The Next Three Years

(IPv4 runout and the motivation for IPv6)

MENOG 5

Beirut

28th October 2009

Early Internet History

- Late 1980s
  - Exponential growth of the Internet
- Late 1990: CLNS proposed as IP replacement
- 1991-1992
  - Running out of “class-B” network addresses
  - Rapid growth of the “default-free” routing table
  - Imminent exhaustion of 32-bit address space
- Two efforts – short-term versus long-term
  - More at “The Long and Windy ROAD”
    http://rms46.vlsm.org/1/42.html

Early Internet History

- CIDR and Supernetting proposed in 1992-3
  - Deployment started in 1994
- IETF “ipng” solicitation – RFC1550, Dec 1993
- Direction and technical criteria for ipng choice – RFC1719 and RFC1726, Dec 1994
- Proliferation of proposals:
  - TUBA – RFC1347, June 1992
  - PIP – RFC1621, RFC1622, May 1994
  - CATNIP – RFC1707, October 1994
  - SIP – RFC1710, October 1994
  - NIMROD – RFC1753, December 1994
  - ENCAPS – RFC1955, June 1996

Early Internet History → 1996

- Other activities included:
  - Development of NAT, PPP, DHCP,…
  - Some IPv4 address reclamation
  - The RIR system was introduced
- → Brakes were put on IPv4 address consumption
- IPv4 32 bit address = 4 billion hosts
  - HD Ratio (RFC3194) realistically limits IPv4 to 250 million hosts

Recent Internet History
The “boom” years → 2001

- IPv6 Development in full swing
  - Rapid IPv4 consumption
  - IPv6 specifications sorted out
  - (Many) Transition mechanisms developed
- 6bone
  - Experimental IPv6 backbone sitting on top of Internet
  - Participants from over 100 countries
- Early adopters
  - Japan, Germany, France, UK,…

Recent Internet History
The “bust” years: 2001 → 2003

- The DotCom “crash”
  - i.e. Internet became mainstream
- IPv4:
  - Consumption slowed
  - Address space pressure “reduced”
- IPv6 Indifference
  - Early adopters surging onwards
  - Sceptics more sceptical
  - Yet more transition mechanisms developed

2004 → Today

- Resurgence in demand for IPv4 address space
  - 11.8% address space still unallocated (10/2009)
  - Exhaustion predictions range from wild to conservative
  - …but late 2011 seems realistic at current rates
  - …but what about the market for address space?
- Market for IPv4 addresses:
  - Creates barrier to entry
  - Condemns the less affluent to multiple NATs
- IPv6 offers vast address space
  - **The only compelling reason for IPv6**

Current Situation

- General perception is that “IPv6 has not yet taken hold”
  - IPv4 Address run-out is not “headline news” yet
    - More discussions plus run-out plans being proposed
  - Private sector requires a business case to “migrate”
    - No easy Return on Investment (RoI) computation
- But reality is very different from perception!
  - Something needs to be done to sustain the Internet growth
  - IPv6 or NAT or both or something else?

Is there a need for a larger address space?

- Internet population
  - ~600 million users in Q4 CY2002
  - ~945M by end CY 2004 – only 10-15%
  - Future Worldwide population? (~9B in 2050)
- US uses 88 /8s – this is 4.8 IPv4 addresses per person
  - Repeat this the world over…
  - 6 billion population could require 29 billion IPv4 addresses
  - (7 times larger than the IPv4 address pool)
- Emerging Internet economies need address space:
  - China uses more than 210 million IPv4 addresses today (12.5 /8s)

Is there a need for a larger address space?

- RFC 1918 is not sufficient for large environments
  - Cable Operators (e.g. Comcast – NANOG37 presentation)
  - Mobile providers (fixed/mobile convergence)
  - Large enterprises
- The Policy Development process of the RIRs turned down a request to increase private address space
  - RIR community guideline is to use global addresses instead
  - This leads to an accelerated depletion of the global address space
- Some want 240/4 as new private address space
  - But how to back fit onto all TCP/IP stacks released since 1995?

