IPv6 Technical & Service Considerations

Philip Smith <pfs@cisco.com>

PacNOG 5

Papeete

French Polynesia

15th-21st June 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 numbers
  - Explosive growth of the “default-free” routing table
  - Eventual exhaustion of 32-bit address space
- Two efforts – short-term vs. 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
  - SIPP – 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 → 2004

- The DotCom “crash”

    i.e. Internet became mainstream

- IPv4:

    Consumption slowed

    Address space pressure “reduced”

- Indifference

    Early adopters surging onwards

    Sceptics more sceptical

    Yet more transition mechanisms developed

2004 → Today

- Resurgence in demand for IPv4 address space

  13.6% address space still unallocated (04/2009)

  Exhaustion predictions ranged from wild to conservative

  …but mid 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 use of NATs

- IPv6 offers vast address space

  The only compelling reason for IPv6

Current Situation

- General perception is that "IPv6 has not yet taken hold"
  - Many discussions and run-out plans 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?

Do we really need a larger address space?

- Internet population

  ~630 million users end of 2002 – 10% of world pop.

  ~1320 million users end of 2007 – 20% of world pop.

  Future? (World pop. ~9B in 2050)

- US uses 81 /8s – this is 3.9 IPv4 addresses per person

  Repeat this the world over…

  6 billion population could require 23.4 billion IPv4 addresses

  (6 times larger than the IPv4 address pool)

- Emerging Internet economies need address space:

  China uses more than 94 million IPv4 addresses today (5.5 /8s)

Do we really need 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 membership 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?

Do we really need a larger address space?

- Large variety of proposals to “make IPv4 last longer” to help with IPv6 deployment
  - NAT444
    - Lots of IPv4 NAT
  - NAT464
    - IPv4 to IPv6 to IPv4 NAT
  - Dual Stack Lite
    - Improvement on NAT464
    - Activity of IETF Softwires Working Group
  - NAT64 & IVI
    - Translation between IPv6 and IPv4
    - Activity of IETF Behave Working Group

IPv6 OS and Application Support

- All software vendors officially support IPv6 in their latest Operating System releases
- Application Support
  - Applications must be IPv4 and IPv6 agnostic
  - User should not have to “pick a protocol”
  - Successful deployment is driven by Applications
  - Successful Application support is driven by Content
- Content Availability
  - Needs to be on IPv4 **and** on IPv6

ISP Deployment Activities

- Several Market segments
  - IX, Carriers, Regional ISP, Wireless
- ISP have to get an IPv6 prefix from their Regional Registry
- Large carriers planning driven by customer demand:
  - Some running trial networks (e.g. Sprint)
  - Others running commercial services (e.g. NTT, FT)
- Regional ISP focus on their specific markets
- Much discussion by operators about transition
  - www.civil-tongue.net/6and4/
  - http://www.nanog.org/mtg-0710/presentations/Bush-v6-op-reality.pdf

Why not use Network Address Translation?

- Private address space and Network address translation (NAT) could be used instead of a new protocol
- But NAT has many serious issues:
    - Breaks the end-to-end model of IP
    - Layered NAT devices
    - Mandates that the network keeps the state of the connections
    - Scaling NAT performance for large networks
    - Makes fast rerouting difficult
    - Service provision inhibited

NAT has many implications

- Inhibits end-to-end network security
- When a new application is not NAT-friendly, NAT device requires an upgrade
- Some applications cannot work through NATs
- Application-level gateways (ALG) are not as fast as IP routing
- Complicates mergers
  - Double NATing is needed for devices to communicate with each other
- Breaks security
- Makes multihoming hard
- Simply does not scale
- RFC2993 – architectural implications of NAT

Conclusion

- There is a need for a larger address space
    - IPv6 offers this – will eventually replace NAT
    - But NAT will be around for a while too
    - Market for IPv4 addresses looming also
- Many challenges ahead

IPv6 Integration/Transition

How will I need to roll out IPv6?

