

addressing planning on ISP - IJ/AS2497 case -

Matsuzaki 'maz' Yoshinobu
<maz@ij.ad.jp>

uncertainness on IPv6 addressing

- variations from IPv4
 - more bits
 - more use of linklocal
- concerns
 - internal routing
- assuming smarter ways if any 😊
 - management
 - policy

IIJ case

- enterprise
 - leased line, data center and etc
- consumer
 - FTTH, ADSL and etc

prefix length

- /128 for loopback interfaces
- /64 for links
 - /127 is also used on inter-router links
- /48 for enterprise
- /56 for consumer
 - changed from /48
- /64 for mobile

global unicast address

- An inter-router link does not require global address inside an AS
 - OSPFv3 uses link-local address to exchange LSAs
 - A loopback interface needs a global address to configure iBGP sessions
- But we configure global address on every interface
 - as ping destination to check availability

linklocal address

- fe80::/64
- AS IS
 - We don't change
 - Most routers use Modified EUI-64 format address
- A virtual address for vrrp/hsrp is another story.
 - Customers might configure a static route at their equipments with this address, so the address should be assigned statically like fe80::1.

route aggregation inside AS

- not so aggressive at IJ
 - similar to IPv4 case
 - does aggregate where it's easily possible
 - consumer, mobile services and so on
- BGP can handle many prefixes
 - BGP carries most prefixes including customers'
 - IGP is just for topology and minimum routes
- of course, we announce aggregated prefixes only to other ASes.

thought

- infrastructure protection
 - access-control and/or policing
- it's better to reserve a block based on purpose to simplify rules
 - inter router links and loopbacks inside AS
 - others
 - inter router links with other parties
 - users, servers

reserving policy

- /40

```
2001:0db8:0000::/40 -- servers and etc 1
2001:0db8:0100::/40 -- infrastructure/backbone
2001:0db8:0200::/40 -- consumer 1
2001:0db8:0300::/40 -- consumer 2
2001:0db8:0400::/40 -- mobile 1
2001:0db8:0500::/40 -- leased line 1
      :
2001:0db8:ff00::/40  --
```

management

- /40 reservation by hand
 - plain text
- infrastructure
 - reserve /56 blocks based on purpose from the /40
- in-house tool to help assignment
 - new assignment
 - return
 - tunnel service bound to IPv4 connectivity services

IJ office network

- /48 was assigned for that
 - split it into 2 /49s - outside/inside firewall
 - this could be wrong, as we don't need that much outside. reconsideration will be needed.
 - reserve /56 from the block based on purpose
 - servers
 - clients
 - and assign /64 on each link

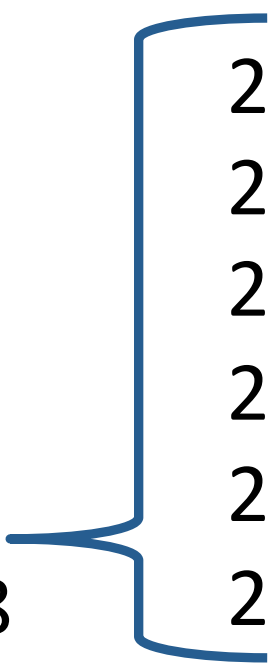
lesson learned #1

- We can estimate needs based on IPv4 experience
 - # of customers, links and server segments
- you need your 'reservation boundary'
 - as an enterprise, we use /56
 - as an ISP, we use /40
 - and break it into /56s if needed

lesson learned #2

- don't forget about 'a'
 - it's HEX

- 2001:db8:1:7::/48
- 2001:db8:1:8::/48
- 2001:db8:1:9::/48
- 2001:db8:1:10::/48



2001:db8:1:a::/48
2001:db8:1:b::/48
2001:db8:1:c::/48
2001:db8:1:d::/48
2001:db8:1:e::/48
2001:db8:1:f::/48

summary

- we don't care about internal routing much
 - BGP can handle it well in our case
- we care
 - infrastructure block
 - human friendly boundary
 - /40, /48, /56 and /64