

Best Current Operational Practice for operators:

IPv6 Prefix Assignment for end-customers – static, dynamic (non-stable) or stable

Jordi Palet

jordi.palet@consulintel.es

Authors:

- Andrew Alston
- Gert Doering
- Jan Žorž
- Jen Linkova
- Jordi Palet
- Kevin Meynell
- Luis Balbinot
- Mark Townsley
- Primož Dražumerič
- Sander Steffann

RIPE BCOP TF

- Work in Progress:
 - <https://www.ripe.net/ripe/mail/archives/bcop/2017-March/000159.html>
 - <https://www.sinog.si/docs/draft-IPv6pd-BCOP-v1.pdf>

BCOP and Why?

- Describe best **actual** practices
- Target: ISPs deploying IPv6
- Lack of experience or following IPv4 practices bring unexpected or unwanted results
 - IPv6 “brokenness” = Content providers rejection of your AS
 - Lack of compliance with new standards such as Homenet
 - Complete **production** network renumbering, etc.

Size of end-customer prefix

- /48, /56 or something else?
- Change your mind, this is not IPv4!
- IPv6 has been designed to assign prefixes not addresses
- Tony Hain maths:
 - IPv6 lifetime over 480 years, and keep doing that several times
 - Scarcity of addresses is not going to be our next problem

Numbering the WAN link

1. /64 out of the end-customer prefix
2. /64 out of a dedicated pool
3. Unnumbered
4. ULA

/64 from customer prefix

- Use the 1st /64 from the customer prefix
 - <https://tools.ietf.org/html/draft-palet-v6ops-point2point>
 - Simplifies routing and provisioning
- Some CPEs may not support RFC6603
 - Prefix exclude option for DHCPv6-PD
- Even being required by RFC7084
 - Basic Requirements for IPv6 CPEs

/64 from dedicated pool

- Most common scenario
 - Dedicated pool for WAN links
- CPE performs router discovery
 - If it is a host (PPPoE), setup is completed
 - If it is a router, will request a prefix (DHCPv6-PD)
- /126, /127, /112 or /64?
 - RFC6164 suggest /127
 - Not all hardware supports it
 - /64 is future proof
 - Hardware limitations for longer than /64 prefixes
 - Allocate /64, use /127 to prevent ND attacks
- If there is **always** a CPE, you can apply security policies w/o harming customers

Unnumbered

- Don't use GUAs
 - Instead use Link-Local
- Doesn't work for all the devices, which can't request DHCPv6-PD
 - No GUAs means no traffic ...
- Complicate troubleshooting
 - Not able to traceroute the point of failure
- Not suitable for unknown CPEs or non-CPEs attached to the WAN link
- End-host will stay unnumbered
- Some hardware may consume additional resources for numbered links

ULA

- Strongly discouraged
- ICMPv6 from the CPE to outside ISP
 - ULA source address will not traverse filters
 - PMTUD will break
 - IPv6 connection will break if Path MTU is not the same

WAN link summary

- /64 GUA is the recommended choice
 - From the customer prefix if RFC6603 is supported
- It may be even required when more than 2 endpoints
 - Managed bridges
 - Repeaters
 - Redundancy (VRRP, multiple routers)
 - Monitoring/troubleshooting devices

Prefix assignment options

- Align the size of the delegated prefix with a nibble boundary (multiples of 4 bits), so it match DNS reverse zone delegations
- A single customer network is /64
 - A single /64 is plain wrong
 - IETF work allows a single /64 for an interface
- Multiple /64 must be the rule
 - RIR policies allow /48

/48 for business, /56 residential

- Some operators do this
 - Rationale -> Marketing/Sales differentiation
- Advanced home users may have problems with this
 - You're not able to use all the 4 digits (/48-/56)
- Some may have already an addressing plan with /48 (ULA, TB, transition, etc.)
 - /56 forces to redo it + renumbering
 - /48 just means changing the prefix
- Alternatively, reserve /48, assign /56
- Are you considering SMEs?

/48 for everybody

- Most practical and pragmatic
- Less call-centre time to sort out problems
- Single “flat” provisioning system
- Same prefix size as ULAs, transition, etc.
 - Direct mapping of existing addressing plans

Less than /56

- Not recommended
 - Technically no reason for that, enough addresses, this is not IPv4!
 - Over 134 million /56 in a /29
 - Over 16 million /56 in a /32
- Ask for more space to your RIR if required
- Never assign a single /64
 - Except for cellular phones (1 /64 for each PDP)
- LTE modems still require /56 or /48

Static, dynamic (non-stable) or stable

- Static typically by means of AAA or custom provisioning system
 - At customer connection they always get the same prefix
- Dynamic by means of a big pool in each termination point
 - At customer connection they get a random prefix
 - If stable, the lease time may provide days, weeks or even months

Dynamic (non-stable) is easier?

- Less effort to deploy
 - Issues come later
 - It comes from IPv4 practices, DHCP
 - But we have NAT!
 - Looks easier for aggregation
 - Not looking for “customer” portability
 - May be an extra service
- Commonly using DHCPv6-PD
 - Each end-customer device has a GUA

However ... dynamic is harmful

- In case of power failure, CPE hang-up, ...
 - Common even in highly-developed countries
- CPE doesn't send prefix valid lifetime = 0
 - End-customer devices keep the old prefix
 - Will try to use it, will fail
 - Customers claims to the call-centre
- Content providers measure IPv6 brokenness
 - Will ignore your IPv6 traffic
- Power outage often happen several consecutive times ...
- Dynamic prefixes force a logging system

Best choice: Static or Dynamic-stable

- Allow broadband services provided by the customer and the ISP
 - Allow stable DNS names
 - camera1.username.ispname.com
 - New business/apps/services, new incomes
- Key for non-residential customers
- Avoid having a logging system
- The WAN link still can be dynamic