IPv4-IPv6 Co-existence/Transition

- A wide range of techniques have been identified and implemented, basically falling into three categories:
  - Dual-stack techniques, to allow IPv4 and IPv6 to co-exist in the same devices and networks
  - Tunneling techniques, to avoid order dependencies when upgrading hosts, routers, or regions
  - Translation techniques, to allow IPv6-only devices to communicate with IPv4-only devices
- Expect all of these to be used, in combination
  - **IPv6 is not compatible with IPv4**

Dual Stack Approach

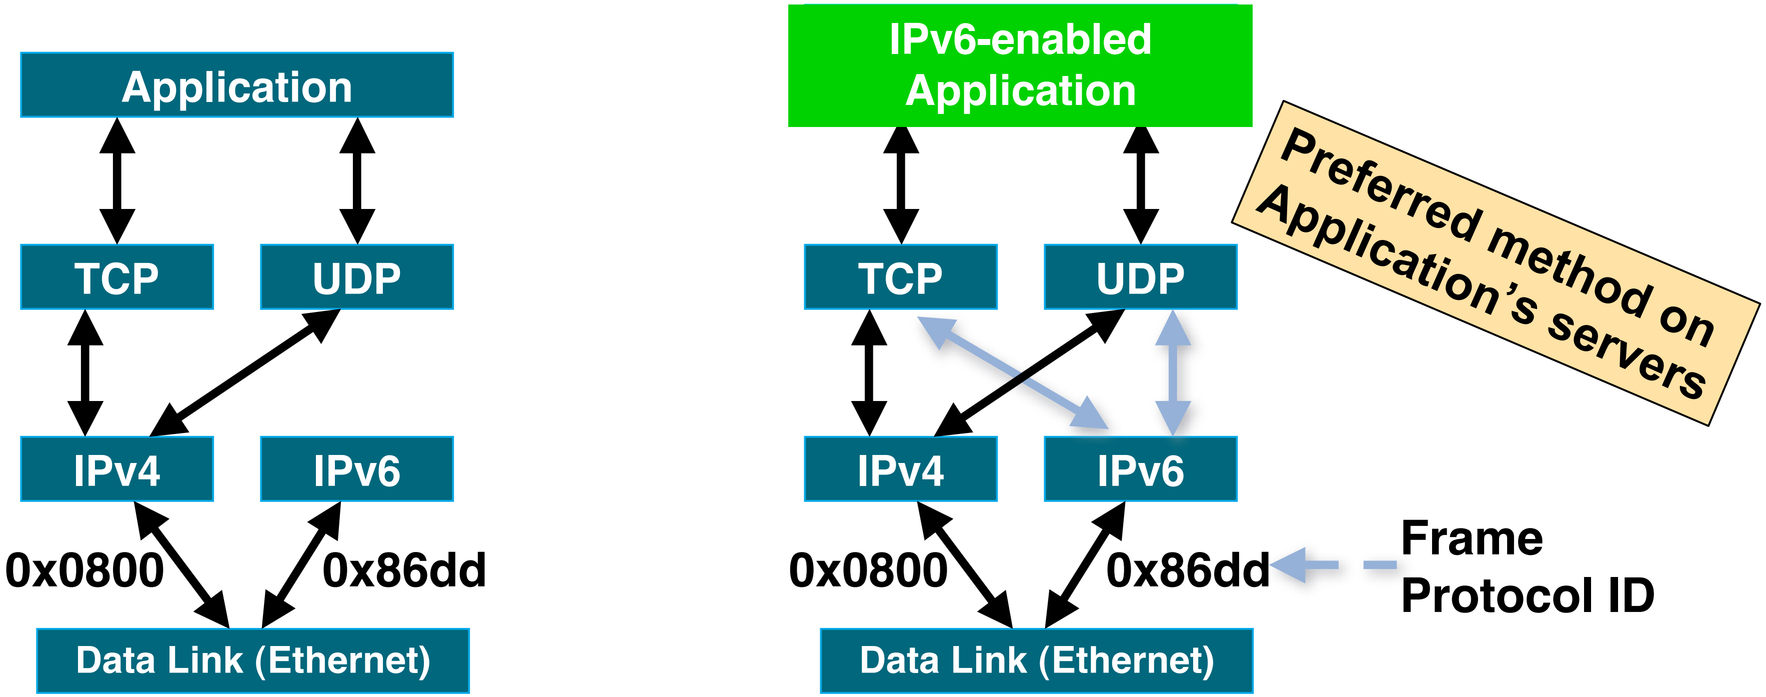

- Dual stack node means:
  - Both IPv4 and IPv6 stacks enabled
  - Applications can talk to both
  - Choice of the IP version is based on name lookup and application preference

