IPv6 Deployment Study

Philip Smith

<pfs@cisco.com>

MENOG 5, Beirut, 29th October 2009

Agenda (& Strategy)

- Network Audit & Optimisation
- Procuring IPv6 Address Space & Creating an Address Plan
- Deploying Addressing & IGP
- Deploying iBGP
- Seeking IPv6 Transit
- Forward and Reverse DNS
- Services & Customers

Network Audit & Optimisation

What can run IPv6 today, what needs to be upgraded, and what do we clean up?

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)
- Analyse each PoP:
  - Documenting router hardware & software specification
  - RANCID (www.shrubbery.net/rancid/) makes this very easy
- Sanity Check Configuration
  - Tidy up and remove unused configurations

Audit: Results

- Upgrade RAM and FLASH for platforms identified as being deficient
- Replace routers which can not run most recent IOS software (12.2S, 12.3 and 12.4)

  This will impact 2600 (non-XM), 3620, elderly 7200s (pre NPE200), &c
- Decide on a software strategy

  Mix of 12.3 and 12.4 -or-

  12.4 everywhere (bigger impact as some platforms which support 12.3 aren't supported for 12.4 - e.g. 2500)

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

Procuring IPv6 address space & Constructing a Deployable address plan

Now we need addresses…

Getting IPv6 address space (RIR)

- If existing Regional Internet Registry account holder with an IPv4 allocation:
  - Just ask for an IPv6 allocation – it really is as simple as that!
- If not an RIR account holder, become one and get your own IPv6 allocation
  - Requires a plan for a year ahead
  - IPv6 allocation policies are documented on each RIR website

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

- 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
  - Summarise between sites if it makes sense
- Customers get **one** /48
  - Unless they have more than 65k subnets in which case they get a second /48 (and so on)

Addressing Plans

- What about Infrastructure 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

Deployable Address Plan

- Documentation
  - IPv4 addresses are probably short enough to memorise
  - IPv6 addresses are unlikely to be memorable at all
- Document the address plan
  - What is used for infrastructure
  - What goes to customers
  - Flat file, spreadsheet, database, etc
  - But documentation is vital
  - Especially when coming to populating the DNS later on

Deployable Address Plan

- Pick the first /48 for our ISP infrastructure
  - Reason: keeps the numbers short
  - Short numbers: less chance of transcription errors
  - Compare:
    - **2001:db8:ef01:d35c::1/128**
      with
    - **2001:db8::1/128**
    - For Loopback interface addresses
- Out of this /48, pick the first /64 for loopbacks
  - Reason: keeps the numbers short

Deployable Address Plan

- For the infrastructure /48:
  - First /64 for loopbacks
  - Remaining 65535 /64s used for internal point to point links
- Second /48:
  - Use for point to point links to customers
  - Unless you use unnumbered interfaces
  - That gives 65536 /64s for 65536 customer links
- Remaining /48s are for delegation to customers

Deploying Addressing and IGP

Let's now touch 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…

Deploying: Router1 in PoP1

- Start with addressing

  Address all the PtP links on Router1

```
interface serial 0/0
  ipv6 address 2001:db8:0:110::1/64
interface hssi 1/0
  ipv6 address 2001:db8:0:111::1/64
```

  Go to the other end of each PtP link and apply the corresponding addressing there also

```
interface serial 2/0/0
  ipv6 address 2001:db8:0:110::2/64
```

  ...and…

```
interface hssi 3/1
  ipv6 address 2001:db8:0:111::2/64
```

Deploying the IGP

- Configure IGP on the links that will run an IGP

```
ipv6 router ospf 100
  log adjacency-changes detailed
  passive-interface default
  no passive-interface serial 0/0
  no passive-interface hssi 1/0
interface serial 0/0
  ipv6 ospf 100 area 0
interface hssi 1/0
  ipv6 ospf 100 area 0
```

- No need to do the OSPF on the other end yet

  Those routers will be done in due course, and saves time jumping back and forth

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

Deploying on PoP LANs

- Example of Point to Point link:

```
interface GigabitEthernet0/0
  description Crossover Link to CR2
  ipv6 address 2001:db8:0:115::1/64
  ipv6 nd suppress-ra
  ipv6 ospf 100 area 0
  ipv6 ospf network point-to-point
```

Deploying on LANs

- Example of local services LAN:

```
interface GigabitEthernet0/1
  description Services LAN
  ipv6 address 2001:db8:0:101::1/64
  ipv6 nd suppress-ra
  ipv6 traffic-filter SERVER-IN in
  ipv6 traffic-filter SERVER-OUT out
```

Where the **`server-in`** and **`server-out`** filters are ipv6 access-lists configured to:

allow minimal access to servers (only ssh for now), **-or-**

to match their IPv4 equivalents

Checks

- Before launching into BGP configuration
  - Sanity check the OSPFv3 configuration
- Are all adjacencies active?
  - Each router should have the same number of OSPFv2 and OSPFv3 adjacencies
- Does each interface with an OSPFv2 configuration have a corresponding OSPFv3 configuration?
- Have interfaces not being used for OSPFv3 been marked as passive?
  - And do they match those marked as passive for OSPFv2?

