A Practical IPv6 Introduction

Philip Smith
philip@nsrc.org

PacNOG 17

Apia

13th – 17th July 2015

Last updated 11th July 2015

Presentation Slides

- Will be available on
  - http://bgp4all.com/ftp/seminars/PacNOG17-IPv6-Practical-Introduction.pdf
  - And on the PacNOG 17 website
- Feel free to ask questions any time

Agenda

- Why IPv6?
- What changes?
- Transition Technologies
- How to plan Transition to IPv6

"The times, They are a' changin'"

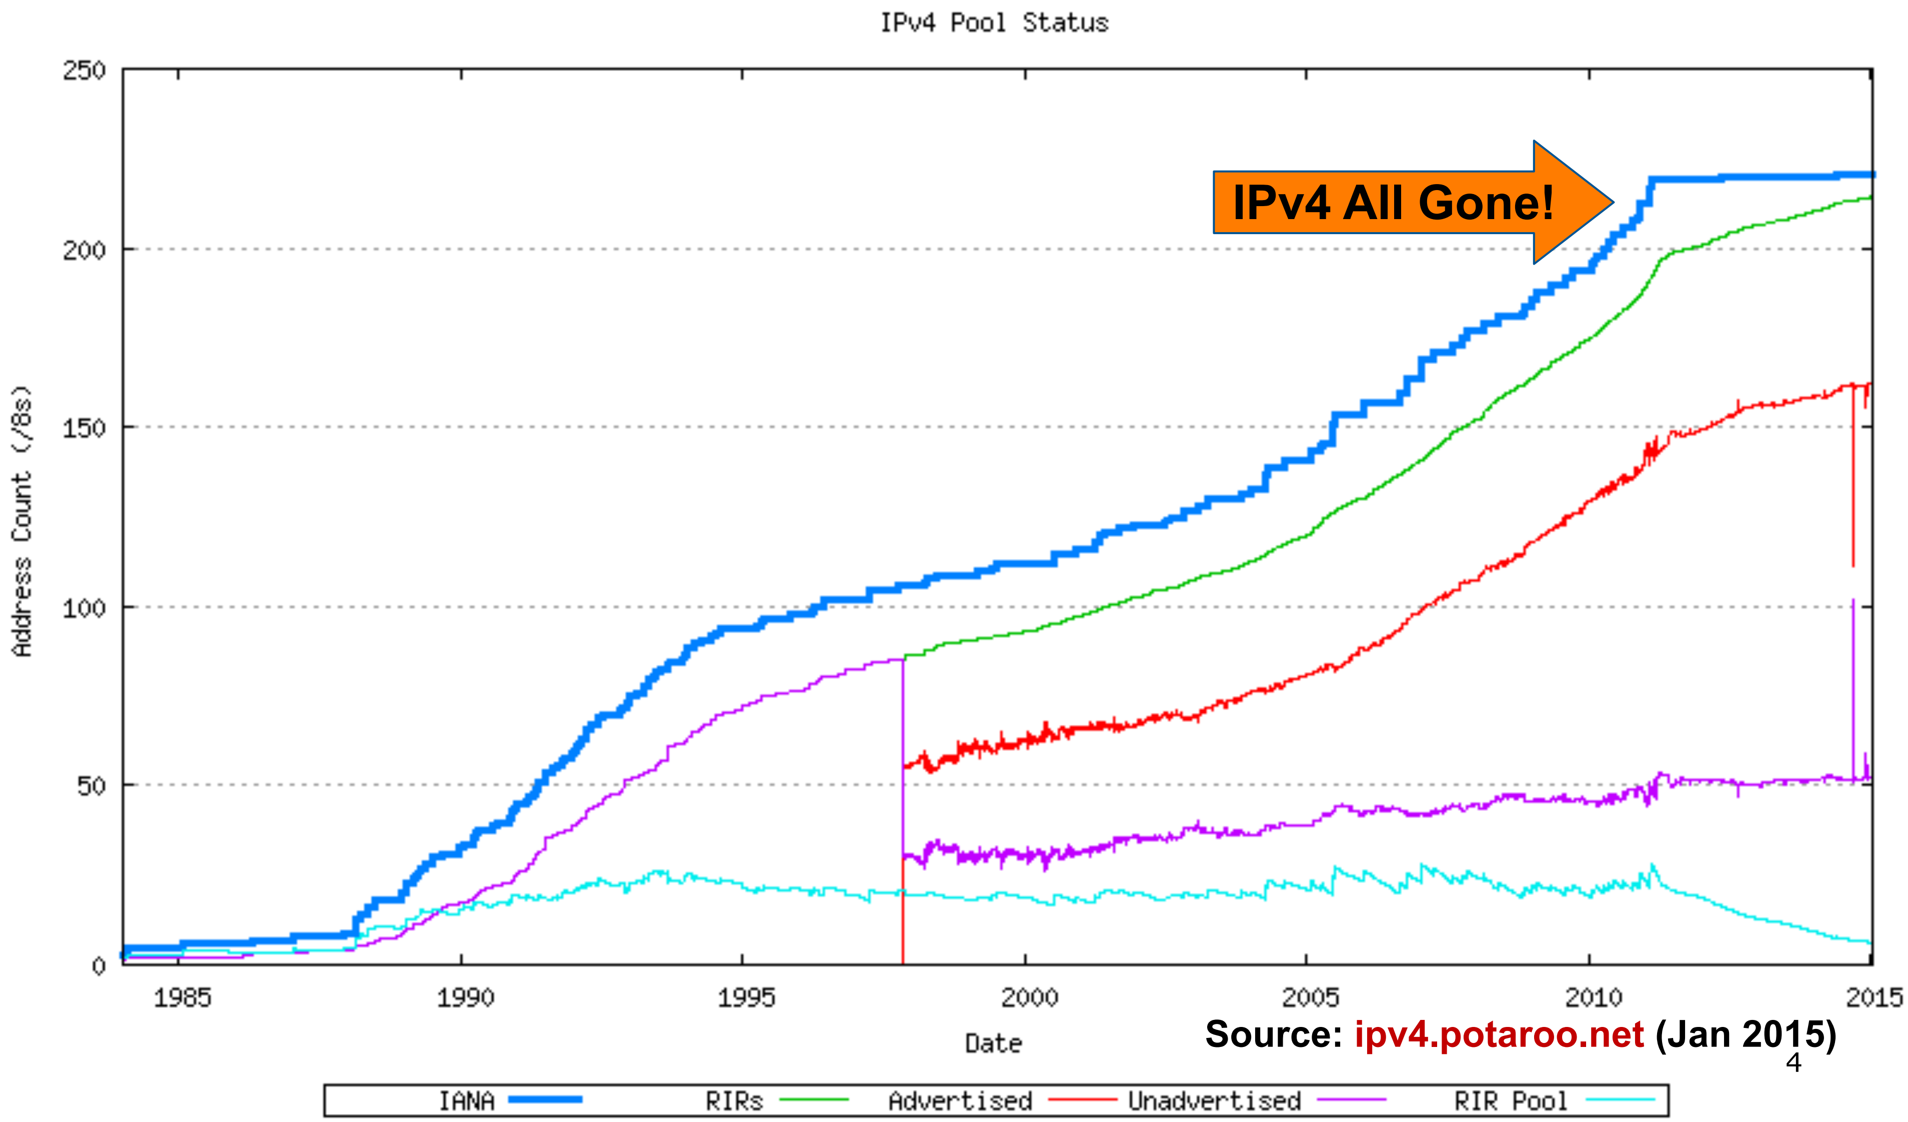

Source: ipv4.potaroo.net (Jan 2015)

Is IPv4 really running out?