Dual Stack Approach & DNS

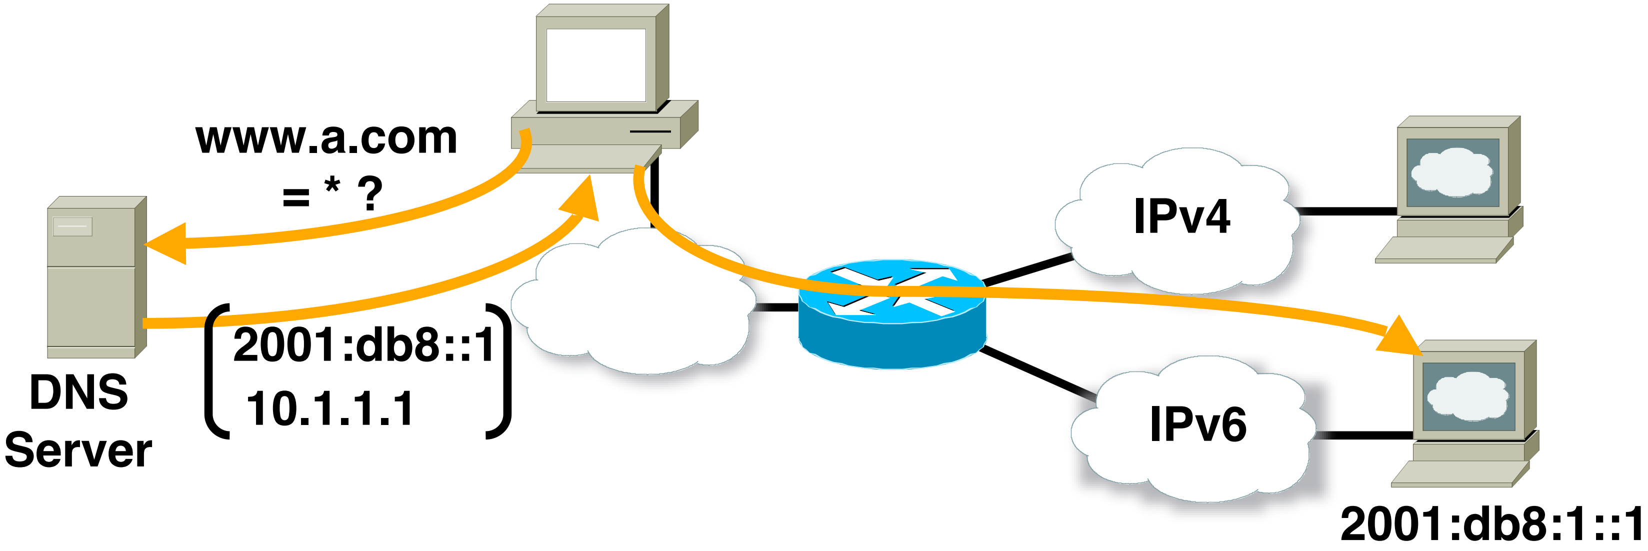

- In a dual stack case, an application that:
  - Is IPv4 and IPv6-enabled
  - Asks the DNS for all types of addresses
  - Chooses one address and, for example, connects to the IPv6 address

Using Tunnels for IPv6 Deployment

- Many techniques are available to establish a tunnel:
  - Manually configured
    - Manual Tunnel (RFC 2893)
    - GRE (RFC 2473)
  - Semi-automated
    - Tunnel broker
  - Automatic
    - 6to4 (RFC 3056)
    - ISATAP

