IPv6 Addressing

ISP Workshops

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 policy 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 still use 6to4
  - Not recommended due to operational problems
  - Read http://datatracker.ietf.org/doc/draft-ietf-v6ops-6to4-to-historic for some of the reasoning why
- 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

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 subnetting address space based on the address numbering
  - Each number in IPv6 represents 4 bits = 1 nibble
  - 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:

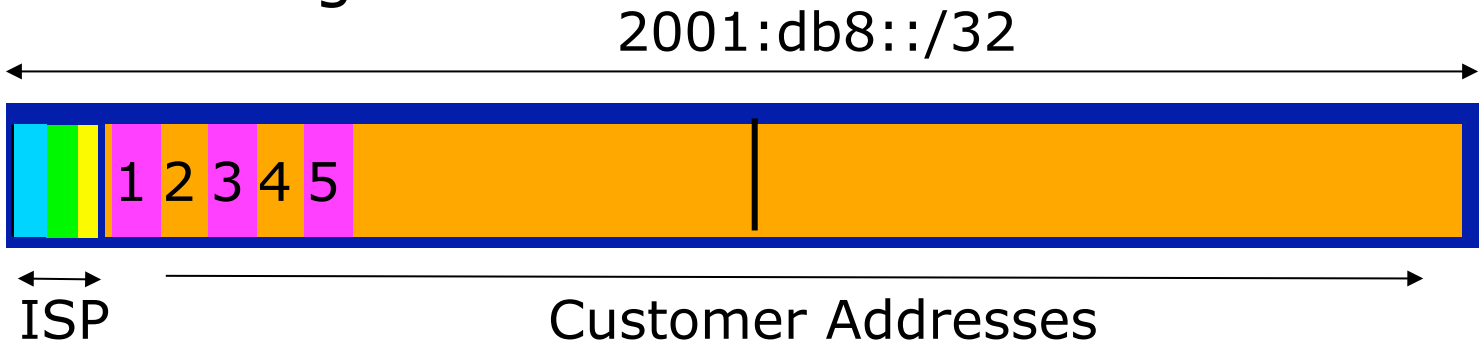

- Customers generate traffic
  - Dividing the range into two pieces will result in one /33 with all the customers and the ISP infrastructure the addresses, and one /33 with nothing
  - No loadbalancing as all traffic will come in the first /33
  - Means further subdivision of the first /33 = harder work

Planned IP addressing

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

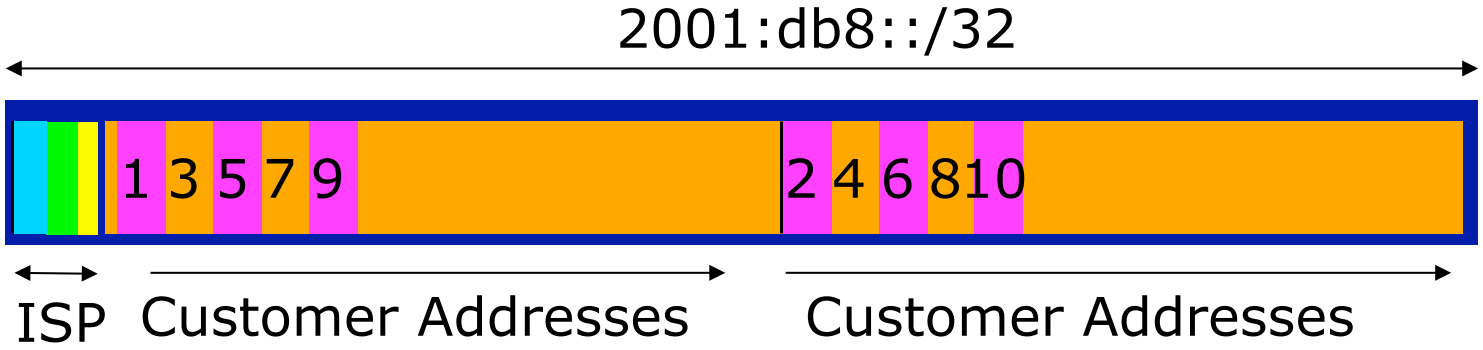

- Scheme then is:
  - First customer from first /33, second customer from second /33, third from first /33, etc
- This works also for residential versus commercial customers:
  - Residential from first /33
  - Commercial from second /33