- Yes!
  - IANA IPv4 free pool ran out on 3rd February 2011
  - RIR IPv4 free pool is starting to run out now
    - www.potaroo.net/tools/ipv4/
    - (depends on RIR soft-landing policies)
- The run-out gadgets and widgets are now watching when the RIR pools will run out:
  - inetcore.com/project/ipv4ec/index_en.html
    - (shows 1 RIR with no IPv4 left, and 3 out of 4 RIRs in run out austerity phase)
  - ipv6.he.net/statistics/

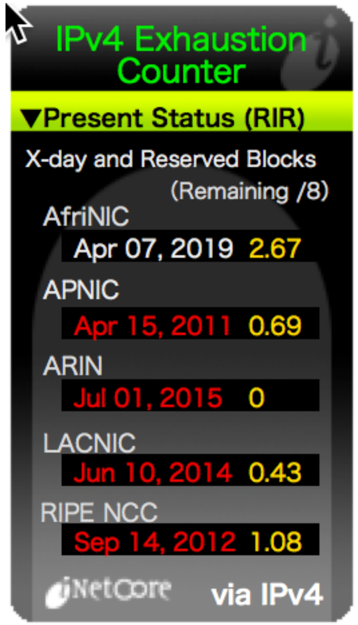

IPv4 run-out

- Policy Development process in each RIR region has discussed and implemented many proposals relating to IPv4 run-out, for example:
  - The Last /8
    - In 2011, all RIRs received one remaining /8 from IANA
  - 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 Service Provider NATs etc)

Agenda

- Why IPv6?
- What changes?
- Transition Technologies
- How to plan Transition to IPv6

So what has really changed?

- IPv6 does not interoperate with IPv4
  - Separate protocol working independently of IPv4
  - Deliberate design intention
  - Simplify IP headers to remove unused or unnecessary fields
  - Fixed length headers to “make it easier for chip designers and software engineers”

What else has changed?

- Expanded address space
  - Address length quadrupled to 16 bytes
- Header Format Simplification
  - Fixed length, optional headers are daisy-chained
  - IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)
- No checksum at the IP network layer
- No hop-by-hop fragmentation
  - Path MTU discovery
- 64 bits aligned
- Authentication and Privacy Capabilities
  - IPsec is integrated
- No more broadcast

Larger Address Space

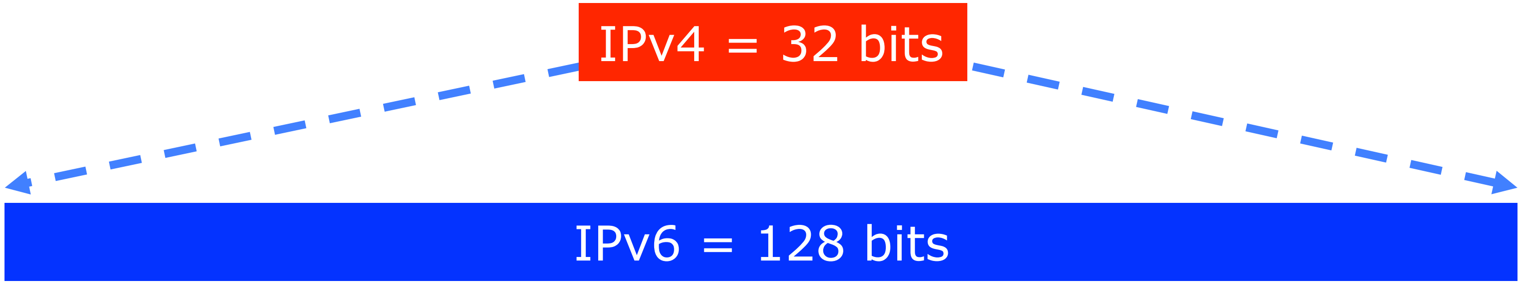

- IPv4

  32 bits

  = 4,294,967,296 possible addressable devices

- IPv6

  128 bits: 4 times the size in bits

  = 3.4 x 10^{38} possible addressable devices

  = 340,282,366,920,938,463,463,374,607,431,768,211,456

  = 4.6 x 10^{28} addresses per person on the planet

IPv6 Address Representation

- 16 bit fields in case insensitive colon hexadecimal representation
  - 2031:0000:130F:0000:0000:09C0:876A:130B
- Leading zeros in a field are optional:
  - 2031:0:130F:0:0:9C0:876A:130B
- Successive fields of 0 represented as ::, but only once in an address:
  - 2031:0:130F::9C0:876A:130B is ok
  - 2031::130F::9C0:876A:130B is **NOT** ok