Status in Internet Operational Community

- Service Providers get an IPv6 prefix from their regional Internet registries
  - Very straight forward process when compared with IPv4
- Much discussion amongst operators about transition:
  - NOG experiments of 2008 – http://www.civil-tongue.net/6and4/
  - What is really still missing from IPv6 – http://www.nanog.org/mtg-0710/presentations/Bush-v6-op-reality.pdf
  - Many presentations on IPv6 deployment experiences

Service Provider Status

- Many transit ISPs have “quietly” made their backbones IPv6 capable as part of infrastructure upgrades
  - Native is common (dual stack)
  - Providers using MPLS use 6PE
  - Tunnels still used (unfortunately)
- Examples:
  - NTT has been long time IPv6 capable
  - OpenTransit/FT, TATA International, Telecom Italia, GlobalCrossing, Telefonica, C&W (EU),…
- OCCAID
  - IPv6-only transit ISP effort (linking Asia, N-America, EU)

OS, Services, Applications, Content

- Operating Systems
  - MacOS X, Linux, BSD Family, many SYS V
  - Windows: XP SP2 (hidden away), Vista, 7
  - All use IPv6 first if available
- Applications
  - Browsers, E-mail clients, IM, bittorrent,…
- Services
  - DNS, Apache WebServer, E-mail gateways,…
- Content Availability
  - Needs to be on IPv4 ***and*** on IPv6

Why are we still waiting…?

- That killer application?
  - Internet Gaming or Peer to Peer applications?
  - Windows Vista or 7 (?)
- Our competitors?
  - Any network deployed in last 3 years will be IPv6 capable
  - Even if not enabled!
- The end-user should not have to choose protocols
  - Remember “Turbo” button on early IBM PC clones?
- The “Chattering Classes”
  - People looking for problems, not solutions

The On-going Debate (1)

- IPv6 Multihoming
  - Same toolset as IPv4 — long term non-scalable
  - ‘Ultimate Multihoming Solution’ no nearer discovery
    - LISP is making interesting progress though
- Early rigid IPv6 address allocation model
  - “One size fits all” barrier to deployment:
    - Only ISPs “should” get IPv6 space from RIRs
    - Enterprises “should” get IPv6 space from ISPs only
  - Routing table entries matter, not the nature of business
    - What is an ISP?

The On-going Debate (2)

- Not every IPv4 device is IPv6 capable
  - Do we really need to replicate all IPv4 capability in IPv6 prior to considering deployment?
- “We have enough IPv4”
  - Those with plenty denying those with little/nothing
- Migration versus Co-existence
  - Realistically IPv6 and IPv4 will co-exist for many years
  - Dual-stack operating systems in network equipment makes this trivial

Why not use Network Address Translation?

- Private address space and Network address translation (NAT) could be used instead of IPv6
- But NAT has many serious issues:
  - Breaks the end-to-end model of IP
  - Breaks end-to-end network security
  - Serious consequences for Lawful Intercept
  - Non-NAT friendly applications means NAT has to be upgraded
  - Some applications don’t work through NATs
  - Layered NAT devices
  - Mandates that the network keeps the state of the connections
  - How to scale NAT performance for large networks??
  - Makes fast rerouting and multihoming difficult
  - How to offer content from behind a NAT?

Is IPv4 really running out?

