Lightweight 4over6: An Extension to DS-Lite Architecture

draft-cui-softwire-b4-translated-ds-lite-09

Y. Cui, Q. Sun, M. Boucadair, T. Tsou, Y. Lee and I. Farrer

IETF 85, Atlanta, Nov. 2012

Motivation

- In DS-Lite, the AFTR maintains per-flow dynamic state. This means:
  - Scaling the AFTR is prohibitive
  - NAT Logging overhead
- Relocate NAPT functionality from AFTR to B4 to:
  - Leverage the existing NAPT44 functions in CPE
  - Reduce the amount of state in the AFTR, to reduce the amount of resources, as well as logging overhead
- Keep IPv4 and IPv6 addressing architectures completely independent

Key points for lw4over6

- State maintenance
  - Subscriber binding table: [IPv6 address, IPv4 address, port-set]
- LwAFTR behavior
  - Holds the binding table: Act as a DHCP server/relay agent, or synchronize the binding with DHCP server
  - Binding table lookup: (IPv4 dst addr, dst port) => IPv6 addr
  - Packet encapsulation/de-encapsulation (RFC2473)
- LwB4 behavior
  - Request IPv4 address + port-set (DHCP)
  - Packet encapsulation/de-encapsulation, NAT within a given external port-set

Lightweight 4over6 features

- A simple extension to DS-Lite [RFC6333]
  - Address sharing mode for Public 4over6 (draft-ietf-softwire-public-4over6)
- Per-subscriber state maintenance on AFTR
  - Good scalability for overwhelming traffic; easy/no logging
- No algorithmic IPv4/IPv6 address coupling
  - Operation simplicity
  - No constraints on address planning
  - Flexible port management: on-demand/ pre-configured IPv4 addr+port-set provisioning
- DS-Lite backward compatibility