IPv6 Address Representation

- Prefix Representation
  - Representation of prefix is just like IPv4 CIDR
  - In this representation you attach the prefix length
  - Like IPv4 address:
    - 198.10.0.0/16
  - IPv6 address is represented in the same way:
    - 2001:db8:12::/40

IPv6 Address Allocation

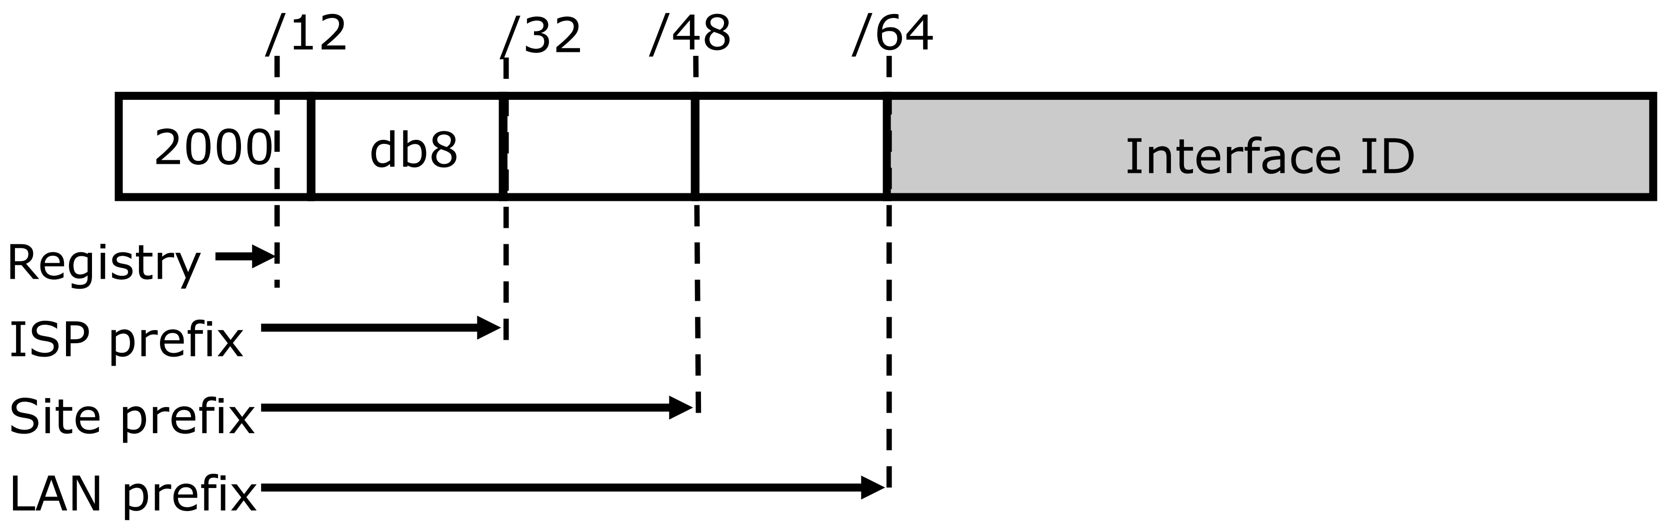

- The allocation process is:
  - The IANA is allocating out of 2000::/3 for initial IPv6 unicast use
  - Each registry gets a /12 prefix from the IANA
  - Registry allocates a /32 prefix (or larger) to an IPv6 ISP
  - Policy is that an ISP allocates a /48 prefix to each end customer