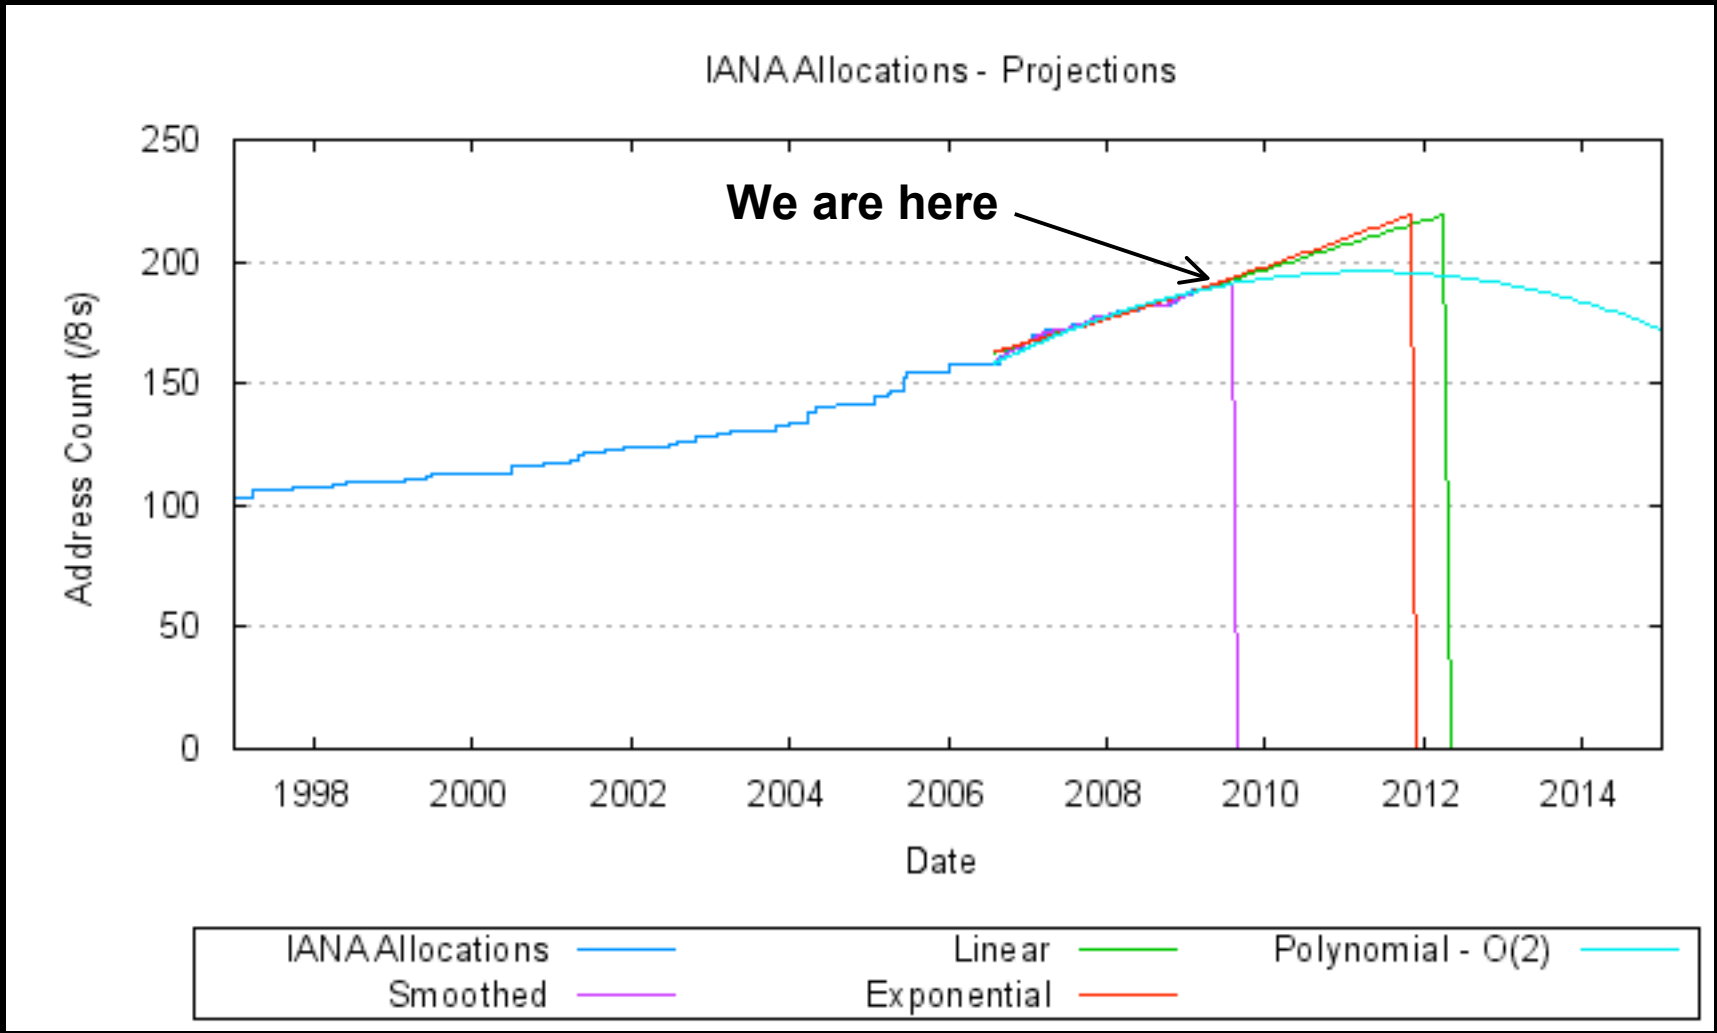

Is IPv4 really running out?