Checks

- Does the number of entries in the OSPFv3 routing table match the number of entries in the OSPFv2 routing table
  - Compare the number of entries in “sh ip route ospf” and “sh ipv6 route ospf”
  - Examine differences and work out the reason why
- Do IPv4 and IPv6 traceroutes through the network
  - Are the paths the same?
  - Are the RTTs the same?
  - Discrepancies must be investigated and fixed

Deploying iBGP

Functioning IGP means all routers reachable…

Deploying iBGP

- Strategy is required here
  - Starting at edge makes little sense
  - Starting at core means route reflector mesh builds naturally
- Modify BGP default assumptions
- Prepare templates
  - Set up peer-groups in master configuration file
  - There should already be a master configuration for IPv4

Creating IPv6 templates

- Typical iBGP peer-groups might be:

core-ibgp	router participates in full mesh iBGP
rr-client	neighbour is a client of this route reflector
rr	neighbour is a route reflector

- These should be replicated for IPv6:

corev6-ibgp	router participates in full mesh iBGP
rrv6-client	neighbour is a client of this route reflector
rrv6	neighbour is a route reflector

  Keep the names the same - just add “v6” in the appropriate place to differentiate

- Peer-groups are to be created within the appropriate address family

Next Steps

- Load all these templates into the routers across the backbone
  - Or simply upload them as each router has IPv6 iBGP deployed on it
- Originate the IPv6 address block on the chosen core routers within the backbone
  - Make sure there is more than one, and the prefix is originated in more than one PoP (for redundancy)
  - BGP network statement and matching static route to Null0 - same as for IPv4

Checks

- Are all the iBGP peers up?
  - Best to check on each route reflector
  - If peerings are still down investigate reasons - usually because a loopback address is missing from OSPFv3
- Are there the same number of IPv6 peers as there are IPv4 peers?
  - If not, what went wrong?
- Prefixes in iBGP
  - There probably will be none apart from the /32 aggregate block and any static LANs which have been introduced into iBGP

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
- 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

Connecting over IPv6 and fixing those traceroutes…

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

Forward DNS

- Operators typically access the router by connecting to loopback interface address
  - Saves having to remember interface addresses or names – and these change anyway
- Setting up the IPv6 entries means adding a quad-A record beside each A record:

```
cr1.pop1      A      192.168.1.1
              AAAA   2001:db8::1:1
cr2.pop1      A      192.168.1.2
              AAAA   2001:db8::1:2
gw1.pop1      A      192.168.1.3
              AAAA   2001:db8::1:10
```

Forward DNS

- Completing the infrastructure zone file as per the example is sufficient
  - Update the SOA record
  - Reload the nameserver software
  - All set
- If connecting from an IPv6 enabled client
  - IPv6 transport will be chosen before the IPv4 transport
  - (Part of the transition process from IPv4 to IPv6)
  - For all connections to IPv6 enabled devices which have entries in the forward DNS zones
  - This could have positive as well as negative consequences!

Reverse DNS

- First step is to have the /32 address block delegated by the RIR
- Prepare the local nameservers to handle the reverse zone, for example in BIND:

```
zone ”8.b.d.0.1.0.0.2.ip6.arpa" in {
        type master;
        file "ip6.arpa-zones/db.2001.0db8;
        allow-transfer {"External"; "NOC-NET";};
};
```

- And then “create and populate the zone file”

Services

Network is done, now let's use it…!

Infrastructure complete

- This was the easy part
  - Network infrastructure generally is very simply to set up as dual stack IPv4 and IPv6
- The next steps are more complex
- Services?
  - Which to make available in IPv6 too?
- Customers?
  - Which can be offered services, and how?

ISP Services

- DNS, Mail, Web
  - Critical customer and Internet facing servers
  - Simple to transition to dual stack
- This involves:
  - Setting up appropriate IPv6 filters on hosting LANs (hint: replicate IPv4 filters)
  - Giving the servers IPv6 addresses
  - Ensuring that the server software is listening on both IPv4 and IPv6 ports
  - Publishing quad-A records along side the regular A records
  - Testing!

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
- Beyond the scope and intention of this presentation
  - But at least now there is no excuse not to explore the above some more

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

Was this hard?

Conclusion

- Was this really all that hard?
- Strategy drawn up & executed
- Did not take too many “man hours”