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:

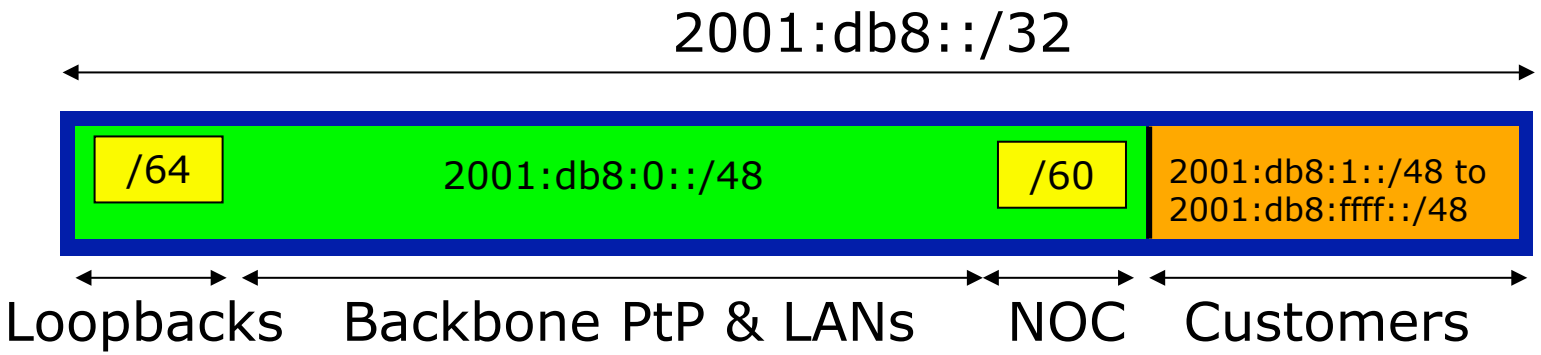

- Alternative:

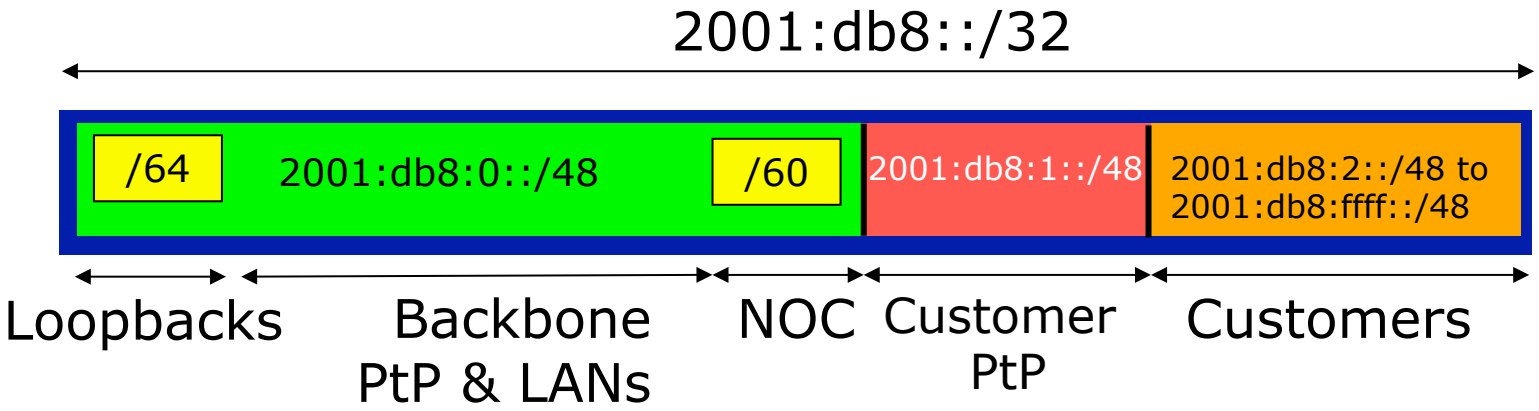

Addressing Plans Planning

- 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
  - Racktables racktables.org
  - 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/

IPv6 Addressing

ISP Workshops