

IPv6 Deployment Planning



Philip Smith

<philip@apnic.net>

APNIC 38

9th – 19th September 2014

Brisbane

Presentation Slides

- Will be available on
 - <http://bgp4all.com/ftp/seminars/APNIC38-IPv6-Deployment-Planning.pdf>
 - And on the APNIC38 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

- ❑ 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
- 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:
 - Audit 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 (1)

- **From your Regional Internet Registry**
 - Become a member of your Regional Internet Registry and get your own allocation
 - Membership usually open to all network operators
 - General allocation policies are outlined in RFC2050
 - 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
 - Multihoming with provider's /48 will be operationally challenging
 - Provider policies, filters, etc

Address Planning

- IPv6 address space available to each network operator is very 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

Why Create an Addressing Plan?

- The options for an IPv4 addressing plan are severely limited:
 - Because of scarcity of addresses
 - Every address block has to be used efficiently
- IPv6 allows for a scalable addressing plan:
 - Security policies are easier to implement
 - Addresses are easier to trace
 - An efficient plan is scalable
 - An efficient plan also enables more efficient network management


Nibble Boundaries

- IPv6 offers network operators more flexibility with addressing plans
 - Network addressing can now be done on nibble boundaries
 - For ease of operation
 - Rather than making maximum use of a very scarce resource
 - With the resulting operational complexity
- A nibble boundary means subdividing address space based on the address numbering
 - Each number in IPv6 represents 4 bits
 - Which means that IPv6 addressing can be done on 4-bit boundaries

Nibble Boundaries – example

- Consider the address block 2001:db8:0:10::/61
 - The range of addresses in this block are:

```
2001:0db8:0000:0010:0000:0000:0000:0000
to
2001:0db8:0000:0017:ffff:ffff:ffff:ffff
```



- Note that this subnet only runs from 0010 to 0017.
- The adjacent block is 2001:db8:0:18::/61


```
2001:0db8:0000:0018:0000:0000:0000:0000
to
2001:0db8:0000:001f:ffff:ffff:ffff:ffff
```

- The address blocks don't use the entire nibble range

Nibble Boundaries – example

- Now consider the address block
2001:db8:0:10::/60
 - The range of addresses in this block are:

2001:0db8:0000:0010:0000:0000:0000:0000
to
2001:0db8:0000:001f:ffff:ffff:ffff:ffff



- Note that this subnet uses the entire nibble range, 0 to f
- Which makes the numbering plan for IPv6 simpler
 - This range can have a particular meaning within the ISP block (for example, infrastructure addressing for a particular PoP)

Addressing Plans – Infrastructure

- ❑ All Network Operators should obtain a /32 from their RIR
- ❑ Address block for router loop-back interfaces
 - Number all loopbacks out of **one** /64
 - /128 per loopback
- ❑ Address block for infrastructure (backbone)
 - /48 allows 65k subnets
 - /48 per region (for the largest multi-national networks)
 - /48 for whole backbone (for the majority of networks)
 - Infrastructure/backbone usually does NOT require regional/geographical addressing
 - Summarise between sites if it makes sense

Addressing Plans – Infrastructure

- 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)
 - Other options:
 - /126s are being used (mimics IPv4 /30)
 - /112s are being used
 - Leaves final 16 bits free for node IDs
 - Some discussion about /80s, /96s and /120s too

Addressing Plans – Infrastructure

□ NOC:

- ISP NOC is “trusted” network and usually considered part of infrastructure /48
 - Contains management and monitoring systems
 - Hosts the network operations staff
 - take the last /60 (allows enough subnets)

□ Critical Services:

- Network Operator’s critical services are part of the “trusted” network and should be considered part of the infrastructure /48
- For example, Anycast DNS, SMTP, POP3/IMAP, etc
 - Take the second /64
 - (some operators use the first /64 instead)

Addressing Plans – ISP to Customer

□ Option One:

- Use ipv6 unnumbered
- Which means no global unicast ipv6 address on the point-to-point link
- Router adopts the specified interface's IPv6 address
 - Router doesn't actually need a global unicast IPv6 address to forward packets

```
interface loopback 0
  ipv6 address 2001:db8::1/128
interface serial 1/0
  ipv6 address unnumbered loopback 0
```

Addressing Plans – ISP to Customer

- Option Two:
 - Use the second /48 for point-to-point links
 - Divide this /48 up between PoPs
 - Example:
 - For 10 PoPs, dividing into 16, gives /52 per PoP
 - Each /52 gives 4096 point-to-point links
 - Adjust to suit!
 - Useful if ISP monitors point-to-point link state for customers
 - Link addresses are **untrusted**, so do not want them in the first /48 used for the backbone &c
 - Aggregate per router or per PoP and carry in iBGP (not ISIS/OSPF)

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
 - This is another very active discussion area
 - Observations:
 - Don't assume that a mobile endsite needs only a /64
 - Some operators are distributing /60s to their smallest customers!!