Dynamic Routing Protocols in IPv6

- Dynamic Routing in IPv6 is unchanged from IPv4:
  - IPv6 has 2 types of routing protocols: IGP and EGP
  - IPv6 still uses the longest-prefix match routing algorithm
- IGP
  - RIPng (RFC 2080)
  - Cisco EIGRP for IPv6
  - OSPFv3 (RFC 5340)
  - Integrated IS-ISv6 (RFC 5308)
- EGP
  - MP-BGP4 (RFC 4760 and RFC 2545)

IPv6 and DNS

- Hostname to IP address:

IPv4	www.abc.test.	A	192.168.30.1
IPv6	www.abc.test.	AAAA	2001:db8:c18:1::2

IPv6 and DNS

- IP address to Hostname:

IPv4	1.30.168.192.in-addr.arpa.	PTR	www.abc.test.
IPv6	2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0.8.b.d. 0.1.0.0.2.ip6.arpa	PTR	www.abc.test.

IPv6 Technology Scope

IP Service	IPv4 Solution	IPv6 Solution
Addressing Range	32-bit, Network Address Translation	128-bit, Multiple Scopes
Autoconfiguration	DHCP	DHCP, Serverless, Reconfiguration
Security	IPsec	IPsec works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality of Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
Multicast	IGMP, PIM, Multicast BGP	MLD, PIM, Multicast BGP, Scope Identifier

Agenda

- Why IPv6?
- What changes?
- Transition Technologies
- How to plan Transition to IPv6

Strategies available for Service Providers

- Do nothing
  - Wait and see what competitors do
  - Business not growing, so don't care what happens
- Extend life of IPv4
  - Force customers to NAT
  - Buy IPv4 address space on the marketplace
- Deploy IPv6
  - Dual-stack infrastructure with or without NATed IPv4 for customers
  - 6rd (Rapid Deploy) with native or NATed IPv4 for customers
  - 464XLAT with native IPv6 and NATed IPv4 for customers

IPv6 Deployment (Realistic Options)

1. Wholesale Backbone Operator
   - Solution: Dual-stack infrastructure
2. Fixed link access Network Operator
   - Cable, ADSL, Wireless, Fixed-Link, FTTx
   - Solution: Dual-stack infrastructure with or without NATed IPv4 for customers
3. Fixed link access Network Operator
   - Cable, ADSL (3rd party or legacy infrastructure)
   - Solution: Dual-stack infrastructure, with 6rd (Rapid Deploy) giving tunnelled IPv6 and native or NATed IPv4 for customers
4. Mobile Operator (3G, 4G, etc)
   - Solution: Dual-stack infrastructure, with 464XLAT giving native IPv6 and NATed IPv4 for customers

Agenda

- Why IPv6?
- What changes?
- Transition Technologies
- How to plan Transition to IPv6

Planning Transition

- Goals
- Network Assessment
- Network Optimisation
- Procuring IPv6 Address Space
- IPv6 Address plan
- Deployment
- Seeking IPv6 Transit
- Customers

Goals

What do we want to achieve?

Goals

- Ultimate aim is to provide IPv6 to our customers:
  - Customers = end users
  - Customers = content providers
  - Customers = retail providers
- Strategy depends on network transport:
  - Native IP backbone
    - Dual Stack is the solution
  - MPLS backbone (tunnels)
    - 6PE or 6VPE is the solution
    - The core infrastructure will remain IPv4 only

Network Assessment

What can run IPv6 today, and what needs to be upgraded?

Audit

- First step in any deployment:
  - Review existing network infrastructure
- Primarily routers across backbone
  - Perhaps also critical servers and services (but not essential as initial focus is on routing infrastructure)

Process

- Analyse each location/PoP
- Document
  - Router or any other L3 device
  - RAM (installed and used)
  - FLASH memory
  - Software release versions
  - Most network operators already keep track of this info
    - If not, RANCID (www.shrubbery.net/rancid/) makes this very easy
- Sanity check
  - Check existing connectivity
  - Remove unused configuration
  - Shutdown and clean up unused interfaces

Software Issues (1)