- Yes
  - IANA IPv4 free pool runs out in June 2011
  - RIR IPv4 free pool runs out approx one year later
  - http://www.potaroo.net/tools/ipv4/
- Small industry producing gadgets and widgets predicting IPv4 run-out
  - http://inetcore.com/project/ipv4ec/index_en.html
  - http://ipv6.he.net/statistics/

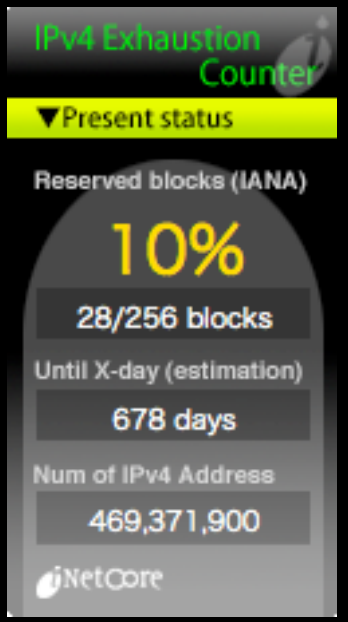

IPv4 run-out

- RIR Policy Development process in each RIR region is now handling many proposals relating to IPv4 run-out
  - The Last /8
    - All RIRs will receive one /8 from the IANA free pool
  - IPv4 address transfer
    - Permits LIRs to transfer address space to each other rather than returning to their RIR
  - Soft landing
    - Reduce the allocation sizes for an LIR as IPv4 pool is depleted
  - IPv4 distribution for IPv6 transition
    - Reserving a range of IPv4 address to assist with IPv6 transition (for Large Scale NATs etc)

Issues Today

- Minimal content is available on IPv6
  - Notwithstanding ipv6.google.com
- Giving IPv6 to customers might confuse
  - Browsers,e-mail clients, etc are smart
  - But increased tech support if IPv6 version of content is ‘down’, but IPv4 version works
- Need to “prolong” IPv4 so there is time for all content to be available on IPv6

Strategies available

- Do nothing
  - Wait and see what competitors do
  - Business not growing, so don’t care
- Extend life of IPv4
  - Push customers to NAT
  - Buy IPv4 address space on the marketplace
- Deploy IPv6
  - Dual stack infrastructure
  - IPv6 and NATed IPv4 for customers
  - Or various other combinations of IPv6, IPv4 and NAT

Prolonging IPv4 to help with IPv6

- Large variety of proposals to “make IPv4 last longer” to help with IPv6 deployment
- All involve Large Scale NAT (LSN)
  - NAT444/SP NAT
    - NAT to customer, NAT’ed core.
  - Dual Stack Lite
    - Private IPv4 to IPv6 to Public IPv4
    - Activity of IETF Softwire Working Group
  - NAT64 & IVI
    - Translation between IPv6 and IPv4
    - Activity of IETF Behave Working Group

