

# IPv6 Deployment Planning



Philip Smith

<philip@nsrc.org>

SANOG 29

Islamabad, Pakistan

23<sup>rd</sup> – 30<sup>th</sup> January 2017

# Presentation Slides

---

- Will be available on
  - <http://bgp4all.com/dokuwiki/conferences/>
  - And on the SANOG29 website
- Feel free to ask questions any time

# Introduction

---

- Presentation introduces the high level planning considerations which any network operator needs to be aware of prior to deploying IPv6
- Content applicable for:
  - Business decision makers
  - Network managers
  - Network engineers
    - Will also require implementation detail

# Agenda

---

1. Goals
2. Network Assessment
3. Network Optimisation
4. Procuring IPv6 Address Space
5. IPv6 Address plan
6. Deployment
7. Seeking IPv6 Transit
8. Customers

# Goals



What do we want to achieve?

# Goals

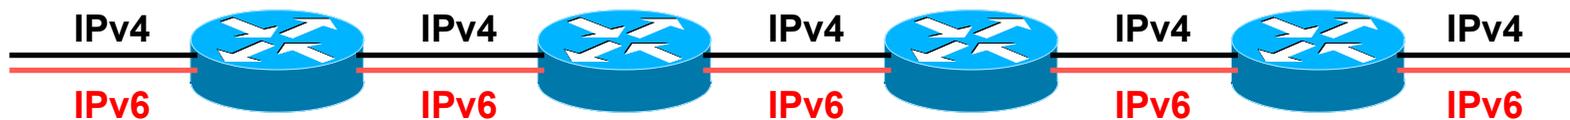
---

- Ultimate aim is to provide IPv6 to our customers:
  - Customers = end users
  - Customers = content 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

# Native IP Backbone

---

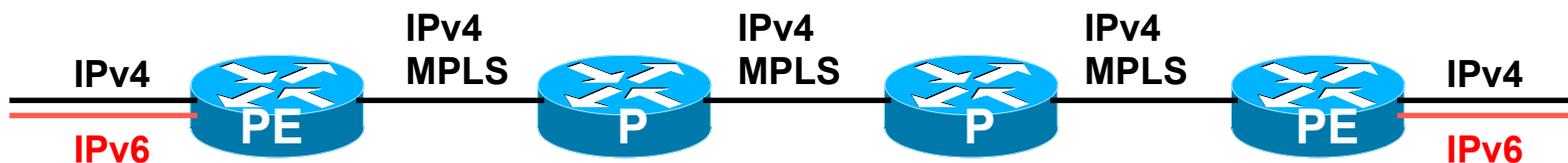
- Routers are the infrastructure
  - Customer connections connect to the native backbone
  - VPN services provided using GRE, IPSEC, IPinIP etc
  - Providing IPv6 for customers means upgrading the native infrastructure to dual-stack



# MPLS Backbone

---

- Routers are the infrastructure
  - Public and Private network access provided within the MPLS cloud
  - The core network does NOT need to be IPv6 aware
  - IPv6 access provided by 6PE or 6VPE
  - Provider Edge routers need dual stack capability



# 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)
  - Non-Volatile Configuration memory
  - Software release versions
  - Most network operators already keep track of this info
    - ❑ If not, RANCID ([www.shrubbery.net/rancid/](http://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, NVRAM, 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



