Connecting IPv4 Islands over IPv6 MPLS Using IPv4 Provider Edge Routers(4PE)

Zhenqiang Li

China Mobile

Scenario to be addressed

- When IPv6-only networks are widely deployed, how to provide t he connectivity for the remaining IPv4-only islands through the I Pv6-only MPLS network will become a problem.
- RFC7439, Gap Analysis for Operating IPv6-Only MPLS Networks, has already pointed out this gap.
- 4PE, IPv4 Provider Edge Routers, is proposed in this draft to mee t the gap.
- RFC4798, IPv6 Provider Edge Routers (6PE), is used to address th e reverse scenario, connecting IPv6-only islands through IPv4-on ly MPLS network.

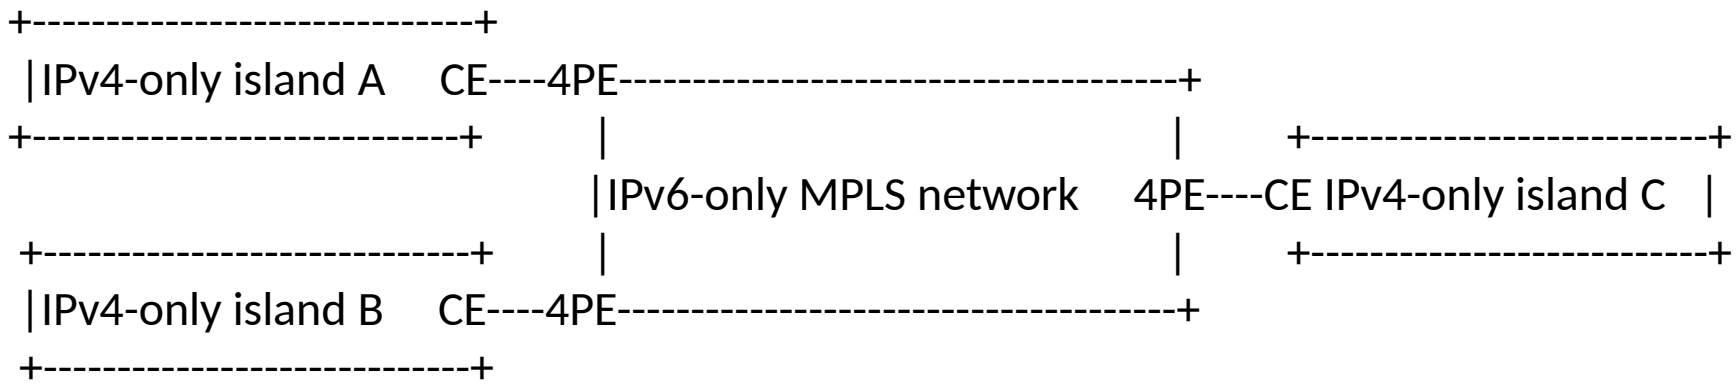

What we have to do

- Exchange IPv4 reachability information among 4PE routers
  - MP-BGP, RFC4760, is extended with a new Subsequence Address Family Identifier (SAFI) to do so. This new SAFI is called 4PE SAFI.
- Transport IPv4 packets from the ingress 4PE router to the egress 4PE router
  - Ingress 4PE router looks up its IPv4 routing table, when it receives packe t from the IPv4 island.
  - If the matched table entry is learned from other 4PE router, ingress 4PE uses the IPv6 next hop to reach the corresponding egress 4PE through t he IPv6 MPLS LSP.

