

# IPv6 Experience at Internode

Development and Deployment

---

*22 Feb 2011, APRICOT/APAN Hong Kong*

# Talk Outline

---

- ❖ Who is Internode?
- ❖ The market in which we operate
- ❖ Our network
- ❖ Implementation
- ❖ Broadband trial
- ❖ Salient points

# Who are we?

---

- ❖ Australian broadband ISP
- ❖ Privately owned
- ❖ Good mix of business customers and about 200,000 residential broadband customers
- ❖ Marketplace reputation as an innovator and thought leader.



# Our Marketplace

---


- ❖ ADSL2+ with PPPoE
- ❖ BYO CPE
- ❖ L2TP wholesale
- ❖ Usage accounting with strong accuracy requirements
- ❖ Significant customer support issues associated with departure from legacy.



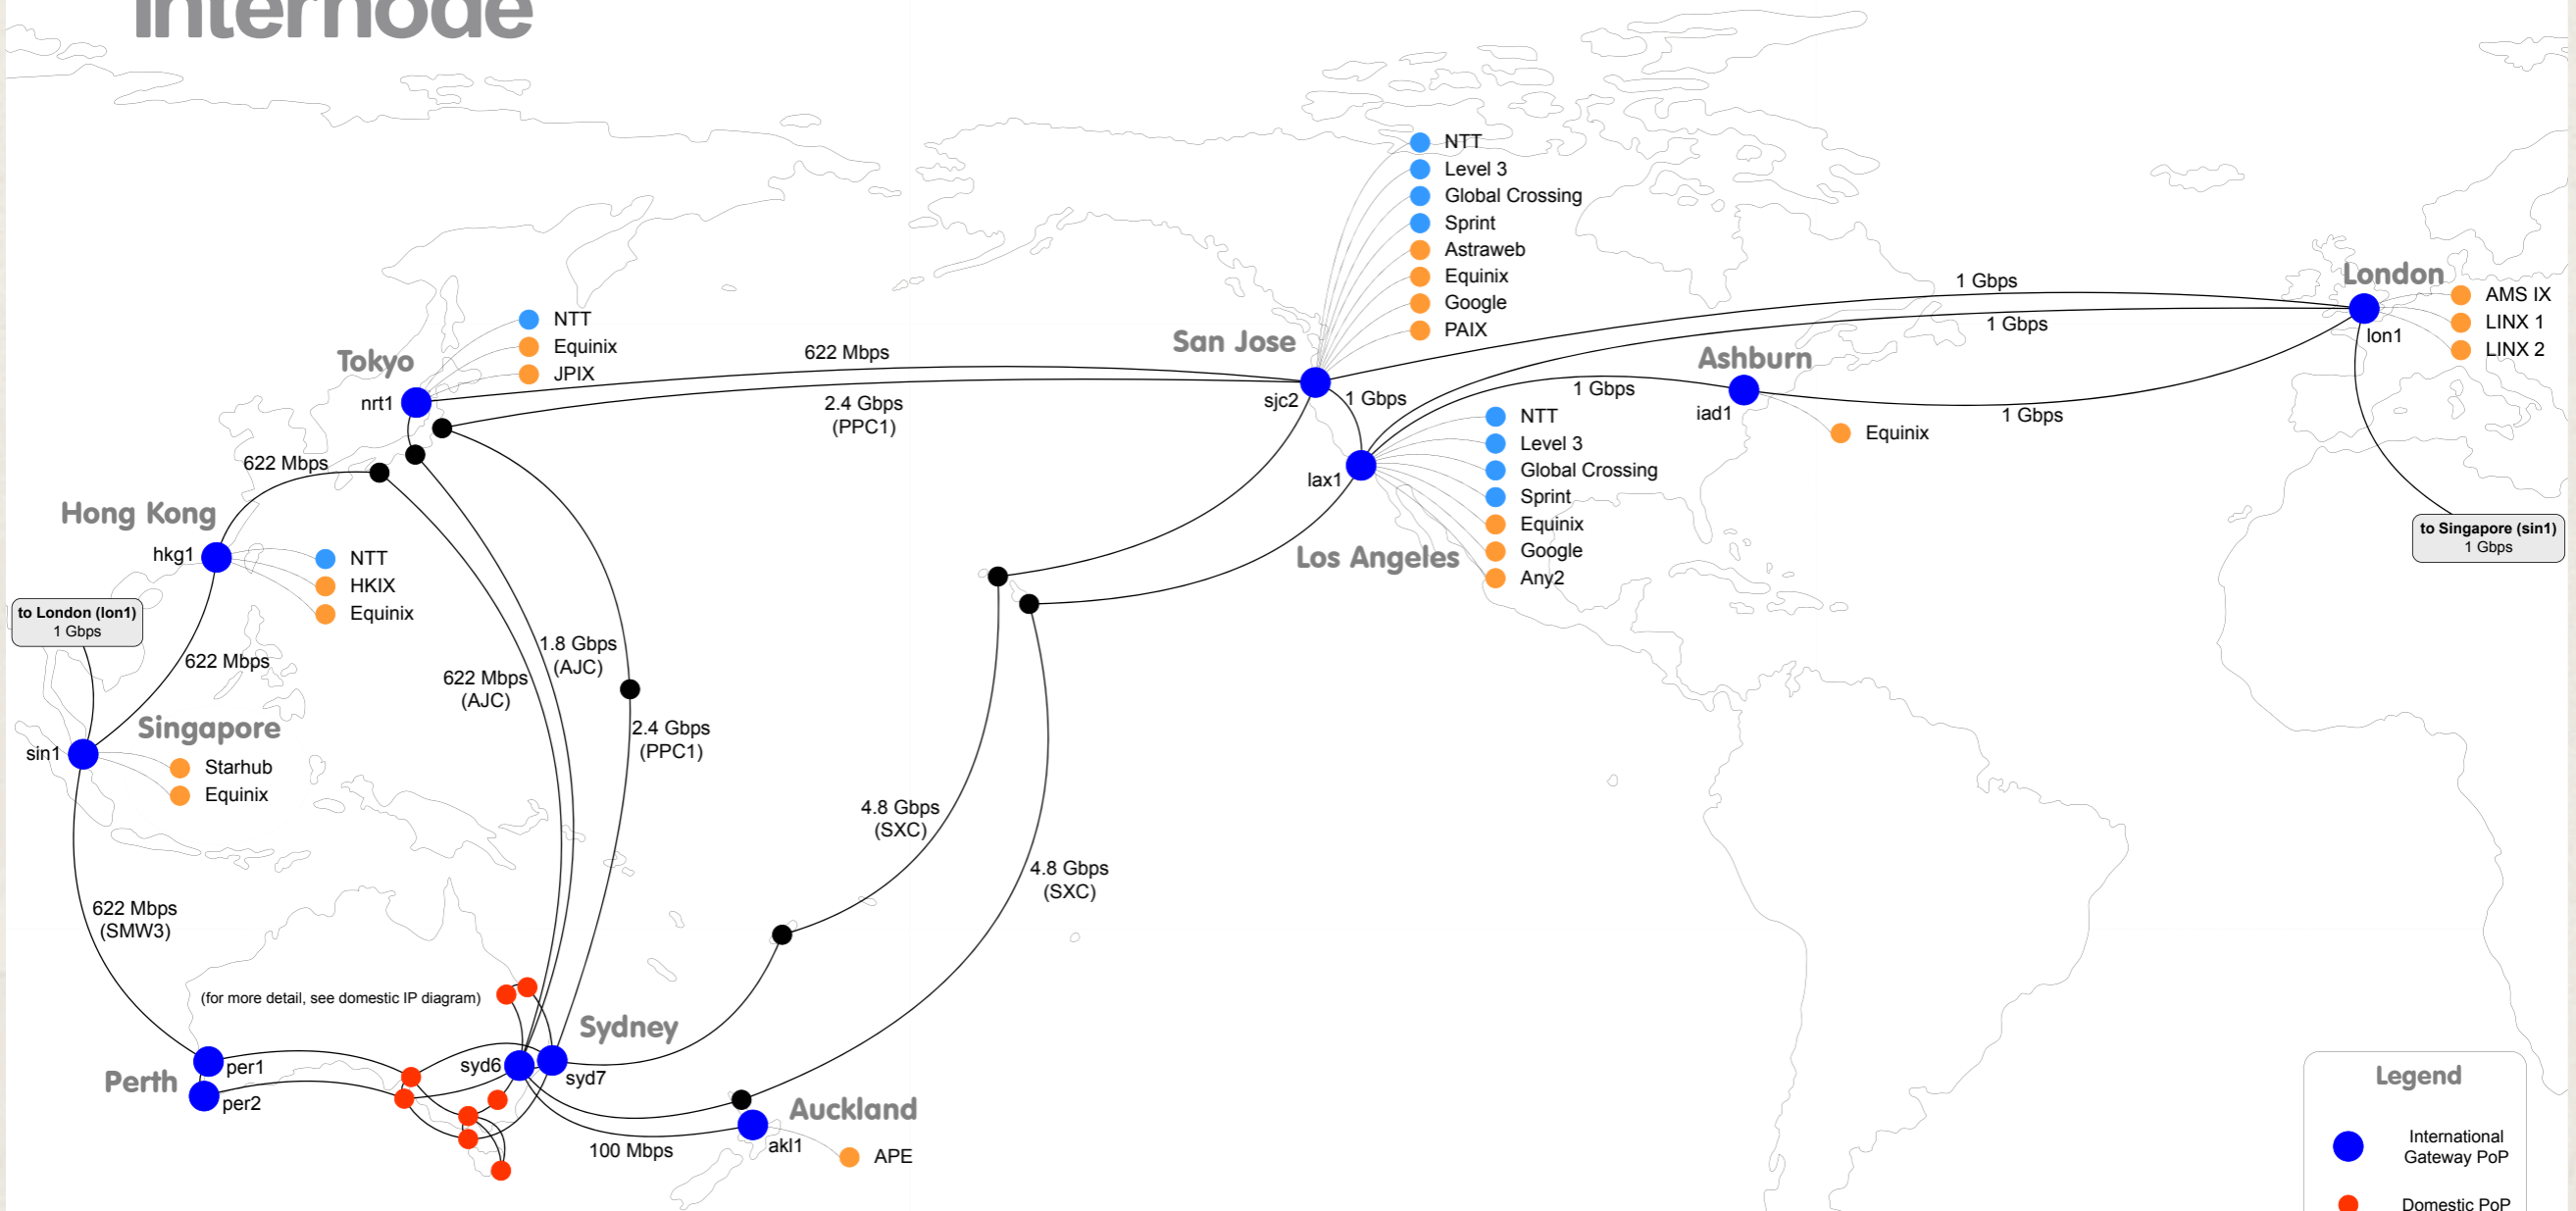
# Our Network

---

- \* Globe-spanning
- \* Australian portion is about 60mS RTT Perth-Brisbane
- \* At least two major POPs per state capital city
- \* At least one BRAS/LNS per major POP.



## International IP Network June 2010



# Internode's network

This diagram is the property of Internode Pty Ltd - AS4739 - Indicative representation only

SXC : Southern Cross Cable  
 AJC : Australia-Japan Cable  
 PPC1 : Pipe Pacific 1 Cable  
 SMW3 : SEA-ME-WE 3 Cable

**Legend**

- International Gateway PoP
- Domestic PoP
- Cable Station
- Transit
- Peer

— Fibre optic link —

# Our IPv6 Transition

---

- ❖ Internode's owner and founder has always wanted to do it.
- ❖ Geoff Huston's presentation about IPv4 utilisation rates at RIPE in July 2007 was the catalyst to get started.
- ❖ Geoff predicted IPv4 exhaustion in 2009, amusing in retrospect.