IPv6 over IPv4 Tunnels

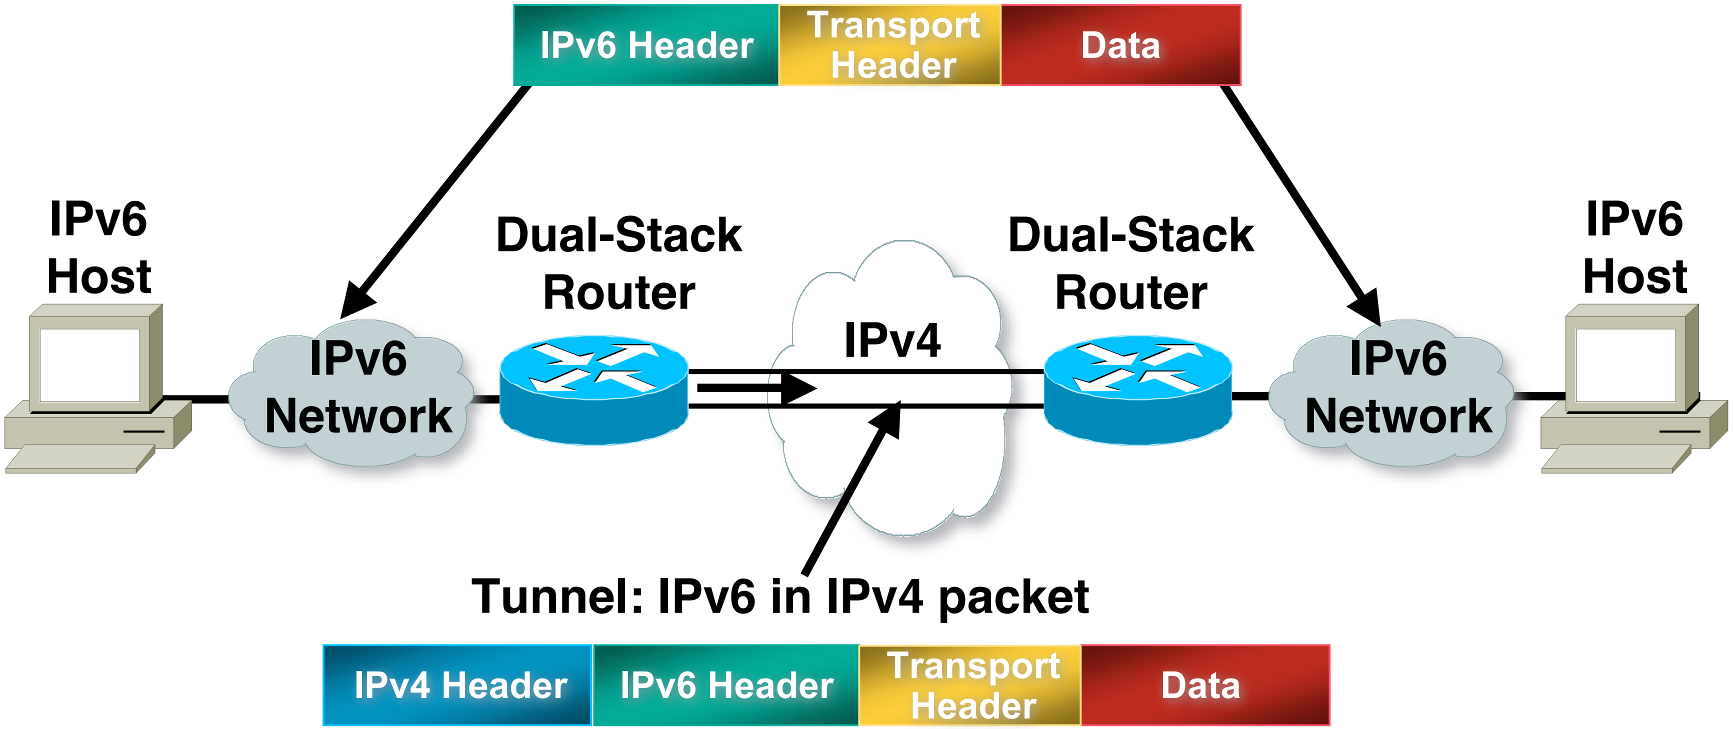

- Tunneling is encapsulating the IPv6 packet in the IPv4 packet
- Tunneling can be used by routers and hosts

ISP Technical Considerations

What does my business need to do?

