

IPv6 prefix length considerations

(or, what is the path to align IETF
with reality?)

Job Snijders

job@ntt.net



IEPG, IETF 98, Chicago

RFC 4291bis (RFC 4291 is from 2006)

Section 2.4 “Unicast Addresses”:

*“However, the Interface ID of all unicast addresses, except those that start with the binary value 000, is required to be **64 bits long.**”*

Reality

```
job@tardis:~$ sudo ip -6 addr add 2001:728:1808:1::21/126 dev eth1
```

```
job@tardis:~$ sudo ip -6 addr show dev eth1
```

```
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qlen 1000
```

```
inet6 2001:728:1808:1::21/126 scope global
```

```
valid_lft forever preferred_lft forever
```

```
inet6 fe80::20d:b9ff:fe41:d4f5/64 scope link
```

```
valid_lft forever preferred_lft forever
```

```
job@tardis:~$ echo HERESY\! a /126\!
```

The war front movements (last 20 years)

RFC 3513 - "only /64 is valid"

RFC 3627 - "don't use /127, use /126 if you must"

RFC 4291 - "reaffirming: only /64 is valid"

RFC 6164 - "a /127 is OK to use too"

RFC 6583 - "there are problems with /64"

RFC 7421 - "/64 is the best!"

RFC 7608/BCP198 - "every prefix length must be forward-able"

RFC 4291bis-07 - "fine, /64 and /127 are valid, but nothing else!"

...

RFC ????? "?????"

AS 2914 – all interfaces

/127 - 22%

/126 - 52% <!-- **woahh!!**

/125 - 0,9%

/124 - 0,5%

/120 - 0,04%

/112 - 0,02%

/64 - 23%

Arguments against

- Android may not be compatible
- By accepting something different then /126 or /64 may exist, maybe somewhere some provider will **A=0, M=1** resulting in single **IA_NA**
- Non-/64 won't work for SLAAC

Arguments in favor

- I can configure whatever suits my needs, no need to discuss with IETF
- **Legitimizes** many existing implementations & deployments
 - Linux, *BSD, Nokia, Juniper, Cisco, etc
- Can expand the **routed** network at the edges, in absence of PD
- **Undisputable protection** against IPv6 ND Cache exhaustion
- Future proof: sometimes you just need to re-slice things

The path forward?

- Nobody is saying that /64 is not a good default length
- When I configure through non-SLAAC means, **anything should be valid**
- If you can't deal with having just a single IA_NA, you have lost already
- Do not encourage current & new vendors to only support for /64

Jared intermezzo

IPv6 was designed with CIDR in mind and pre-DHCPv4

- /64 for everyone
 - One prefix length to rule them all
- Plenty of space (compared to ipv4)
- SLAAC will tell me the top 64-bits for me
 - MAC address will do the low part
- Users are smart and memorized their DNS server
 - 2001:418:3ff::53 is easy to remember
 - And my mother can remember it too
- If IPv6 is just IPv4 with more bits, DHCP should perform similarly
 - Unclear to enterprise folks and barrier to entry
 - Does IETF believe enterprise networks exist?