Addressing Plans – Customer

- Consumer Broadband Example:
 - DHCPv6 pool is a /48
 - DHCPv6 hands out /60 per customer
 - Which allows for 4096 customers per pool
- Business Broadband Example:
 - DHCPv6 pool is a /48
 - DHCPv6 hands out /56 per customer
 - Which allows for 256 customers per pool
 - If BRAS has more than 256 business customers, increase pool to a /47
 - This allows for 512 customers at /56 per customer
 - Increasing pool to /46 allows for 1024 customers
 - BRAS announces entire pool as one block by iBGP

Addressing Plans – Customer

- Business “leased line”:
 - /48 per customer
 - One stop shop, no need for customer to revisit ISP for more addresses until all 65k subnets are used up
- Hosted services:
 - One physical server per vLAN
 - One /64 per vLAN
 - How many vLANs per PoP?
 - /48 reserved for entire hosted servers across backbone
 - Internal sites will be subnets and carried by iBGP

Addressing Plans – Customer

- Geographical delegations to Customers:
 - Network Operator subdivides /32 address block into geographical chunks
 - E.g. into /36s
 - Region 1: 2001:db8:1xxx::/36
 - Region 2: 2001:db8:2xxx::/36
 - Region 3: 2001:db8:3xxx::/36
 - etc
 - Which gives 4096 /48s per region
 - For Operational and Administrative ease
 - Benefits for traffic engineering if Network Operator multihomes in each region

Addressing Plans – Customer

- Sequential delegations to Customers:
 - After carving off address space for network infrastructure, Network Operator simply assigns address space sequentially
 - Eg:
 - Infrastructure: 2001:db8:0::/48
 - Customer P2P: 2001:db8:1::/48
 - Customer 1: 2001:db8:2::/48
 - Customer 2: 2001:db8:3::/48
 - etc
 - Useful when there is no regional subdivision of network and no regional multihoming needs

Addressing Plans – Routing Considerations

- ❑ Carry Broadband pools in iBGP across the backbone
 - Not in OSPF/ISIS
- ❑ Multiple Broadband pools on one BRAS should be aggregated if possible
 - Reduce load on iBGP
- ❑ Aggregating leased line customer address blocks per router or per PoP is undesirable:
 - Interferes with ISP's traffic engineering needs
 - Interferes with ISP's service quality and service guarantees

Addressing Plans – Traffic Engineering


- Smaller providers will be single homed
 - The customer portion of the ISP's IPv6 address block will usually be assigned sequentially
- Larger providers will be multihomed
 - Two, three or more external links from different providers
 - Traffic engineering becomes important
 - Sequential assignments of customer addresses will negatively impact load balancing

Addressing Plans – Traffic Engineering

- ❑ ISP Router loopbacks and backbone point-to-point links make up a small part of total address space
 - And they don't attract traffic, unlike customer address space
- ❑ Links from ISP Aggregation edge to customer router needs one /64
 - Small requirements compared with total address space
 - Some ISPs use IPv6 unnumbered
- ❑ Planning customer assignments is a very important part of multihoming
 - Traffic engineering involves subdividing aggregate into pieces until load balancing works

Unplanned IP addressing


- ISP fills up customer IP addressing from one end of the range:



- Customers generate traffic
 - Dividing the range into two pieces will result in one /33 with all the customers and the ISP infrastructure the addresses, and one /33 with nothing
 - No loadbalancing as all traffic will come in the first /33
 - Means further subdivision of the first /33 = harder work

Planned IP addressing

- If ISP fills up customer addressing from both ends of the range:



- Scheme then is:
 - First customer from first /33, second customer from second /33, third from first /33, etc
- This works also for residential versus commercial customers:
 - Residential from first /33
 - Commercial from second /33

Planned IP Addressing


- ❑ This works fine for multihoming between two upstream links (same or different providers)
- ❑ Can also subdivide address space to suit more than two upstreams
 - Follow a similar scheme for populating each portion of the address space
- ❑ Consider regional (geographical) distribution of customer delegated address space
- ❑ Don't forget to always announce an aggregate out of each link

Addressing Plans – Advice


- ❑ Customer address assignments should not be reserved or assigned on a per PoP basis
 - Follow same principle as for IPv4
 - Subnet aggregate to cater for multihoming needs
 - Consider regional delegation
 - ISP iBGP carries customer nets
 - Aggregation within the iBGP not required and usually not desirable
 - Aggregation in eBGP is very necessary
- ❑ Backbone infrastructure assignments:
 - Number out of a **single** /48
 - ❑ Operational simplicity and security
 - Aggregate to minimise size of the IGP

Addressing Plans – Scheme

□ Looking at Infrastructure:



□ Alternative:



Addressing Plans

Planning

- Registries will usually allocate the next block to be contiguous with the first allocation
 - (RIRs use a sparse allocation strategy – industry goal is aggregation)
 - Minimum allocation is /32
 - Very likely that subsequent allocation will make this up to a /31 or larger (/28)
 - So plan accordingly

Addressing Plans (contd)

- Document infrastructure allocation
 - Eases operation, debugging and management
- Document customer allocation
 - Customers get /48 each
 - Prefix contained in iBGP
 - Eases operation, debugging and management
 - Submit network object to RIR Database

Addressing Tools

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

Deploying IPv6



Now we put it onto the network

Deploying addressing and IGP

- Strategy needed:
 - Start at core and work out?
 - Start at edges and work in?
 - Does it matter?
- Only strategy needed:
 - Don't miss out any PoPs
 - Connectivity is by IPv4, so sequence shouldn't matter
 - Starting at core means addressing of point to point links is done from core to edge (many ISPs use strategy of low number towards core, high number towards edge)
 - But it really doesn't matter where you start...

