IPv6 Deployment Planning

ISP Workshops

Last updated 27th July 2018

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
- Use of these materials is encouraged as long as the source is fully acknowledged and this notice remains in place
- Bug fixes and improvements are welcomed
  - Please email *workshop (at) bgp4all.com*

Philip Smith

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

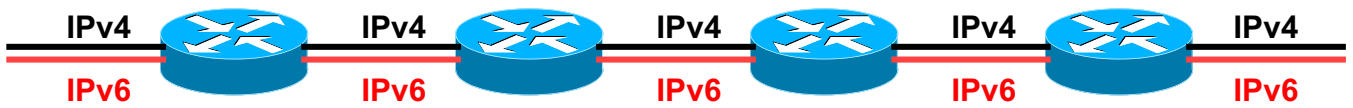

MPLS Backbone

- 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

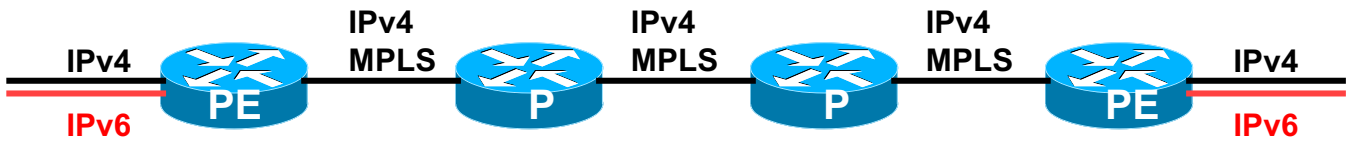

Network Assessment

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

Assessment

- 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/) 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…

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
    - 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
  - 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

- Industry standard for customer assignments today:
  - /64 for just one LAN (hosted server)
  - /64 for mobile handset (3GPP Release 9 & earlier)
  - /60 for mobile handset tethering (3GPP Release 10 onwards)
  - /56 for a small network (home user / small business)
  - /48 for a large network (enterprise)

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 is not 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 replace with IS-IS?)
- 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

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

IPv6 to Mobile Customers

- Access technologies are predominantly 3G and LTE
- End-sites could range from handsets to major corporations
- Strategy depends on infrastructure and device capability:
  - Dual-stack using 464XLAT – Android
  - IPv6-only with NAT64 – Apple iOS
  - Dual-stack – reported to be in Apple iOS 11.3 onwards
  - Mobile operators need to support both popular consumer devices

IPv6 to Mobile Customers (1)

- Dual-stack:
  - Native IPv4 (private addresses) and IPv6 to handset
  - Infrastructure is dual-stack
  - Handsets support:
    - All Android phones
    - Apple iOS phones from iOS 11.3 (reported by several mobile operators)
  - Operator needs CGNAT to handle IPv4 NAT needs
  - Notify Google and Apple for carrier update to be pushed to handsets

- Tethering
  - Bridging (/64 for handset is shared with tethered devices)
  - 3GPP Release 10 adds DHCP-PD support

IPv6 to Mobile Customers (2)

- Dual-stack:
  - Native IPv6 and IPv4-NAT
    - IPv6 native from handset to content
    - IPv4 is carried within IPv6
  - Infrastructure is IPv6 only
  - Handsets support 464XLAT (CLAT)
    - Most Android phones (4.4.4 and 5.1 onwards)
  - Operator needs CGNAT to handle PLAT function for handset access to IPv4 legacy sites

- Tethering
  - Bridging (/64 for handset is shared with tethered devices)
  - 3GPP Release 10 adds DHCP-PD support

IPv6 to Mobile Customers (3)

- IPv6-only with NAT64:
  - Native IPv6 only
  - Infrastructure is IPv6 only
  - Handsets are IPv6-only
    - Apple iPhone (iOS 9 onwards, iPhone6S onwards)
  - Operator needs CGNAT to handle NAT64 function for handset access to IPv4 legacy sites
- Tethering
  - Bridging (/64 for handset is shared with tethered devices)
  - 3GPP Release 10 adds DHCP-PD 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
  - IPv4: PPPoE (IPCP)
  - IPv6: PPPoE (IPv6CP) + DHCPv6 PD
  - Used by ISPs such as Internode (AU) and XS4ALL (NL)
- Issues:
  - IPv6 CPE in some markets are generally more expensive
    - Customised "country versions" often delete IPv6 support
  - Cheapest CPE have no IPv6 – need to be replaced/upgraded
    - Customers are very cost conscious – $1 can sway a purchasing decision
- Solution:
  - Operator publishes recommended list of CPE (which support dual-stack)!
    - And sample configurations

IPv6 to Broadband Customers

- Method 2: use 6rd
  - 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 IPCP for native IPv4 link
    - IPv6 address 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 need to be replaced/upgraded to support 6rd
- Solution:
  - Operator publishes recommended list of CPE (which support 6rd)!
    - And sample configurations!

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
  - PPP with IPv6CP (or with DHCPv6 PD (??))
  - 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 (tethered devices are bridged on to /64)
    - /60 = 16 subnets for tethering (DHCP-PD with 3GPP release 10)
  - Home/Small Organisation:
    - /56 = 256 subnets
    - Reserve the whole /48
    - There is no IPv6 address shortage!
  - 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?

Customer End-Site

- Re-run of this presentation, but:
  - Do all devices need IPv6?
  - Realistically, IPv6 needed on:
    - End-user devices (handset, tablet, laptop, desktop) need IPv6
      - Already turned on by default, remember ☺
    - External facing servers need IPv6 initially (website, mail relay, public DNS)
    - Corporate Firewalls, Routers and IDS
  - Other internal systems have no urgent need to deploy:
    - Internal facing servers
    - IP phone systems & Printers
    - Management access to IP enabled devices
    - Network and Building security monitoring systems

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

ISP Workshops