IPv6 Addressing

AfNOG 2012 AR-E Workshop

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)

- 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

Using 6to4 for IPv6 address space

- Some entities use 6to4
  - Not recommended due to operational problems
  - Read http://datatracker.ietf.org/doc/draft-ietf-v6ops-6to4-to-historic
- FYI: 6to4 operation:
  - Take a single public IPv4 /32 address
  - 2002:<ipv4 /32 address>::/48 becomes your IPv6 address block, giving 65k subnets
  - Requires a 6to4 gateway
  - 6to4 is a means of connecting IPv6 islands across the IPv4 Internet

Addressing Plans – ISP Infrastructure

- ISPs should receive /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)
  - Summarise between sites if it makes sense

Addressing Plans – ISP 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 – 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 give small customers a /56 or a /60 and single LAN end-sites a /64, e.g.:

    /64 if end-site will only ever be a LAN
    /60 for small end-sites (e.g. consumer/broadband)
    /56 for medium end-sites (e.g. small business)
    /48 for large end-sites
  - (This is another very active discussion area)

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":
  - /56 per customer
  - Reserve the /48 – allows for growth of customer network
- 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 – Miscellaneous

- NOC:
  - ISP NOC is "trusted" network and usually considered part of infrastructure /48
    - Contains management and monitoring systems
    - Hosts the network operations staff
- Infrastructure Suggestions:
  - Loopbacks:
    - take the first /64
  - NOC & Management systems:
    - take the last /60 (allows enough subnets)
  - Backbone point to point links:
    - Occupy the remaining /48 space

Addressing Plans – ISP to Customer

- ISP to Customer links
  - Use ipv6 unnumbered
    - Which means no ipv6 address
    - Router adopts the specified interface's IPv6 address
      - Router doesn't actually need an IPv6 address to forward packets

Or

  - Use the second /48 for point to point links
    - 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 and carry in iBGP (not ISIS/OSPF)

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 singled 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:

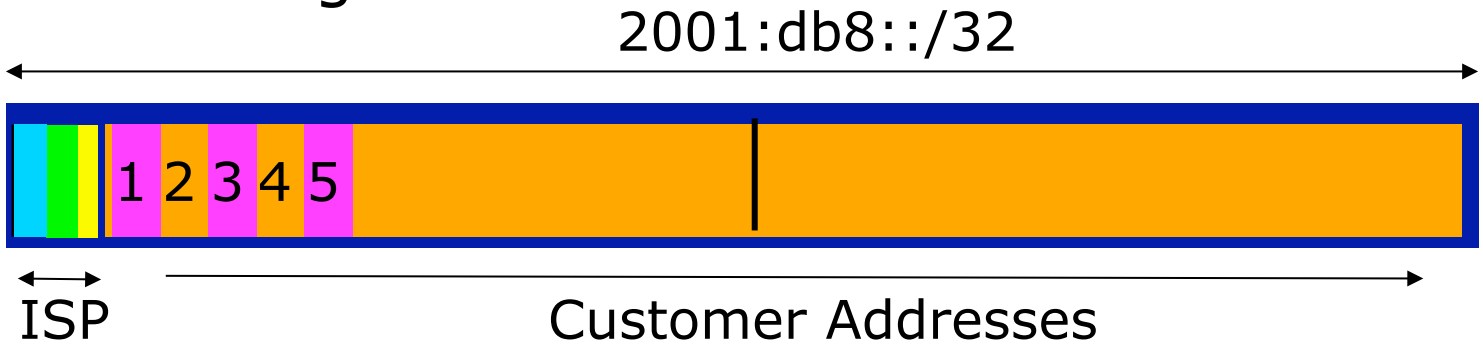

- Customers generate traffic
  - Dividing the range into two pieces will result in one /33 with all the customers and the ISP infrastructure the addresses, and one /33 with nothing
  - No loadbalancing as all traffic will come in the first /33
  - Means further subdivision of the first /33 = harder work

Planned IP addressing

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

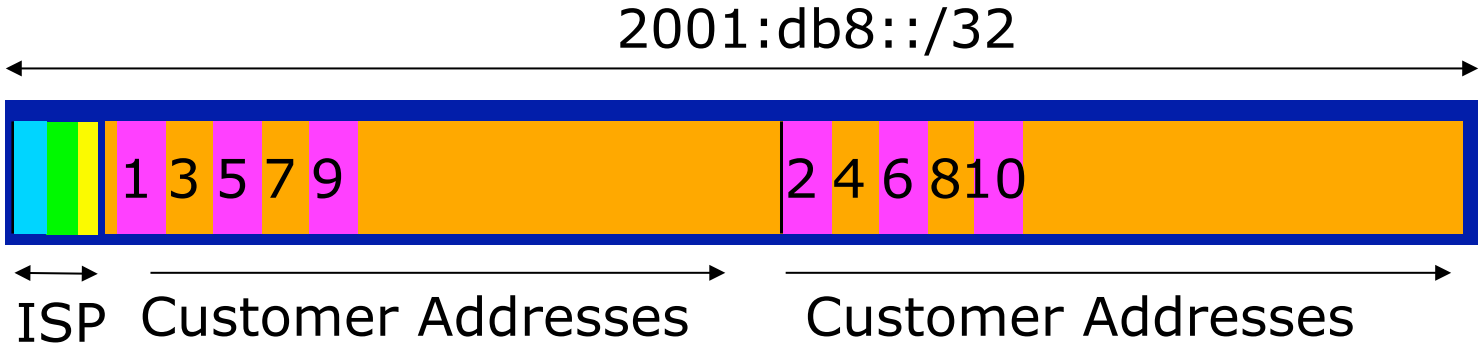

- Scheme then is:
  - First customer from first /33, second customer from second /33, third from first /33, etc
- This works also for residential versus commercial customers:
  - Residential from first /33
  - Commercial from second /33