- Does the existing software have IPv6 support?
  - Yes: deployment is straightforward
  - No: investigate cost of upgrade
- Is a software upgrade available?
  - Yes: is hardware suitably specified?
  - No: hardware replacement
- Implement software upgrade
  - Budget, purchase & schedule installation

Software Issues (2)

- If existing software supports IPv6:
  - Are deployed software versions consistent across infrastructure?
    - Recommend maximum of two variations (easier troubleshooting, bug tolerance, etc)
- If existing software does not support IPv6:
  - Cost of upgrade to a version which does?
  - Testing for existing feature compatibility:
    - A software image with IPv6 may have "lost" features required for the existing operational network

Hardware Issues

- Can hardware specification be upgraded (eg RAM, FLASH etc)?
  - Yes: budget, purchase, installation
  - No: hardware replacement
- Hardware replacement:
  - Assess suitable replacement product
  - Analyse impact on operating network, existing services and customer

Result

- Once the previous steps are completed, entire network is running IPv6 capable software
- Deployment of IPv6 can now begin

Network Optimisation

Is the IPv4 network the best it can be?

Optimisation

- IPv4 networks have been deployed and operational for many years
  - Your network may fall into this category
- Optimisation means:
  - Does the interior routing protocol make sense?
  - Do all routing protocols have the latest best practices implemented?
  - Are the IGP metrics set so that primary and backup paths operate as expected?

Motivation for Optimisation

- IPv6 deployment (apart from MPLS cores) will be dual stack
  - Which means sitting alongside existing IPv4 configurations
- Aim is to avoid replicating IPv4 “shortcuts” or “mistakes” when deploying IPv6
  - IPv6 configuration will **replicate** existing IPv4 configuration
- Improvements in routing protocol BCPs should be deployed and tested for IPv4
  - Take the opportunity to “modernise” the network

Procuring IPv6 address space

Now we need addresses…

Getting IPv6 address space

- From your Regional Internet Registry
  - Membership open to all organisations who are operating a network
  - Minimum allocation is a /32
  - Assignment of /48 is also available for end-sites, multihoming organisations, and IXPs
- From your upstream ISP
  - Receive a /48 from upstream ISP's IPv6 address block

Address Planning

- IPv6 address space available to each network operator is large compared with IPv4
  - Design a scalable plan
  - Be aware of industry current practices
  - Separation of infrastructure and customer addressing
  - Distribution of address space according to function

Addressing Plans – Infrastructure

- Address block for router loop-back interfaces
  - Number all loopbacks out of **one** /64
  - /128 per loopback
- Address block for infrastructure (backbone)
  - /48 for whole backbone, allows 65k subnets
- What about LANs?
  - /64 per LAN
- What about Point-to-Point links?
  - Protocol design expectation is that /64 is used
  - /127 now recommended/standardised
    - http://www.rfc-editor.org/rfc/rfc6164.txt
    - (reserve /64 for the link, but address it as a /127)

Addressing Plans – Customer

- Customers get one /48
  - Unless they have more than 65k subnets in which case they get a second /48 (and so on)
- In typical deployments today:
  - Several ISPs are giving small customers a /56 and single LAN end-sites a /64, e.g.:

/64	if end-site will only ever be a LAN
/56	for small end-sites (e.g. home/office/small business)
/48	for large end-sites

- Observations:
  - Don't assume that a mobile endsite needs only a /64
  - Some operators are distributing /60s to their smallest customers!!

Deploying IPv6

Now we put it onto the network

IPv6 Deployment

- Number all the infrastructure interfaces according to the established addressing plan
  - No customers yet
- Secure routers and L3 devices for IPv6 access
- Enable IPv6 internal routing protocols
  - First IGP – care needed not to break IPv4 connectivity
  - iBGP – should replicate IPv4 iBGP
- Check that operation compares with IPv4 operation
  - Fix any problems – in a dual stack network the protocols must function the same way

Seeking IPv6 Transit

Hello World, I'd like to talk to you…

Seeking Transit

- Most transit providers now offer native IPv6 transit
- Next step is to decide:
  - To give transit business to those who will accept a dual stack connection

    or
  - To stay with existing IPv4 provider and seek a tunnelled IPv6 transit from an IPv6 provider

