Analysis of the 64-bit Boundary in IPv6 Addressing

draft-carpenter-6man-why64-01

Brian Carpenter
Tim Chown
Fernando Gont
Sheng Jiang
Alexandru Petrescu
Andrew Yourtchenko

March 2014

Why this document?

- RFC7136 (updating RFC4291) states that "For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long."
- Therefore the *de facto* IPv6 subnet prefix is (almost) always /64.
- Why? This draft analyses that question, but **does not propose any change.**
  - although some people do operate with longer prefixes (see following scenarios)

Scenarios for subnet prefixes >/64

- Insufficient address space delegated by ISP
  - Cases like home or small office networks, vehicles, building services, transport services.
  - Use prefixes like /80 to create subnets
  - Homenet architecture draft forbids this
- Desire for a hierarchical prefix scheme that won't fit in /48.../64
- ND cache exhaustion attack
  - Use prefixes like /120 to reduce scope of attack
  - Note that RFC6583 describes other mitigations.

Interaction with IPv6 specifications

- The 64-bit IID size is widely mentioned (a survey found about 30 RFCs and several current drafts).
  - This dates back to 1996.
- In some cases the mention is incidental, but in most the text would certainly need a formal update if longer prefixes and shorter IIDs were allowed.

Breakage with IID < 64 bit

- We have identified definite breakage in the following cases
  - Multicast address formation
  - CGA
  - NAT64 (if prefix >/96)
  - NPTv6
  - ILNP
  - SHIM6 HBA

Experimental observations (1)

- Prefix Information Option (PIO) [RFC4861] with the A bit set and prefix ≠ /64
  - ignored as an error on all O/S tested
- PIO with L or A&L bits set and prefix ≠ /64
  - prefix is considered "on-link" on all o/s tested
- Route Information Option (RIO) [RFC4191] with prefix ≠ /64
  - routed on Windows XP SP2 & 7 Home Premium
  - RIO is not supported on Linux and *BSD variants

Experimental observations (2)

- Longest prefix match based forwarding should work for any prefix length. Some forwarding devices have been shown to work correctly with masks such as /80 or /96.
- DHCPv6 is widely used with no dependency on the /64 boundary. There are deployments of /120 subnets configured using DHCPv6.
- At least one type of switch has a content-addressable memory narrow enough to prevent filtering of long prefixes.

Privacy issues

- IID length has privacy implications.
- More bits allow more randomness.
  - e.g. 40 bits means trillions of guesses needed
  - prefixes longer than ~/80 might make the IID guessable
- By contrast, a /120 would create exactly the same privacy situation as IPv4.
  - hosts forced to pick new IIDs when roaming
  - IIDs cannot be constant across networks

Implementation and deployment (1)

- The /64 assumption is built into an unknown amount of code.
- In practice today, deployments of subnet prefixes >/64 cannot make use of SLAAC.
  - To fix this, all specifications depending on /64 would need to be modified, with due regard to interoperability with unmodified stacks
  - All stacks would need to be updated in due course

Implementation and deployment (2)

- It has been argued that:
  - longer prefixes would allow hierarchical routing in enterprise networks.
  - short IIDs would make discovery and audit of nodes easier.
  - fixed subnet size makes planning, configuration, documentation and training simpler.
  - with /64 there are always free addresses for new devices.

Next steps

- Reminder: this draft **does not propose any change.**
- Conclusions still to be written.
- We want input. What have we missed?
- Does the WG want to document this topic?