SLAAC's Reaction to Renumbering Events

Fernando Gont

IEPG
Prague, Czech Republic. March 24, 2019

Intro

- Operational problem related to renumbering events in SLAAC
- Problem discussed in draft-gont-6man-slaac-renum
  - Previously discussed in draft-linkova-6man-default-addr-selection-update
- Triggered 300+ messages on 6man/v6ops lists!

© 2019 SI6 Networks. All rights reserved

Common scenario

- Sample scenario:

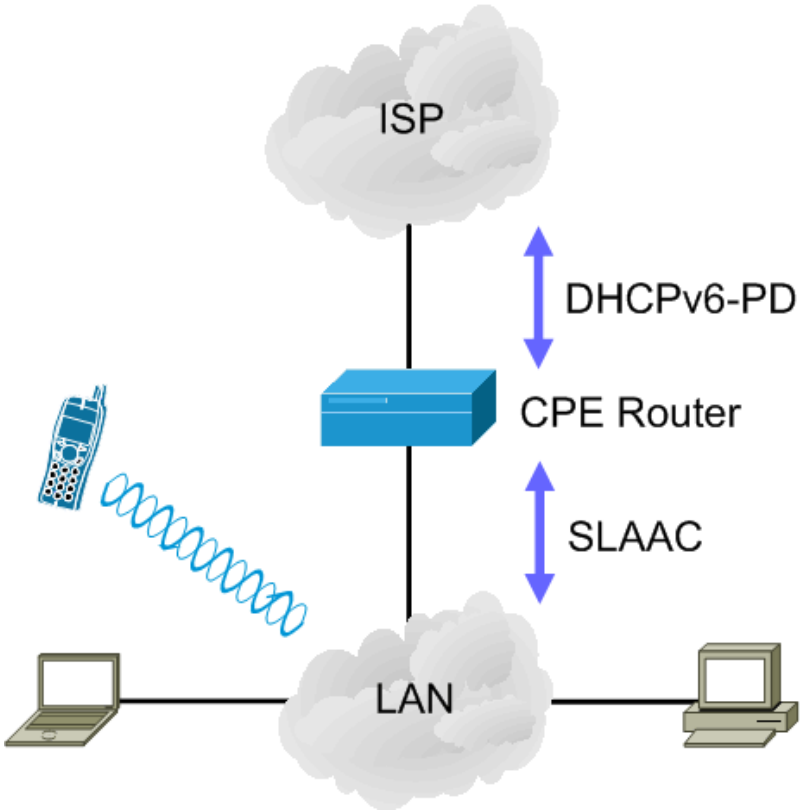

© 2019 SI6 Networks. All rights reserved

Common scenario (II)

- Our typical scenario works as follows:
  - CPE router gets a prefix leased via DHCPv6-PD
  - CPE router announces a sub-prefix via SLAAC
- Typical parameters:
  - Router Lifetime in the order of half an hour
  - Lease times quite usually in the order of several days to months

© 2019 SI6 Networks. All rights reserved

Problem statement

- Problem scenario
  - CPE router is hard-rebooted
  - CPE router crashes and reboots
- What happens when the CPE router comes back to life?
  - Quite frequently it has no state of previously-leased prefix
  - It thus request a new prefix via DHCPv6-PD
  - The new prefix is announced on the LAN
- What about the previous prefix?
  - It is still there!
  - Announced lifetimes allow continued use for days to months

© 2019 SI6 Networks. All rights reserved

Problem statement (II)

- Result:
  - Old addresses are maintained
  - Quite frequently, such addresses are preferred
  - Old routes are maintained
- What does this mean?
  - Connectivity with new owner of prefix not possible
  - IPv6 connectivity may fail
  - In dual-stack scenarios, it may mean more IPv4 traffic
    - Due to Happy Eyeballs

© 2019 SI6 Networks. All rights reserved

Deployments that avoid the problem

- **Sites that use stable prefixes**
- Pro's
  - Nice for law-enforcement – prefix identifies the user!
  - Upon reboots CPE gets same prefix so… no problem!
- Con's
  - Some provisioning systems reportedly don't support this
  - Bad for user privacy – **RFC4941 mostly useless with stable prefixes!**
  - Some ISPs want to charge extra for stable prefixes – ala IPv4

***There is no spoon*. The network should be resilient!**

© 2019 SI6 Networks. All rights reserved

Deployments that avoid the problem (II)

- **CPEs that record leased prefixes on stable storage**
- Many (most?) simply don’t
- It’s tricky, anyway
  - Still cannot invalidate the stale prefix!
  - They have to be able to record many prefixes
  - Lease times of days/months, ant reboots may be frequent
  - And should announce them for remaining leased time
- You cannot rely on the CPE recording prefixes on stable storage

***There is no spoon*. The network should be resilient!**

© 2019 SI6 Networks. All rights reserved

How NOT to solve the problem

- **Update to IPv6 Source Address Selection (RFC6724)**
- Option #1: Prefer prefix with longer Preferred Lifetime
  - Does not make sense
- Option #2:
  - Prefer last advertised prefix
  - Address flapping **guaranteed**!
  - Cannot communicate with new “owners” of the prefix
  - You may hit limit on max number of configured addresses

Hint: If prefix is stale… get rid of it!

© 2019 SI6 Networks. All rights reserved

How we think it should be solved

- **Get rid of stale addresses and router in a timelier manner**
- If the same router advertises a new prefix (but not the previous one), assume the prefix has become stale
- Count number of consecutive RAs from same router with PIOs that do not include the previous prefix:
  - After two such RAs, unprefer the addresses
  - After two additional ones, remove the addresses and routes

This solves the problem at the hosts themselves

© 2019 SI6 Networks. All rights reserved

How we think it should be solved (II)

- This issue begs a number of questions...
- Does it really make sense for Prefix Lifetime > Router Lifetime?
  - In the context of RFC8028, it doesn’t make much sense
  - Announce the prefix for the whole lease time, but never with lifetimes larger than the Router Lifetime.
- What’s the point of announcing a prefix with a lifetime of one month?
  - Just keep the addresses in the event of dead router?

Making appropriate usage of timers can help legacy hosts

© 2019 SI6 Networks. All rights reserved

Questions?

© 2019 SI6 Networks. All rights reserved