Audit

- First step in any deployment:

  Audit existing network infrastructure
- Primarily routers across backbone

  Perhaps also critical servers and services (but not essential as initial focus is on routing infrastructure)

Audit Process

- Analyse each PoP
- Document
  - Router platform
  - RAM (installed and used)
  - FLASH memory
  - Software release versions
  - RANCID (www.shrubbery.net/rancid/) makes this very easy
- Sanity check
  - Check existing connectivity
  - Remove unused configuration
  - Shutdown and clean up unused interfaces

Optimisation

- IPv4 networks have been deployed and operational for many years
  - Your network may fall into this category
- Optimisation means:
  - Does the iBGP design make sense?
  - Are the OSPF areas in the right places?
  - Does the ISIS backbone 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 will be dual stack
  - So 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

iBGP considerations

- Full mesh iBGP still?
  - Perhaps consider migration to route reflectors
- Route reflector configuration
  - Proper redundancy in place?
  - Overlapping clusters, one reflector per cluster
  - Direct path between client and reflector
- BGP best practices deployed
  - Peer-group strategy? (Will have to be replicated for IPv6)
  - Full routes in core iBGP?
  - Partial routes in edge/rr client iBGP
  - Community strategy for internal and external announcements?

Getting IPv6 address space (RIR)

- Become a member of your Regional Internet Registry and get your own allocation
  - Requires a plan for a year ahead
  - IPv6 allocation policies are documented on each RIR website
  - The following slides describe considerations when constructing such a plan
- Note Well: There is plenty of IPv6 address space
  - The RIRs require high quality documentation

Getting IPv6 address space (non-RIR)

- From your upstream ISP
  - Get one /48 from your upstream ISP
  - More than one /48 if you have more than 65k subnets
- Use 6to4
  - Take a single public IPv4 /32 address
  - 2002:<ipv4 /32 address>::/48 becomes your IPv6 address block, giving 65k subnets
  - Requires a 6to4 gateway
- These two options aren't really viable for service providers though – a /32 from an RIR is the way to go

Addressing Plans – ISP Infrastructure

- Address block for router loop-back interfaces
  - Number all loopbacks out of **one** /64
  - /128 per loopback
- Address block for infrastructure
  - /48 allows 65k subnets
  - /48 per PoP or region (for large networks)
  - /48 for whole backbone (for small to medium networks)
  - Summarise between sites if it makes sense

Addressing Plans – ISP Infrastructure

- What about LANs?
  - /64 per LAN
- What about Point-to-Point links?
  - Expectation is that /64 is used
  - People have used /126s
    - Mobile IPv6 Home Agent discovery won't work
  - People have used /112s
    - Leaves final 16 bits free for node IDs
  - Some people are considering /80s or /96s
  - See RFC3627 for more discussion

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)
- Should not be reserved or assigned on a per PoP basis
  - ISP iBGP carries customer nets
  - Aggregation within the iBGP not required and usually not desirable
  - Aggregation in eBGP is very necessary

Seeking Transit

- ISPs offering native IPv6 transit are still in the minority
- Next step is to decide:

  Whether to give transit business to those who will accept a dual stack connection

  or

  Whether to stay with existing IPv4 provider and seek a tunnelled IPv6 transit from an IPv6 provider

- Either option has risks and challenges

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

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

Forward and Reverse DNS

- Populating the DNS is an often omitted piece of an ISP operation

    Unfortunately it is extremely vital, both for connectivity and for troubleshooting purposes

- Forward DNS for IPv6

    Simply a case of including suitable AAAA records alongside the corresponding A records of a host

- Reverse DNS for IPv6

    Requires getting the /32 address block delegated from the RIR, and then populating the ip6.arpa fields

Conclusion

- IPv6 deployment for an ISP replicates existing IPv4 infrastructure

    Need to consider IPv6 support on existing hardware

    Need to procure and deploy IPv6 resources

- What about customer services?

    IPv6 over <*??*> ?

    IPv6 on customer router?

- What about content?

    Dual stack service provision should be easy

IPv6 Technical & Service Considerations

Philip Smith <pfs@cisco.com>

PacNOG 5

Papeete

French Polynesia

15th-21st June 2009