also sufficient to make things work. Where things started to break down was we were getting requests for /48s from businesses and while it could have worked, it was proving more difficult to manage.
- With no basic difference between /48s and /56s we decided to standardise on /48s to keep everyone the same, removing the need to have different ranges for different customers.

Advertising to the internet

- How did we advertise ipv6 out to the world?
- Looking at bgp.he.net for 4764, you can still see the /36's being advertised, as well as the parent /32. So out of our /32 allocation we de-aggregated down to a /36 thereby increasing everyones route table size by another 6 entries...

https://bgp.he.net/AS4764#_prefixes6

HURRICANE ELECTRIC
INTERNET SERVICES

Search

AS4764 Aussie Broadband

Quick Links

- BGP Toolkit Home
- BGP Prefix Report
- BGP Peer Report
- Exchange Report
- Bogon Routes
- World Report
- Multi Origin Routes
- DNS Report
- Top Host Report
- Internet Statistics
- Looking Glass
- Network Tools App
- Free IPv6 Tunnel
- IPv6 Certification
- IPv6 Progress
- Going Native
- Contact Us

AS Info | Graph v4 | Graph v6 | Prefixes v4 | Prefixes v6 | Peers v4 | Peers v6 | Whois | IRR | IX

Prefix	Description
2403:5800::/28	Wideband Networks Pty Ltd
2403:5800::/32	Wideband Networks Pty Ltd
2403:5800:1::/48	Aussie Broadband
2403:5800:3::/48	Aussie Broadband
2403:5800:5::/48	Aussie Broadband
2403:5800:3000::/36	Wideband Networks Pty Ltd
2403:5800:4000::/36	Wideband Networks Pty Ltd
2403:5800:5000::/36	Wideband Networks Pty Ltd
2403:5800:7000::/36	Wideband Networks Pty Ltd
2403:5800:9000::/36	Wideband Networks Pty Ltd
2403:5800:c000::/36	Wideband Networks Pty Ltd
2406:c500:ffd2::/48	INFORMATION TECHNOLOGY PROFESSIONALS ASSOCIATION LTD

- We advertised the /36s by creating a null route and injecting into bgp to be advertised out, but the /40s per bng were internal only and not advertised to the wider world.
- This staying in place until we removed the beta and became mainstream in October 2021.

```
router static
 address-family ipv6 unicast
  2403:5800:c000::/36 null 0 250
!
router bgp 4764
 address-family ipv6 unicast
  aggregate-address 2403:5800:c000::/36 summary-only route-policy
SET-OUTBOUND-COMMUNITIES
```

- Why did it take so long?

- Well... Rome wasn't built in a day. And yes ipv6 has been around for a long time now. But..

- The initial issue discussed above with the bng software incorrectly marking ipv6 dhcp dscp bits kind of made it really difficult to rollout on a mass scale until that was fixed.

CPEs (modems/routers)

- Mark Smith (who some of you may know) has spent a huge amount of time working with our CPE vendors to make them work. Yes they should be working out of the box, but they weren't, even in 2021. The issues Mark discovered and got them (vendors) to fix were:

- ipv6 not enabled by default. This is a big one, ie out of the box, they just weren't setup to do ipv6, which makes plug and play impossible.
- ipv6/dhcpv6 stopping on wan interfaces after a few days
- ipv6 protocol parameters being set lower than RFC defaults
- ipv6 only working on one wan interface rather than multiple, eg works on the vdsl but not ethernet wan

CPEs

- Marks observations was CPEs are actually worse now than what they were in 2011! Refer to Mark's Ausnog presentation

  https://www.ausnog.net/sites/default/files/ausnog-05/presentations/ausnog-05-d02p02-mark-smith.pdf

- That being said, Mark has been able to get our 2 main suppliers up to speed and ipv6 works again out of the box!

Network changes

- Much discussion ensued as to how to setup our network to support all of our customers getting ipv6 addresses.
- We had 2 options, continue with /56s, or change to /48s per customer. Why the possible change to /48s?
- Discussed earlier, to keep everything between types of customers consistent. The actual cost of a /48 or a /56 is irrelevent in the scheme of things as ipv4 is the true cost that we pay apnic fees for.

More space...

- So if we use /48s how much ipv6 space do we need?
- Assuming 32k /48s per BNG, that equals half a /32 per BNG. (a /33)
- We have 25 BNGs in use at the moment, so 25 x /33s = 13 /32s = /28 with some spare.
- So we can have 31 BNGs worth of /33s (1 was the original /32 we had at the start)
- Getting the /28 though from Apnic proved to be easier said than done, although we did get there in the end.

- Using apnic's own blog https://blog.apnic.net/2017/07/10/isps-simplifying-customer-ipv6-addressing-part-2/ we requested the /28 as it would provide all our customers with a simpler address scheme of a /48 each regardless of whether they were a residential or business customer.