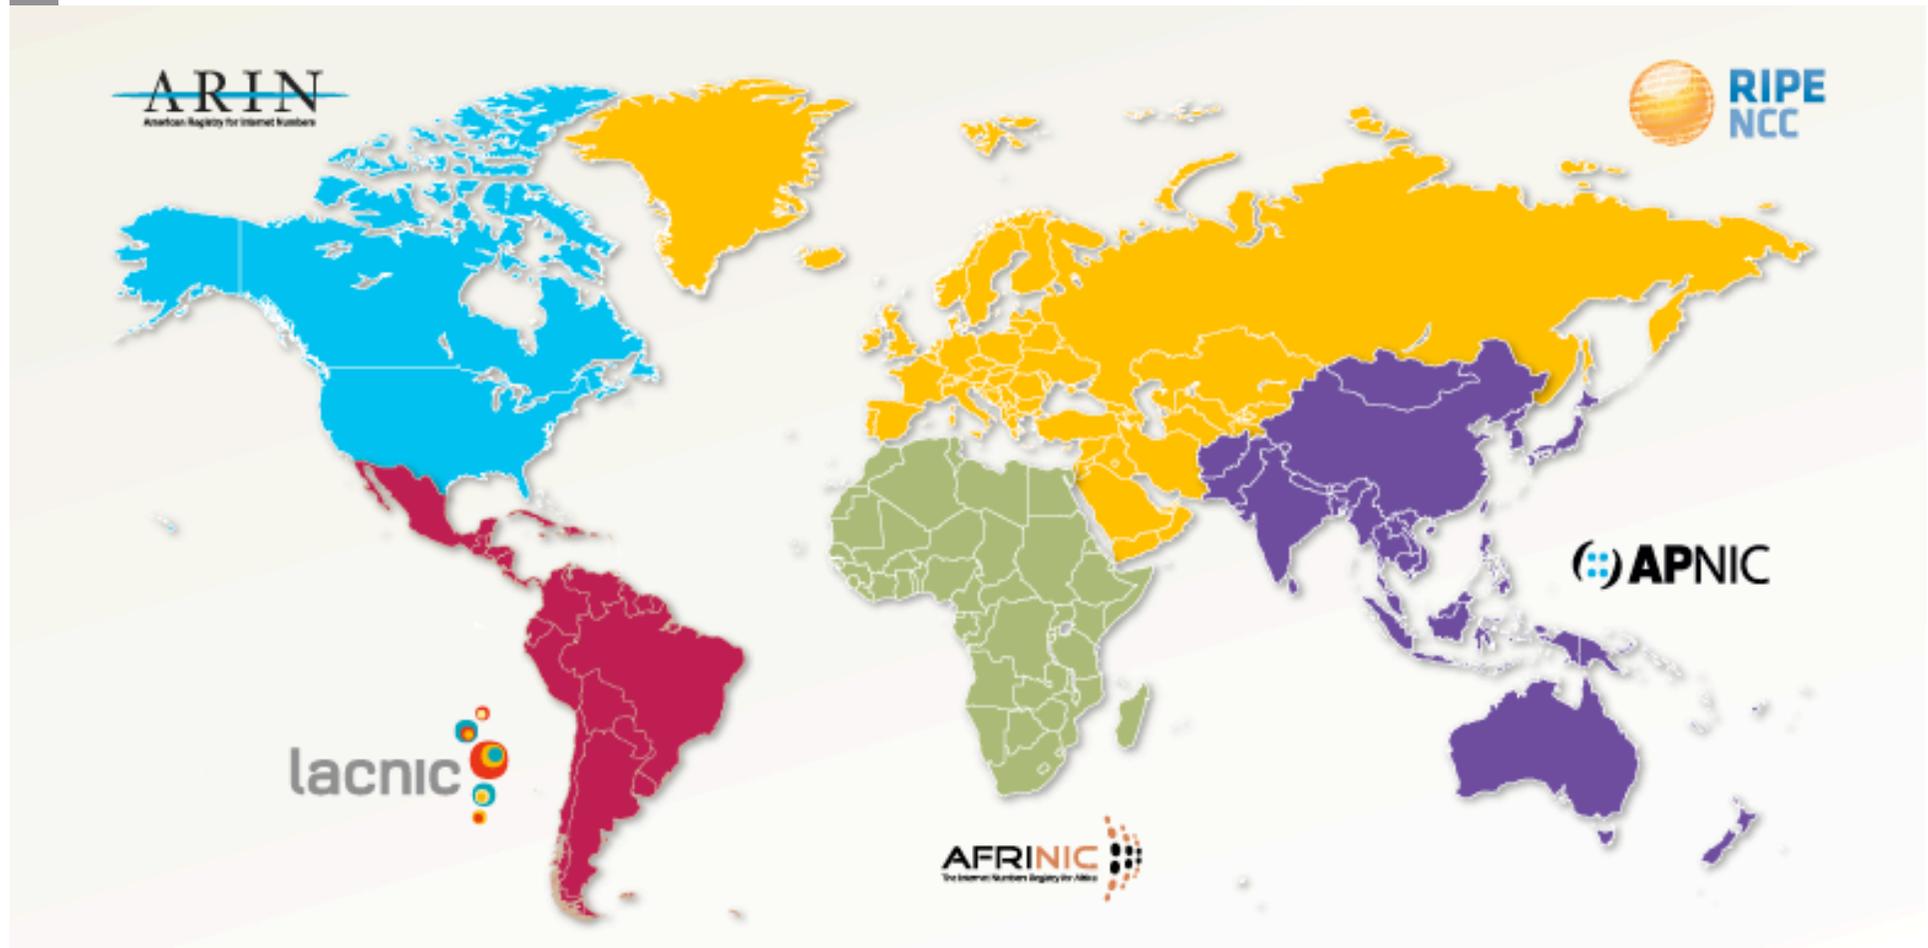
Now we need addresses...