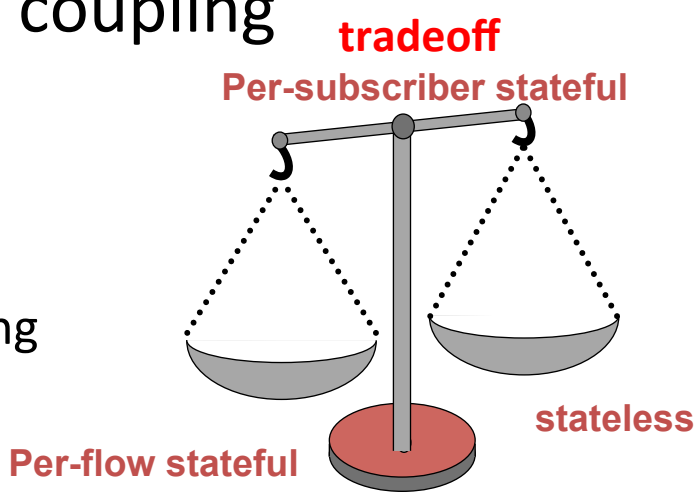

IPv4 / IPv6 Addressing Dependency Problem – Initial State

MAP-E Mesh/Hub&Spoke

lw4o6

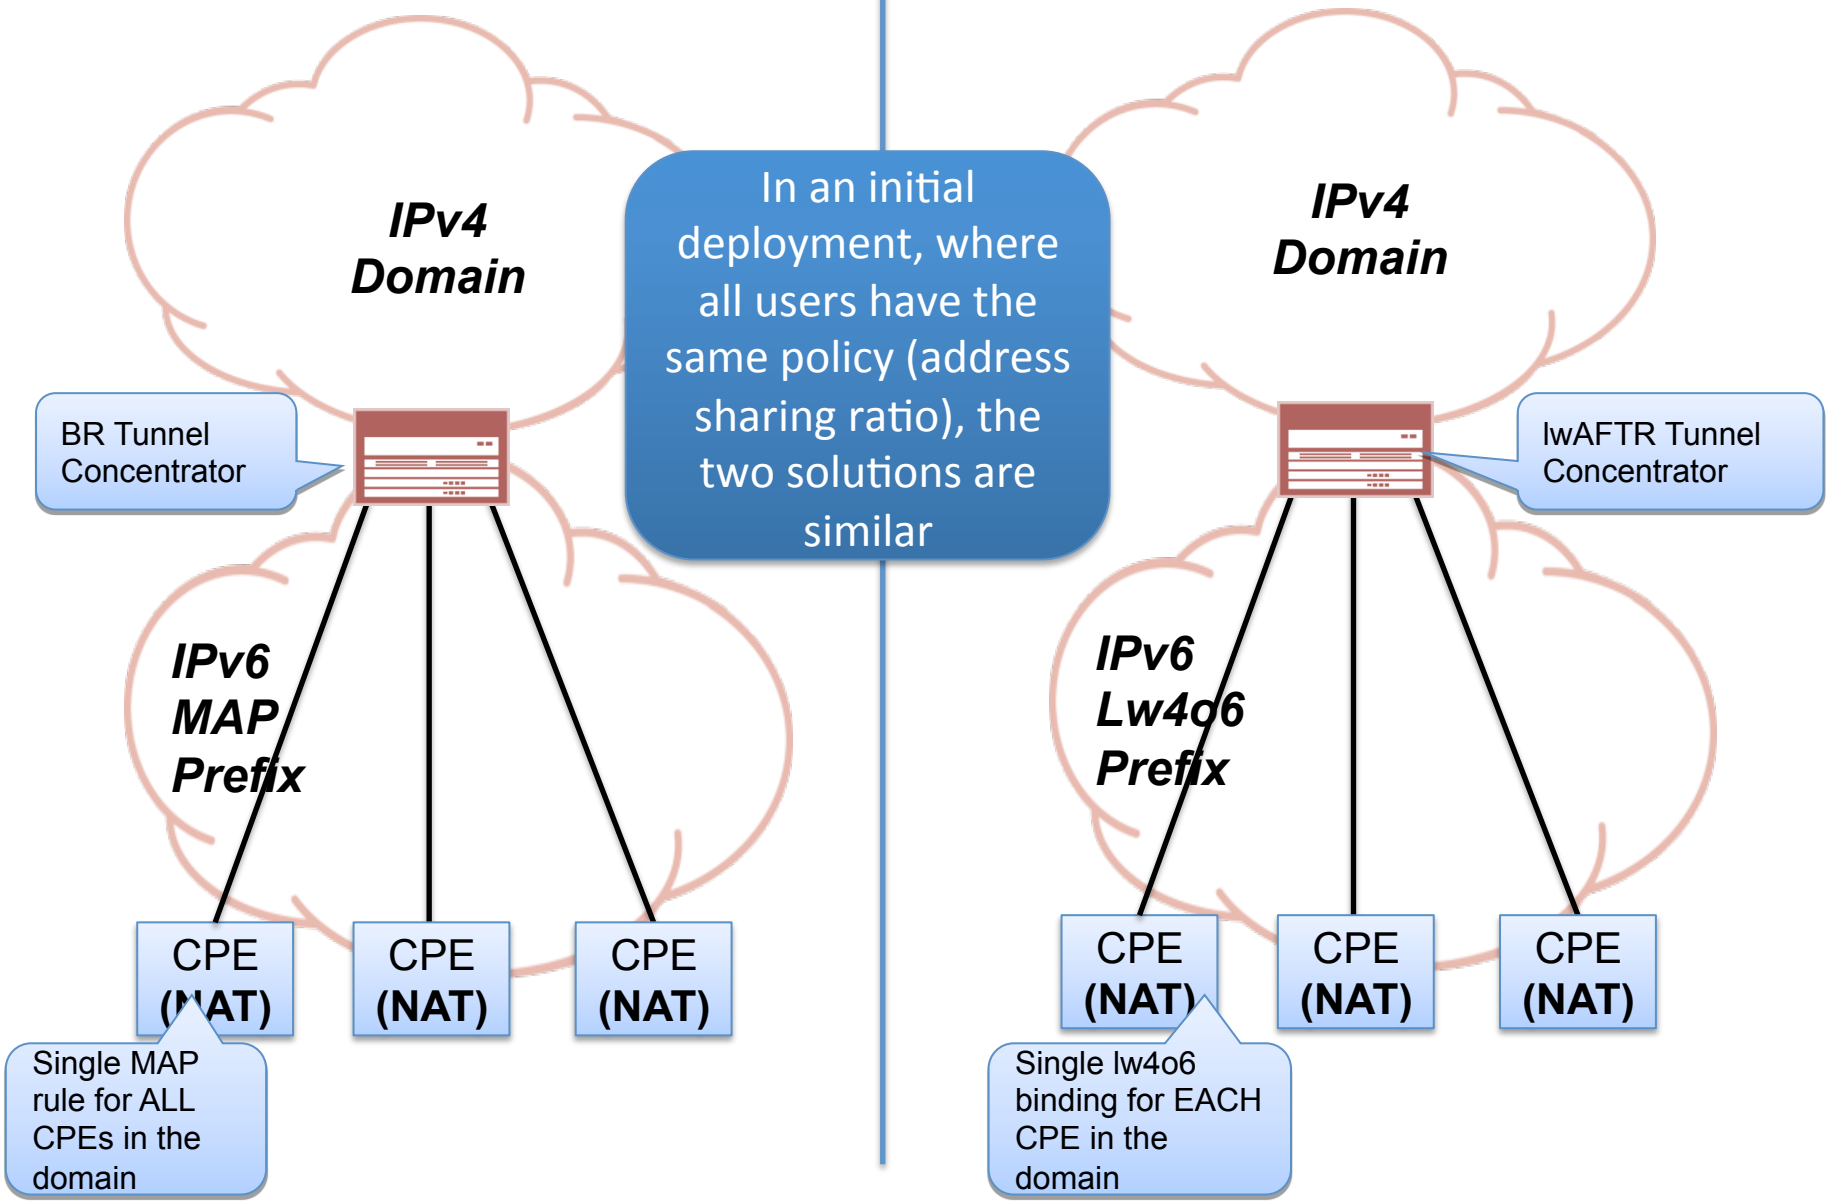

IPv4 / IPv6 Addressing Dependency Problem – Over time

MAP-E Mesh/Hub&Spoke lw4o6

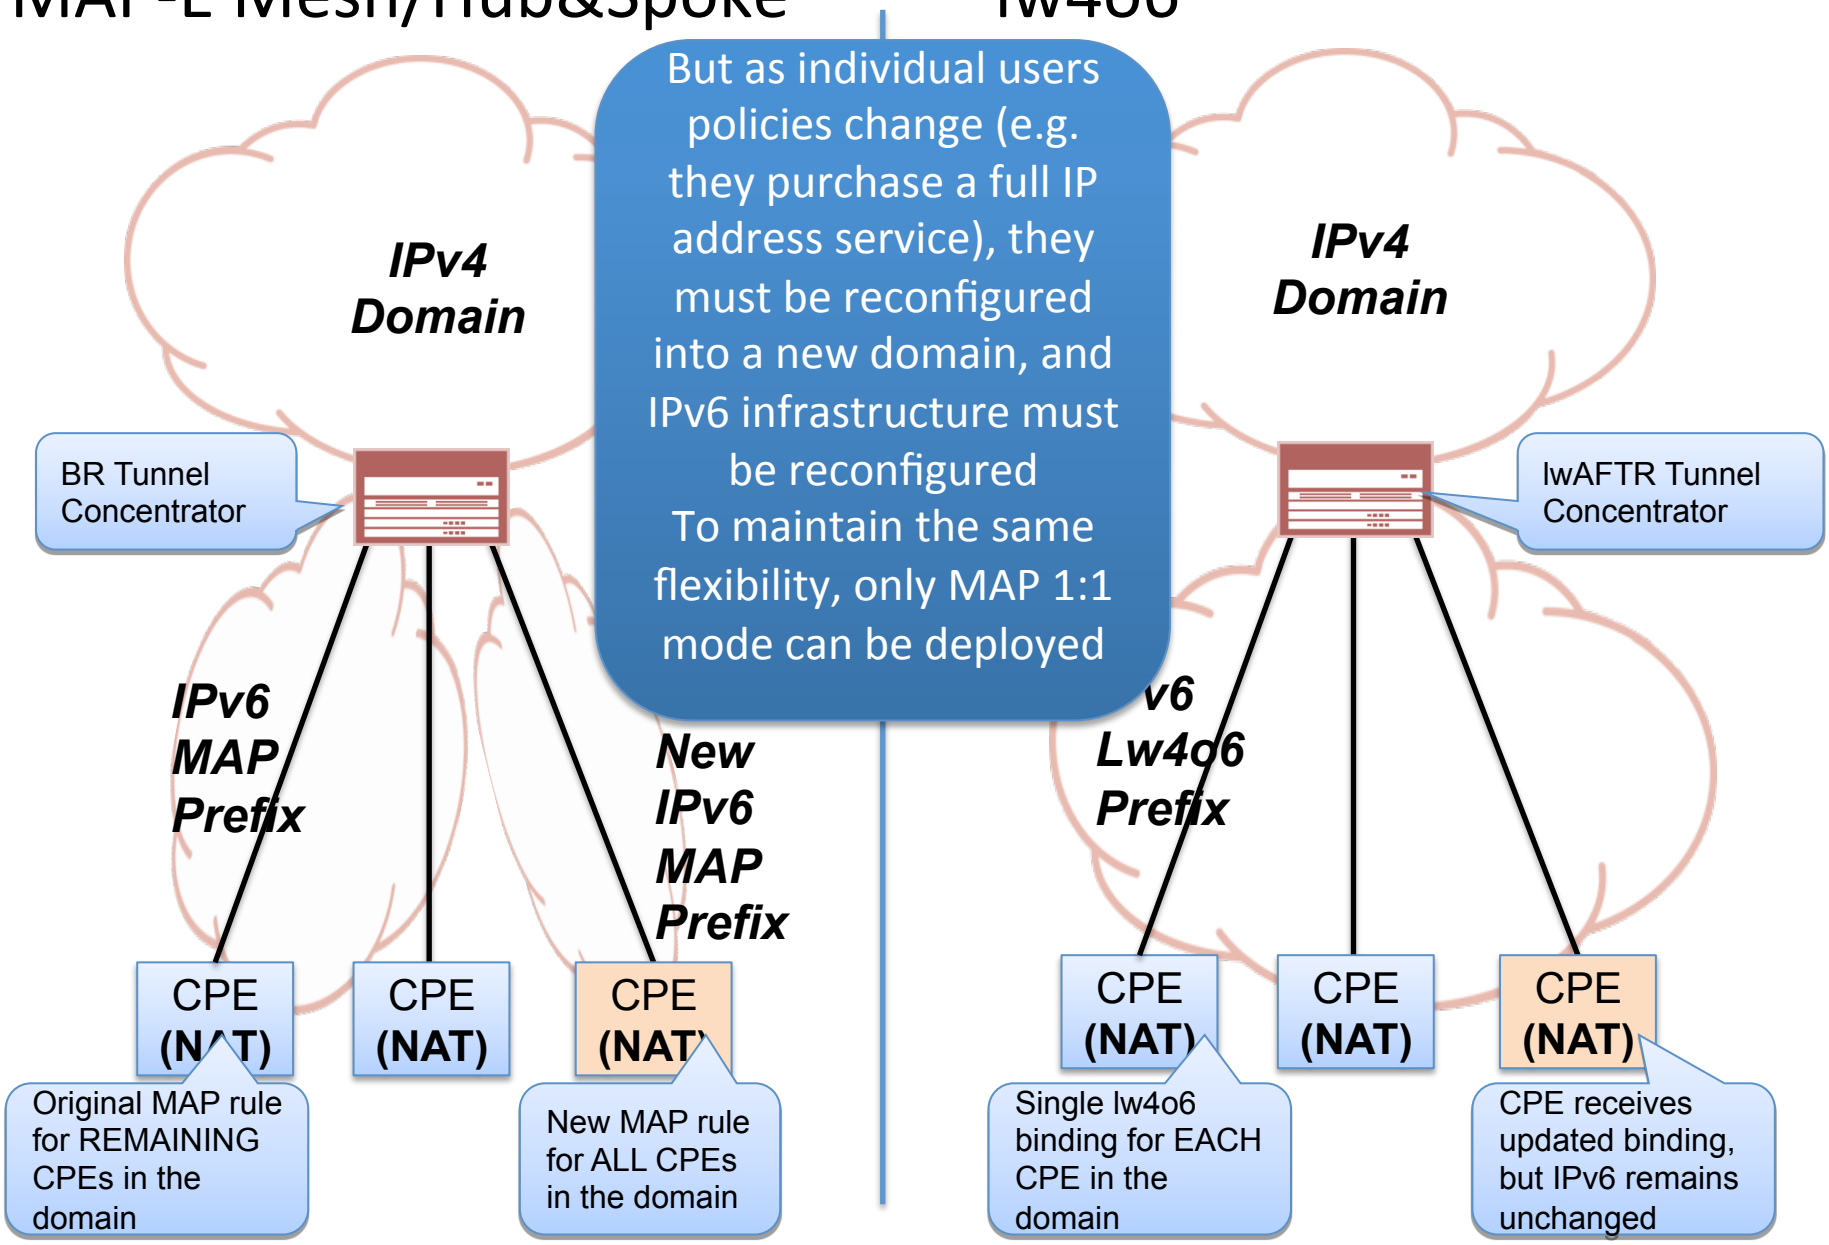

MAP/lw4o6

- Key benefits of MAP-E
  - If direct mesh communication between CPEs is required
  - For implementations that have larger subscriber groups with communities common port overload requirement

- Key benefits of lw4o6
  - If variance and change in subscriber port overload requirements is necessary