The 4PE SAFI

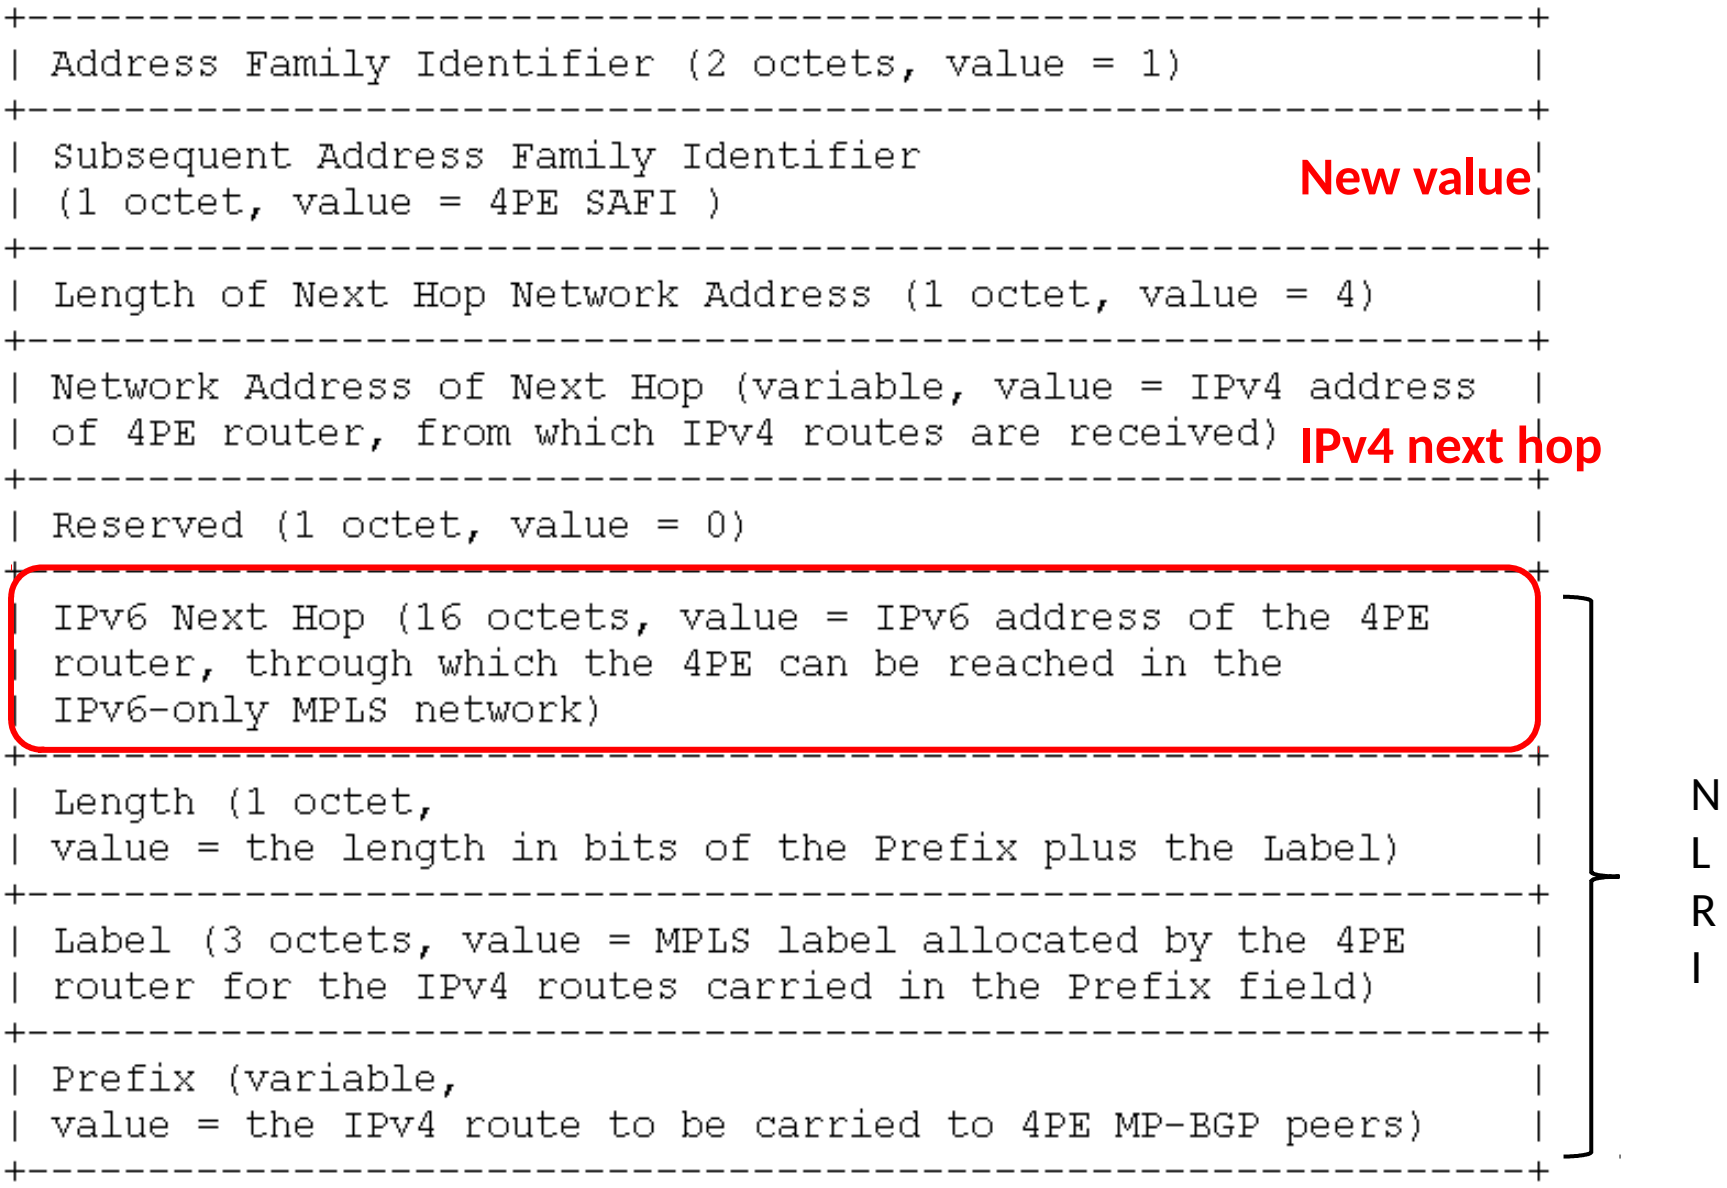
Address Family Identifier (2 octets, value = 1)
Subsequent Address Family Identifier
(1 octet, value = 4PE SAFI)
New value
Length of Next Hop Network Address (1 octet, value = 4)
Network Address of Next Hop (variable, value = IPv4 address
of 4PE router, from which IPv4 routes are received)
IPv4 next hop
Reserved (1 octet, value = 0)
IPv6 Next Hop (16 octets, value = IPv6 address of the 4PE
router, through which the 4PE can be reached in the
IPv6-only MPLS network)
Length (1 octet,
value = the length in bits of the Prefix plus the Label)
Label (3 octets, value = MPLS label allocated by the 4PE
router for the IPv4 routes carried in the Prefix field)
Prefix (variable,
value = the IPv4 route to be carried to 4PE MP-BGP peers)
NLRI