Dual Stack Network

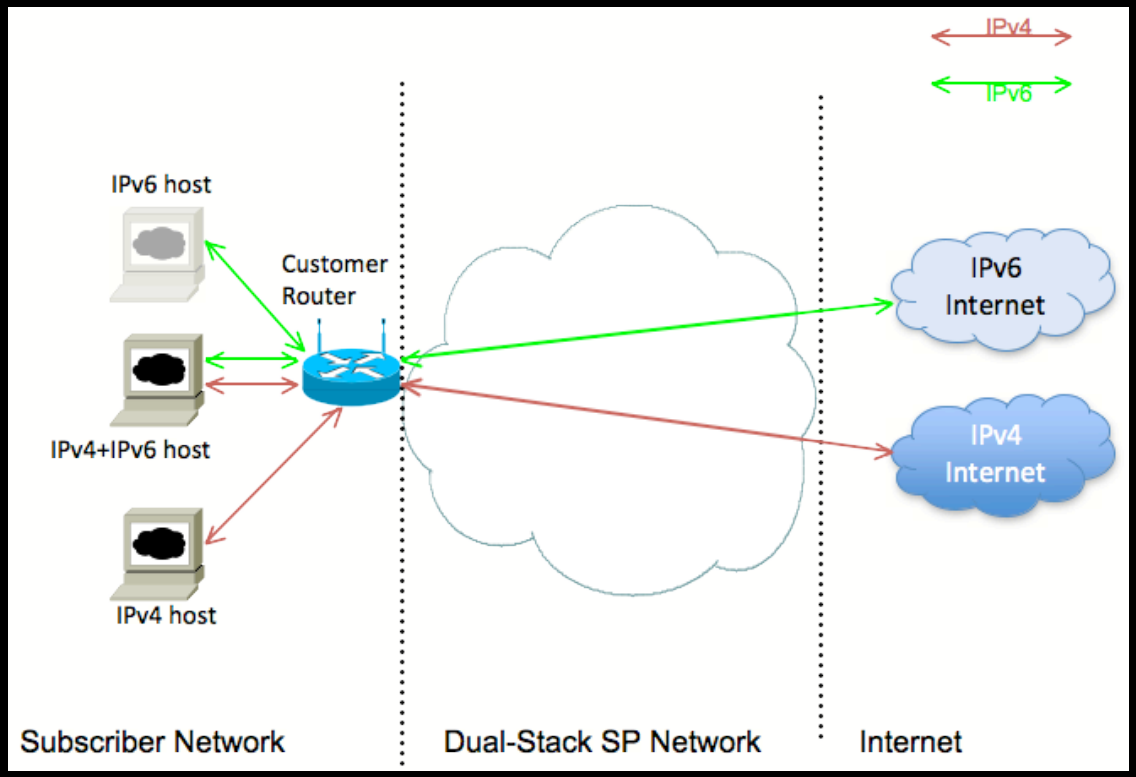

- The original transition scenario, but dependent on:
  - IPv6 being available all the way to the consumer
  - Sufficient IPv4 address space for the consumer