Dual Stack Transit Provider

- Fall into two categories:
  A. Those who sell you a pipe over which you send packets
  B. Those who sell you an IPv4 connection and charge extra to carry IPv6
- ISPs in category A are much preferred to those in category B
- Charging extra for native IPv6 is absurd, given that this can be easily bypassed by tunnelling IPv6
  - IPv6 is simply protocol 41 in the range of IP protocol numbers

Dual Stack Transit Provider

- Advantages:
  - Can align BGP policies for IPv4 and IPv6 – perhaps making them more manageable
  - Saves money – they charge you for bits on the wire, not their colour
- Disadvantages:
  - Not aware of any

Separate IPv4 and IPv6 transit

- Retain transit from resolute IPv4-only provider
  - You pay for your pipe at whatever $ per Mbps
- Buy transit from an IPv6 provider
  - You pay for your pipe at whatever $ per Mbps
- Luck may uncover an IPv6 provider who provides transit for free
  - Getting more and more rare as more ISPs adopt IPv6

Separate IPv4 and IPv6 transit

- Advantages:
  - Not aware of any
  - But perhaps situation is unavoidable as long as main IPv4 transit provider can't provide IPv6
  - And could be a tool to leverage IPv4 transit provider to deploy IPv6 – or lose business
- Disadvantages:
  - Do the $$ numbers add up for this option?
  - Separate policies for IPv4 and IPv6 – more to manage

Customer Connections

Network is done, now let's connect paying customers…

Customer Connections

- Giving connectivity to customers is the biggest challenge facing all ISPs
- Needs special care and attention, even updating of infrastructure and equipment
  - Cable/ADSL
  - Dial
  - Leased lines
  - Wireless Broadband

IPv6 to Broadband Customers

- Method 1: Use existing technology and CPE
  - This is the simplest option – it looks and feels like existing IPv4 service
  - PPPoE + DHCPv6 PD
  - Used by ISPs such as Internode (AU) and XS4ALL (NL)
- Issues:
  - IPv6 CPE are not as widely available yet as “throw-away” IPv4-only CPE

IPv6 to Broadband Customers

- Method 2: Use 6rd
  - This is for when Broadband infrastructure cannot be upgraded to support IPv6
  - Used by ISPs such as FREE (FR) and Softbank (JP)
  - Example:
    - 2001:db8:6000::/48 assigned to 6rd
    - Customer gets 192.168.4.5/32 by DHCP for IPv4 link
    - IPv6 addr is 2001:db8:6000:0405::/64 for their LAN (taking last 16 bits of IPv4 address)
    - DHCPv6 PD can be used here too (eg to give /56s to customers)
- Issues:
  - All CPE needs to be replaced/upgraded to support 6rd

IPv6 to Fixed Link Customers

- Use existing technology:
  - Most access routers (PE) and Customer routers (CPE) are easily upgradeable or replaceable to include IPv6 support
  - Service looks and feels like existing IPv4 service
- Configuration options:
  - IPv6 unnumbered on point to point links (or address them)
  - Static routes, subnet size according to business size
  - Or use BGP with private or public (multihomed) ASN
  - Whatever is done for IPv4 should be repeated for IPv6
- Fixed link Customers are probably the easiest to roll IPv6 out to
  - Customer deploying IPv6 within their own networks is a separate discussion (rerun of this presentation!)

Customer Connections

- What about customer end systems?
  - Is IPv6 available on all their computers and other network connected devices?
  - How to migrate those which aren't?
  - How to educate customer operations staff
  - What about their CPE?
  - What about the link between your edge device and their CPE?
  - What about security?

Conclusion

We are done...!

Conclusion

- When deploying IPv6 for the first time, a strategy and planning are of paramount importance
- Presentation has highlighted the steps in the planning and presentation process
  - Variations on the theme are quite likely – there is no single correct way of proceeding

A Practical IPv6 Introduction

PacNOG 17