- This very handy chart from Ripe:

  https://www.ripe.net/about-us/press-centre/ipv6-chart_2015.pdf

Prefix	/48s	/56s	/64s	Bits
/24	16M	4G	1T	104
/25	8M	2G	512G	103
/26	4M	1G	256G	102
/27	2M	512M	128G	101
/28	1M	256M	64G	100
/29	512K	128M	32G	99
/30	256K	64M	16G	98
/31	128K	32M	8G	97
/32	64K	16M	4G	96
/33	32K	8M	2G	95
/34	16K	4M	1G	94
/35	8K	2M	512M	93
/36	4K	1M	256M	92

- shows how many /48s per prefix we can use.
- We have currently about 467,000 customers meaning we're at around 50% utilisation of our allocated /28. However as we're using a /33 per BNG, we have 6 BNGs left before we run out. It probably means I need to re-apply to expand our allocation to a /27 which would double the space, and give us much more breathing room.
- 2403:5810::/28 is unallocated at the moment so we have a chance...
- Postscript, we now have 2403:5800::/27

- Looking back at https://bgp.he.net/AS4764#_prefixes6 you can see we have a single /27 advertisement for 2403:5800::/27

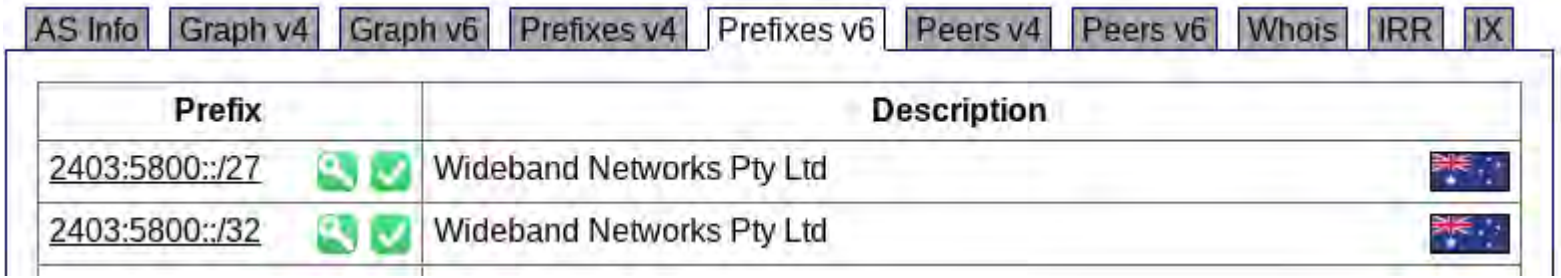

- Ultimately we'll be dropping it down to the /28~~7~~ only :-)

- None of the BNG /33s are advertised to the outside world. We have our core routers doing a null route for the /27 and advertising that out, while the BNGs use communities to prevent our border routers from advertising those prefixes out.
- Each BNG aggregates the individual /48s into a single /33 to advertise to the rest of our network, thereby minimising the size of our internal route table.
- Traffic bound for us will hit our border routers, and still know how to get to the BNG which then sends to the end user.

2403:5800::/32	Legacy		
2403:5801::/33	Vic/Tas	2403:5809:8000::/33	NSW/ACT
2403:5801:8000::/33	Vic/Tas	2403:580A::/33	QLD
2403:5802::/33	Vic/Tas	2403:580A:8000::/33	QLD
2403:5802:8000/33	Vic/Tas	2403:580B::/33	QLD
2403:5803::/33	Vic/Tas	2403:580B:8000::/33	QLD
2403:5803:8000::/33	Vic/Tas	2403:580C::/33	QLD
2403:5804::/33	Vic/Tas	2403:580C:8000::/33	
2403:5804:8000::/33	Vic/Tas	2403:580D::/33	SA/NT
2403:5805::/33	Vic/Tas	2403:580D:8000::/33	SA/NT
2403:5805:8000::/33		2403:580E::/33	WA
2403:5806::/33	NSW/ACT	2403:580E:8000::/33	WA
2403:5806:8000::/33	NSW/ACT	2403:580F::/33	WA
2403:5807::/33	NSW/ACT	2403:580F:8000::/33	
2403:5807:8000::/33	NSW/ACT		
2403:5808::/33	NSW/ACT		
2403:5808:8000::/33	NSW/ACT		
2403:5809::/33	NSW/ACT		

- The only issue we have with this solution is if a customer gets moved to another BNG, they lose their /48 prefix and have to get a new one. We don't have a proper fix for this yet.
- We are investigating aggregating at our state routers, eg VIC/Tas from the previous slide and allow customers to move between bngs within the state. This would end up with 288,000 ipv6 routes in Vic, but aggregated to the 9 /33s.

Conclusion

- So what do we see in our network now?
- All new customers are assigned an ipv6 range by default.
- We have retro-activated approximately 60,000 customers with ipv6 that we were seeing actively requesting ipv6 via dhcp and were running CPEs we could manage.
- The next step is to allocate the rest of our customers with ipv6 so everyone has an allocation and getting them using it.
- We are implementing more monitoring to be able to see actual ipv6 traffic use.
- From stats from Google, we are seeing about 5% ipv6, and from Netflix about 0.8% usage. This needs to go up! But TVs mostly don't have ipv6 enabled :-(

aussiebroadband.com.au

Aussie Broadband

Bloody Good Broadband