NAT444/SP NAT

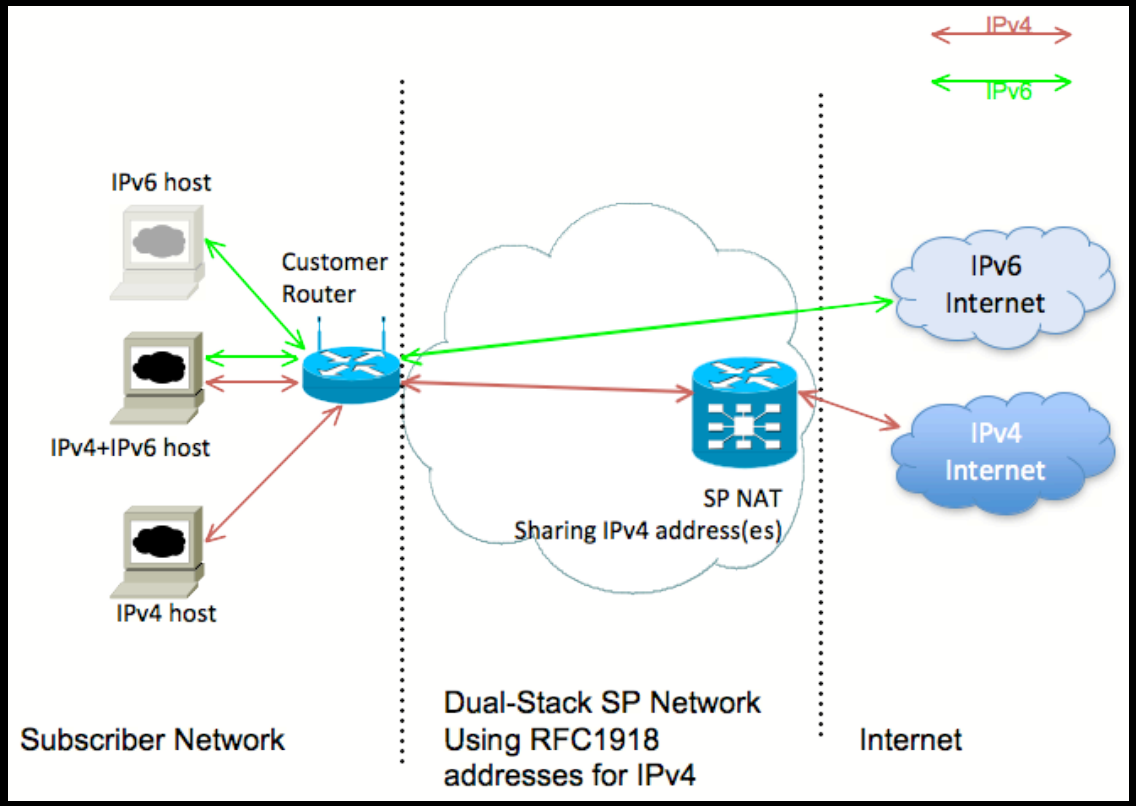

- Consumer uses private IPv4 and native IPv6
- SP uses private IPv4 and native IPv6 for backbone

DualStack-Lite

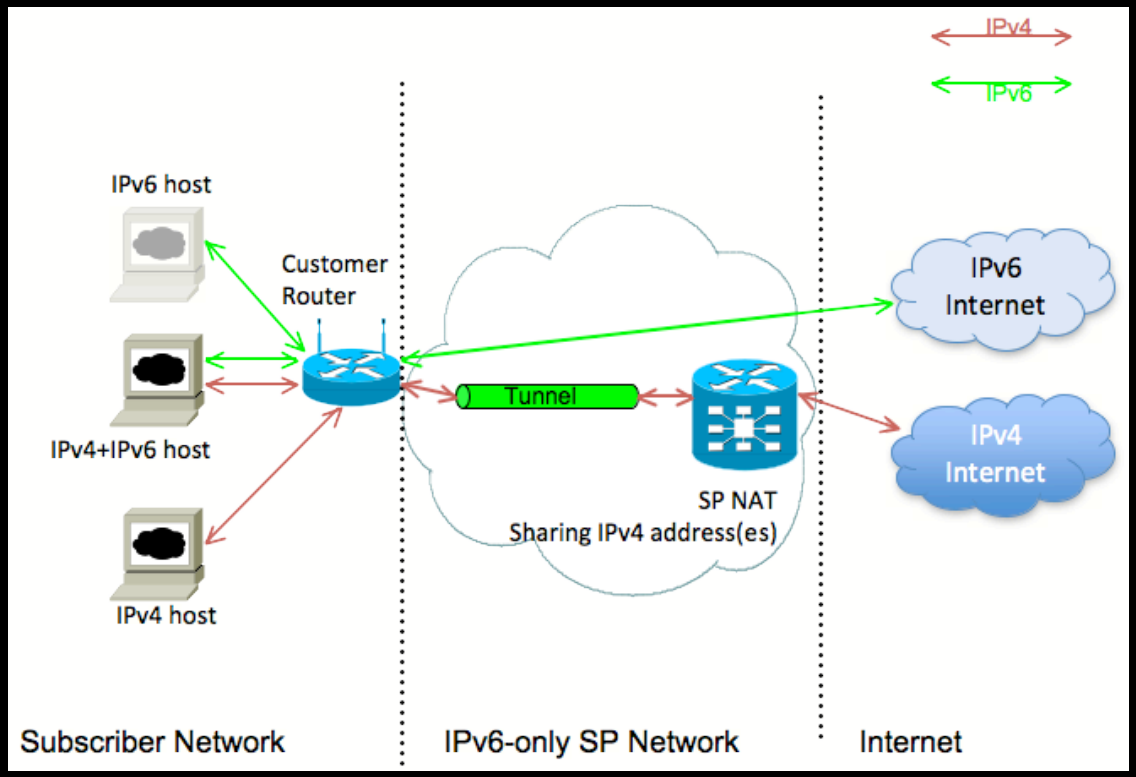

- SP has IPv6 only infrastructure
- For consumer, IPv4 tunnel to SP NAT, IPv6 native