# Our IPv6 Transition

---

- ❖ “Outside-in”
- ❖ Started at our peering edge router in San Jose, CA.
- ❖ Progressively brought it in to the rest of the network over the next few weeks.
- ❖ Killed bugs as we progressed (e.g.: C7200 accidentally exporting IPv6 prefixes to IPv4 RIB)
- ❖ *Automated BGP config helped a lot!* Change one script to fix the whole network



# Our IPv6 Transition

---

- ❖ Core network done, first access customer in January 2008.
- ❖ IPv6 dual-stack for network design/operations staff at home on ADSL provided via multihop L2TP to a private LNS
- ❖ Early access for technical staff enables skills development, training.

# Digression: Is a /32 big enough?

---

- ❖ Some folks advocate /48 assignments for everyone.
- ❖ There are 65536 /48's in a /32 -- a bit small for an ISP with 200,000-odd customers.
- ❖ Perhaps not every customer is going to get a /48.
- ❖ Dynamic IP isn't going away.
- ❖ Hope /32 is big enough, 'cos we're kinda committed to it now!

# Broadband Customer Edge

---

- ❖ Australia uses access methods not widely considered by IPv6 boosters.
- ❖ PPPoE IP6CP, SLAAC with ND / RA to allocate addresses for broadband CPE.
- ❖ /64 for the link drawn from a dynamic pool.
- ❖ DHCPv6-PD for prefix assignment
  - ❖ Static or dynamic depending on service type, config by Radius.
  - ❖ Mostly dynamic.

# Vendor Support

---

- ❖ As it turns out, that deployment model is rather difficult.
- ❖ Cisco 10Ks are IPv6-challenged.
  - ❖ Switched to ASR 1000 series, IOS-XE 2.6.x
  - ❖ IOS-XE 2.6.x has had its own bugs, which we've been working through with Cisco.
- ❖ Customer-facing opt-in trial commenced November 2009 on Cisco 7200-series LNSs, 12.2(33)SRDx.

# IPv6 Broadband Trial

---

- ❖ 7200-series platform
- ❖ Multihop L2TP
  - ❖ PPP login as *foo@internode.on.net* to get IPv4, *foo@ipv6.internode.on.net* to get dual-stack.
- ❖ Several purposes:
  - ❖ CPE vendors: “Test against this”, solves chicken and egg problem.
  - ❖ Geek users: “Here’s your playpen.”
  - ❖ Internode: Operational experience, debugging, testbed, etc.

# Bugs and Misfeatures

---

- ❖ IPv6 accounting: Not so good (e.g., no accounting if v4 or v6 not negotiated)
- ❖ DHCPv6-PD DoS: Radius query every time PD request received, no caching.
- ❖ Cisco DHCPv6-PD server occasionally “forgetting” delegations.
- ❖ Bug-roulette: Find IOS that works well with both v4 *and* v6...

# Current status

---

- ❖ Intending to migrate the trial platform from 7200s to ASRs running 2.6.2-ES this week.
- ❖ Will leave it that way for two weeks.
- ❖ Move to our production BRAS/LNS systems (stop L2TP multihopping) early March 2011.

# Still to do

---

- \* Our IPAM system is, in the Aussie vernacular, “agricultural.”
  - \* Needed a rewrite anyway, now the rewrite will include IPv6.
- \* No ISG support for IPv6
- \* DNS: A great big barrel of bad. Thankfully another team’s problem :)
- \* *vpn6* address family in our MPLS layer-3 VPN product
  - \* All configs automated out of a provisioning system, this shouldn’t be too hard.



# Lessons and recommendations

---

- ❖ **Automate everything.**
  - ❖ But don't use IPv6 as the reason to automate, else you won't be able to do it until the automation is finished
- ❖ **Be practical:** Religious wars about implementation details are boring and pointless at this stage.
- ❖ To vendors: *No more excuses. Stop stalling and just do it, okay?!*
- ❖ **Shipped products and running code;** everything else is detail.
- ❖ **Be “incremental.”** Doesn't have to be all-or-nothing.



*Any questions?*

Mark Newton, [newton@internode.com.au](mailto:newton@internode.com.au), @NewtonMark

---

*22 February 2011*