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
- 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
  - 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:

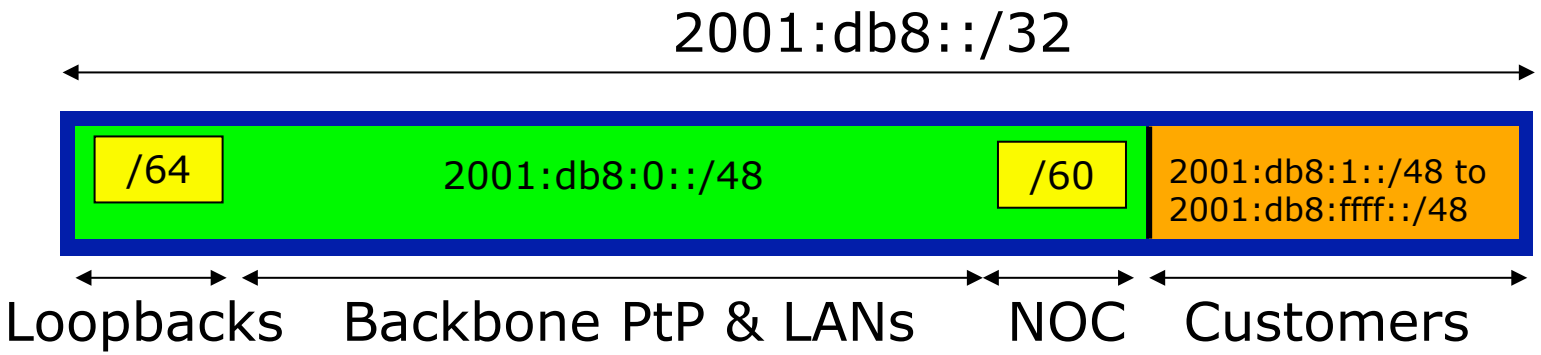

- Alternative:

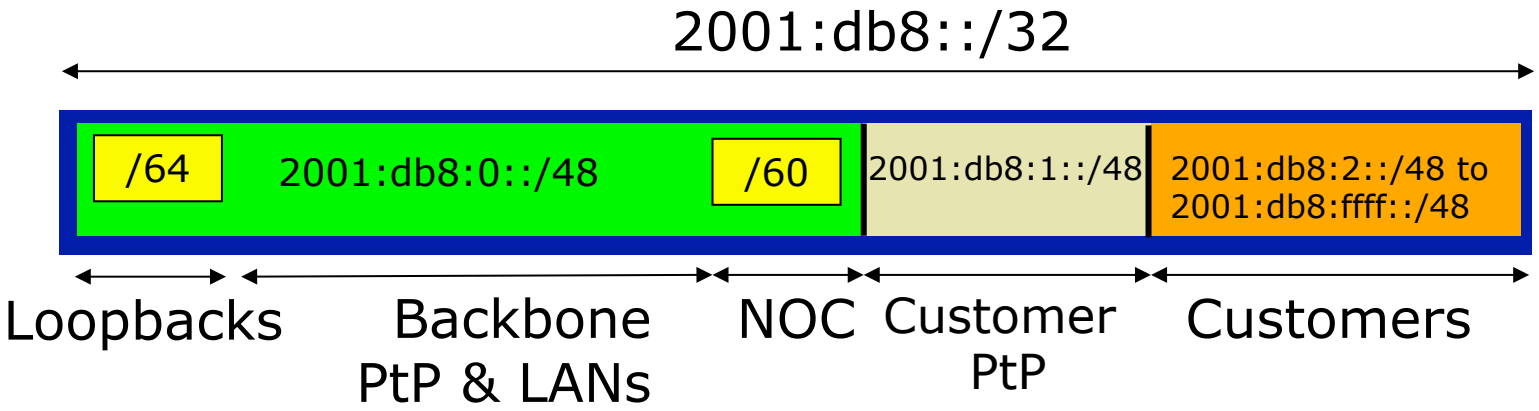

Addressing Plans Planning

- 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
  - 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 tools (which support IPv6 too):
  - NetDot — netdot.uoregon.edu
  - HaCi — sourceforge.net/projects/haci
  - IPAT — nethead.de/index.php/ipat
  - ipv6gen — techie.devnull.cz/ipv6/ipv6gen/
  - sipcalc — www.routemeister.net/projects/sipcalc/
  - freeipdb — home.globalcrossing.net/~freeipdb/

IPv6 Addressing

AfNOG 2012 AR-E Workshop