NAT64

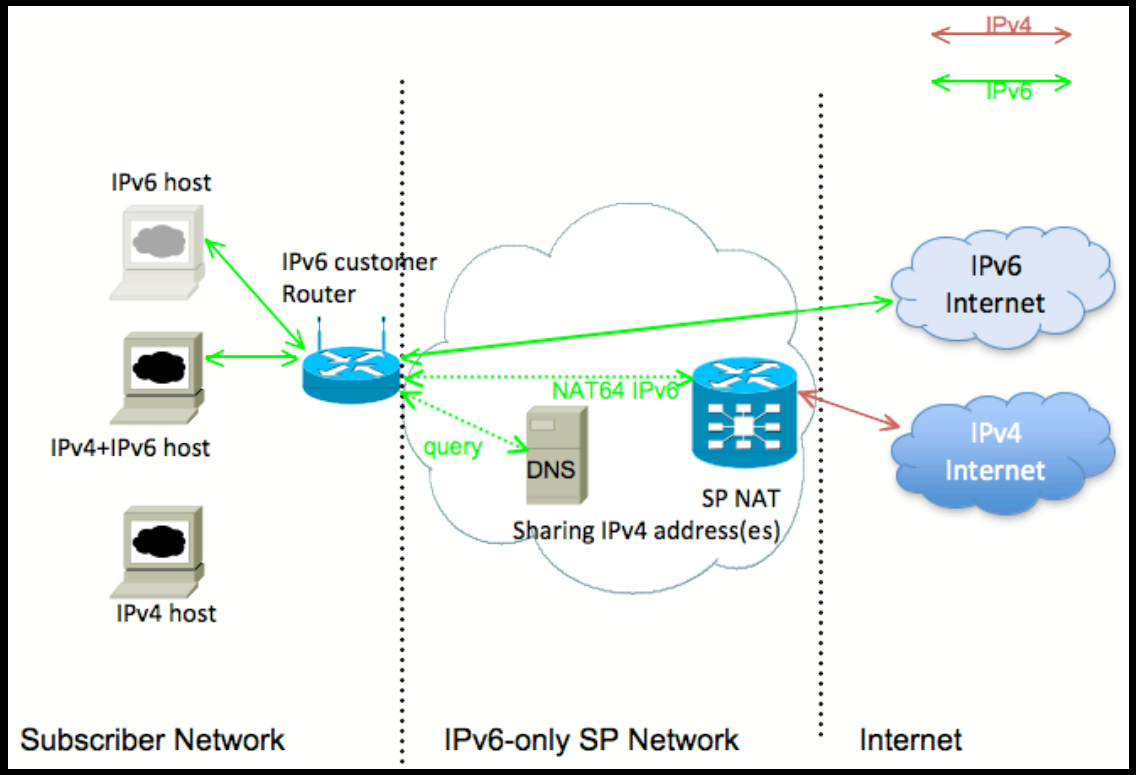

- Consumer uses only IPv6 plus Protocol Translation to reach IPv4
- Service provider uses only IPv6

IPv4 Address Markets

- Address Market:
  - When organisations don't return unused address space to their RIR (as they are supposed to do)
  - But give it to other organisations (in exchange for some form of compensation)
- If markets happen:
  - Organisations will "sell" unused portions of IPv4 address space to other organisations
    - e.g. have a /16, but two /24s are unused
  - Bypasses their RIR (but RIR will still have to register address space so that it can be routed by ISPs)

Routing Table Implications

- Assuming markets happen
  - e.g. organisation with /16 disposes of two /24s
  - Can no longer announce just the /16
  - Have to announce component parts, excluding two /24s
  - One routing announcement replaced by many
- What will happen to the IPv4 Routing Table?
  - Table today is 302k prefixes, of which 158k are /24s
  - Growth is faster than it has been since introduction of CIDR
  - Deaggregation is growing too – Routing Table could theoretically be reduced to 139k prefixes today
  - Source: http://thyme.apnic.net/current/

Closing Thoughts

- IPv6 is part of our lives now
  - IPv6 means the Internet can carry on growing
  - Not totally clear exactly how pervasive it will become
  - IPv4 is not going away any time soon