IPv6 Deployment

- ❑ Number all the infrastructure interfaces according to the established addressing plan
 - No customers yet
- ❑ Care needed on LANs
- ❑ Secure routers and L3 devices for IPv6 access
 - Once a device is enabled for IPv6, it must have all the same security policies applied as for IPv4

Deploying on PoP LANs

- LANs need special treatment
 - Even those that are only point to point links
- Issues:
 - ISPs don't want to have Router Advertisements active on network infrastructure LANs
 - Activating IPv6 on a LAN which isn't adequately protected may have security consequences
 - Servers may auto configure IPv6
 - No firewall filtering means no security ⇒ compromise

IPv6 Interior Routing Protocols

- ❑ Make a decision about which IGP to use
 - (continue with OSPF vs deploy ISIS?)
- ❑ Enable chosen IPv6 IGP
 - Care needed not to break IPv4 connectivity
 - Adjacencies in IPv6 should match existing adjacencies in IPv4
 - IGP v6 routing table should match v4 routing table
- ❑ Check that the IPv6 network's operation compares with IPv4 operation
 - Fix any problems
 - In a dual stack network the protocols must function the same way

IPv6 Routing Protocol Deployment

- Enable IPv6 BGP
 - iBGP – should replicate IPv4 iBGP
 - Same number of active neighbours
 - IPv6 version of the IPv4 configuration
 - Modify existing templates
 - eBGP comes next
- Check that the IPv6 network's 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

- 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
- Operators 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

Managing and Monitoring the Network



Watching the Infrastructure...

Managing and Monitoring the Network

- Existing IPv4 monitoring systems should not be discarded
 - IPv4 is not going away yet
- How to Monitor IPv6?
 - Netflow
 - MRTG
 - Syslog
 - Commercial systems?
 - Others?

Netflow for IPv6

- ❑ Public domain flow analysis tool NFSEN (and NFDUMP) support Netflow v5, v7 and v9 flow records
 - IPv6 uses v9 Netflow
 - NFSEN tools can be used to display and monitor IPv6 traffic
 - More information:
 - ❑ <http://nfdump.sourceforge.net/>
 - ❑ <http://nfsen.sourceforge.net/>
- ❑ ISPs using existing IPv4 netflow monitoring using NFSEN can easily extend this to include IPv6

MRTG

- ❑ MRTG is widely used to monitor interface status and loads on SP infrastructure routers and switches
- ❑ Dual stack interface will result in MRTG reporting the combined IPv4 and IPv6 traffic statistics
- ❑ MRTG can use IPv6 transport (disabled by default) to access network devices

Other Management Features

- A dual stack network means:
 - Management of the network infrastructure can be done using either IPv4 or IPv6 or both
 - ISPs recognise the latter is of significant value
- If IPv4 network breaks (e.g. routing, filters, device access), network devices may well be accessible over IPv6
 - Partial “out of band” network
- IPv6 is preferred over IPv4 (by design) if AAAA and A records exist for the device
 - So remote logins to network infrastructure will use IPv6 first if AAAA record provided

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
 - Mobile
 - Cable/ADSL
 - Dial
 - Leased lines
 - Wireless Broadband

IPv6 to Mobile Customers

- ❑ Access technologies include 3G/LTE, Wifi (802.11) and WiMax
- ❑ End-sites could range from handsets to major corporations
- ❑ Strategy depends on infrastructure and device capability:
 - Dual-stack
 - IPv4-only with NAT46
 - IPv6-only with NAT64

IPv6 to Mobile Customers

- Dual-stack:
 - Most probably IPv4-NAT and native IPv6
 - Handset / device / infrastructure support?
- IPv4-only with NAT46:
 - Availability of IPv4 to IPv6 protocol translators?
 - Are there IPv6-only sites as yet?
- IPv6-only with NAT64:
 - Deployment of CGN
 - Handset / device / infrastructure support?

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 v6 + 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 Handset:
 - /64 = 1 subnet
 - Home/Small Organisation:
 - /60 = 16 subnets
 - Reserve the whole /56
 - Reserve a /48 for small orgs = 256 small orgs per /48
 - Medium Organisation:
 - /56 = 256 subnets
 - Reserve the whole /48
 - 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?
 - What needs to be available on IPv6?
 - 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