# Where to get IPv6 addresses

---

- Your upstream ISP
- Africa
  - AfriNIC – <http://www.afrinic.net>
- Asia and the Pacific
  - APNIC – <http://www.apnic.net>
- North America
  - ARIN – <http://www.arin.net>
- Latin America and the Caribbean
  - LACNIC – <http://www.lacnic.net>
- Europe and Middle East
  - RIPE NCC – <http://www.ripe.net/info/ncc>

# Internet Registry Regions



# Getting IPv6 address space (1)

---

- If existing Regional Internet Registry account holder with an IPv4 allocation:
  - Just ask for an IPv6 allocation and it will be given – it really is as simple as that!
- Become a member of your Regional Internet Registry and get your own allocation
  - Membership usually open to all network operators
  - RIR specific details for IPv6 allocations are listed on the individual RIR website
  - Open to all organisations who are operating a network
  - Receive a /32 (or larger if you will have more than 65k /48 assignments)

# Getting IPv6 address space (2)

---

- **From your upstream ISP**
  - Receive a /48 from upstream ISP's IPv6 address block
  - Receive more than one /48 if you have more than 65k subnets
- **If you need to multihome:**
  - Apply for a /48 assignment from your RIR
  - Trying to multihome with provider's /48 will be operationally challenging
    - Provider policies, filters, etc

# 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 is critical
  - Distribution of address space according to function

# Addressing Plans – Infrastructure

---

- ❑ Network Operators should procure a /32 from their RIR
- ❑ Address block for infrastructure
  - /48 allows 65k subnets in the backbone
- ❑ Address block for router loop-back interfaces
  - Number all loopbacks out of one infrastructure /64
  - /128 per loopback
- ❑ Point-to-point links
  - /64 reserved for each, address as a /127
- ❑ LANs
  - /64 for each LAN

# 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)
- Industry standard for customer assignments today:
  - /64 for just one LAN
  - /56 for a small network
  - /48 for a large network

# Addressing Tools

---

- Examples of IP address planning tools:
  - NetDot [netdot.uoregon.edu](http://netdot.uoregon.edu) (recommended!!)
  - OpenNetAdmin [opennetadmin.com](http://opennetadmin.com)
  - HaCi [sourceforge.net/projects/haci](http://sourceforge.net/projects/haci)
  - Racktables [racktables.org](http://racktables.org)
  - IPAT [nethead.de/index.php/ipat](http://nethead.de/index.php/ipat)
  - freeipdb [home.globalcrossing.net/~freeipdb/](http://home.globalcrossing.net/~freeipdb/)
- Examples of IPv6 subnet calculators:
  - ipv6gen [code.google.com/p/ipv6gen/](http://code.google.com/p/ipv6gen/)
  - sipcalc [www.routemeister.net/projects/sipcalc/](http://www.routemeister.net/projects/sipcalc/)

# 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 ISPs 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 generally more expensive (not the “throwaway” consumer devices yet)
  - Cheaper CPE have no IPv6 yet – need to be replaced/ upgraded

# 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)
  - 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 Dialup Customers

---

- Use existing technology:
  - Most dialup access routers are easily upgradable to support IPv6
  - Service looks and feels like the IPv4 service
  - PPPv6 (with DHCPv6 PD perhaps)
  - CPE is usually PC or laptop (and most OSes have supported IPv6 for many years)
  - Service already offered for several years by many ISPs

# 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!)

# IPv6 to Customers

---

- What about addressing? Here is a typical strategy:
  - Mobile Device:
    - /64 = 1 subnet
  - Home/Small Organisation:
    - /56 = 256 subnets
    - Reserve the whole /48
  - Enterprise/Large Organisation:
    - /48 = 65536 subnets

# 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 deployment process
  - Variations on the theme are quite likely – there is no single correct way of proceeding

# IPv6 Deployment Planning



Philip Smith

<philip@nsrc.org>

SANOG 29

Islamabad, Pakistan

23<sup>rd</sup> – 30<sup>th</sup> January 2017