When 4PE receives 4PE SAFI

- Since 4PE SAFI is a kind of MP-BGP message, the 4PE router trea ts it as per [RFC4760] and [RFC3107].
- 4PE routers MUST distinguish the IPv4 routes learned from othe r 4PE routers and those from the IPv4-only island directly conne cted to it.
- 4PE router MUST establish the relation between IPv4 Next Hop, IPv6 Next Hop and MPLS label carried in the 4PE SAFI. Through this relation, 4PE routers can get MPLS label and IPv6 Next Hop using IPv4 Next Hop of the matching routing table entry when it forwards received IPv4 packet. The method or data structure us ed to do this is an implementation issue.

When 4PE receives IPv4 packet

- 4PE router treats the IPv4 packet as normal IPv4 router does except for the following steps.
  - If the matching IPv4 route for this packet is learned from other 4PE routers, the 4PE router has further to get the IPv6 Next Hop and MPLS label using the matching IPv4 next hop.
  - Then, 4PE router uses the IPv6 Next Hop to lookup in its IPv6 routing table to get the IPv6-signaled LSP to reach the egress 4PE router.
  - Next, 4PE router encapsulates the received IPv4 packet using two labels and forwards it toward the egress 4PE router through the IPv6-signaled LSP.

IANA Requirements

- IANA is requested to assign a new SAFI for the 4PE SAFI. Numbe r 9 is suggested

Discussion

- Eric thinks RFC5549 is an alternative solution.
- RFC5549 encodes IPv6 next hop in the Network Address of Next Hop field dir ectly and the value of Length of Next Hop Network Address field is used to ju dge IPv4 or IPv6 next hop is encoded in the Network Address of Next Hop fie ld.
- Since RFC5549 doesn't carry IPv4 next hop for IPv4 route, how to install thos e IPv4 routes in the IPv4 routing table is a problem. Eric insists this is a imple mentation issue.

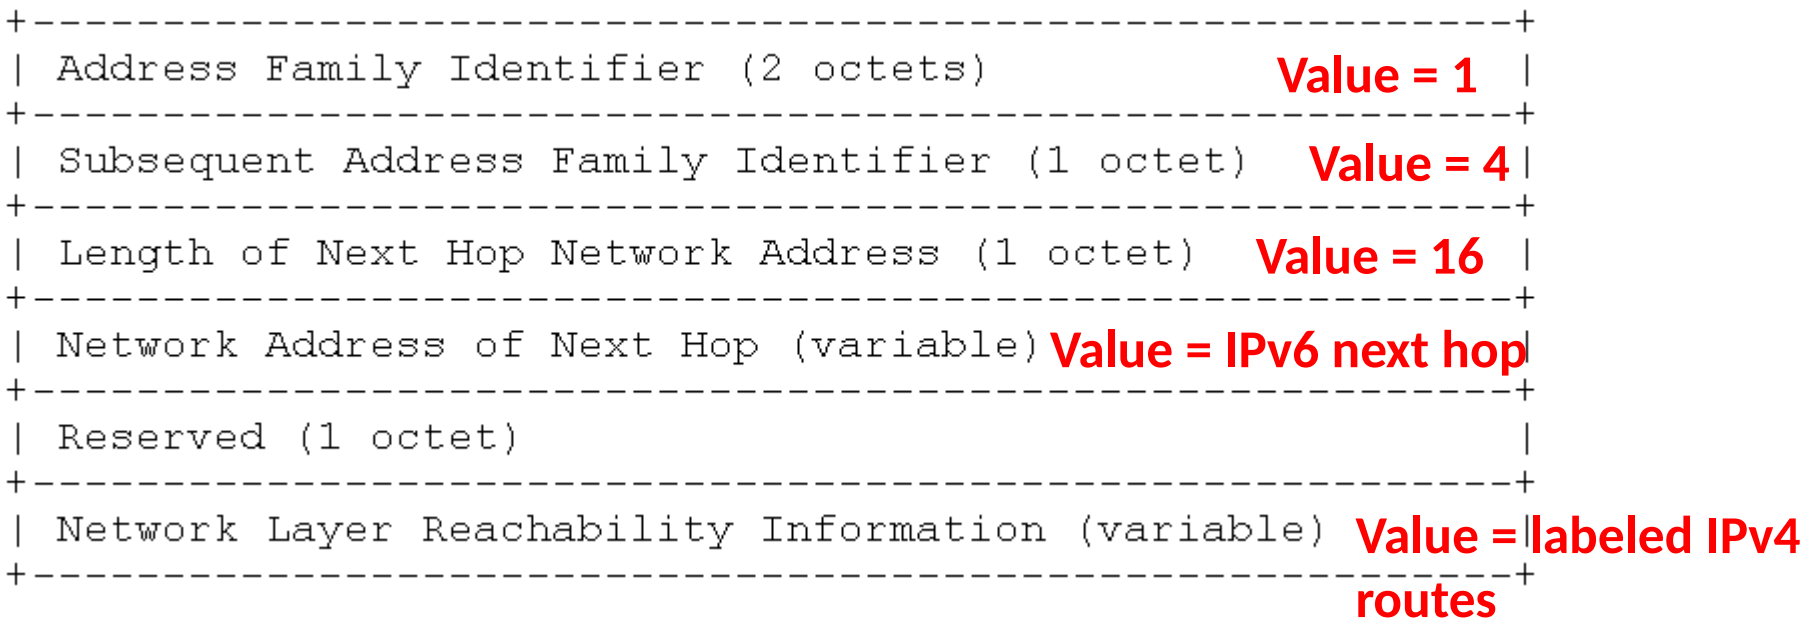

- Thanks
- Comments to
  - li_zhenqiang@hotmail.com
  - bess@ietf.org