  - Optimized for the per-subscriber mapping / address sharing for 1:1
  - As only the 1:1 mode is supported, no addition operational complexity is required for other functional modes

Running Code/Field Trial/Interop

- Implementations by:
  - Tsinghua University
  - China Telecom
  - Freebit
  - Deutsche Telekom
  - Huawei (lwAFTR/lwB4)
  - Yamaha (lwB4)
  - Fiberhome (lwB4)
  - GreenNet (lwAFTR/lwB4)

- Field trialed or tested successfully in:
  - China Telecom
  - Deutsche Telekom
  - France Telecom
  - CERNET
- Successful Interop testing between:
  - Tsinghua University, China Telecom, Huawei, Yamaha, NetDominator, BII

Outcome of Vancouver meeting

- Clarify the problem space
  - Solution Optimized Per-subscriber state only
  - Solution keeps IPv4 and IPv6 addressing architectures completely separate
- WG guidance on draft revision
  - Focusing on the differences from DS-Lite and its motivation/advantages

Updates from -07

- Draft rewritten following the WG guidance

- Terminology changes
  - Lightweight 4over6 initiator => lwB4
  - Lightweight 4over6 concentrator => lwAFTR
  - Add the definition of Port-Set
- Change the reference to DHCP port-set draft
  - [I-D.bajko-pripaddrassign] => [I-D.sun-dhc-port-set-option]
  - [I-D.sun] updates [I-D.bajko] and [I-D.wu]

Updates from -07

- Change the provisioning methods description
  - DHCPv4ov6 is RECOMMENDED and its DHCPv6 option is SHOULD
  - Define the behavior for full port set case as per public 4over6 for lwB4 & lwAFTR
- Improve ICMPv4 handling for the lwAFTR
  - Support inbound ICMP echo requests: check for a valid port in the ID field and forward if present, drop if not
  - Modify the alternative from 'discard all inbound ICMP requests' to 'discard all inbound ICMP messages'

Next steps

- The current version reflects the requirements of the WG
- The implementation and trial result are good
- Call For Working group Adoption?

Backup

Lightweight 4 over 6 Overview