Supporting BIER in non-MPLS IPv6 Networks

Sandy Zhang, Jeffrey Zhang
IJsbrand Wijnands, Hooman Bidgoli

IETF 109

IPv6 is Nothing Special to BIER

- Existing procedures defined for IPv4 non-MPLS networks apply to IPv6 with no need for any changes or enhancements
  - BIER header follows L2/tunnel header as L2/tunnel payload
    - Do need a new “next header” type defined for BIER in case of native IPv4/v6 tunnel
- IPv4/IPv6 tunneling/encap, if used, is just a transport means to BIER
  - Just like any other means, e.g. MPLS/GRE/whatever
  - Between BFRs not directly connected
    - BIER header is the beginning part of IP payload
    - IP encapsulation not needed between directly connected BFRs
  - IPv4/IPv6 and BIER are independent of each other

BIERin6

- draft-zhang-bier-bierin6 reflects the concepts in previous slide
- IPv6 encapsulation needed only between BFRs not directly connected
  - BIER header follows L2 header directly between directly connected BFRs
- It does ***optionally*** use IPv6 encapsulation even between directly connected BFRs
  - To allow certain platforms that does not yet support BIER Ethertype
  - To allow software-based BIER forwarding in certain deployment scenarios
- Purposes of this draft:
  - Specify the above ***optional*** feature
  - Explain how existing BIER procedures can work for IPv6 networks

BIERv6

- draft-xie-bier-ipv6-encapsulation
- BIER header is encoded in IPv6 Destination Options Header (DOH)
  - BFIR->BFER IPv6 encapsulation end-to-end
  - SRv6-style handling, with special overlay (MVPN/EVPN) procedures
- WG requested to have a requirements draft first
  - To justify this additional solution
  - Multiple solutions allowed ***if there are significant advantages***

Requirements for Supporting BIER in IPv6

- draft-ietf-bier-ipv6-requirements
- Mandatory Requirements
  - Basically, support BIER architecture
- Optional Requirements
  - Fragmentation
  - IPsec ESP
- ***Nothing IPv6 specific***

BIERin6 Satisfies All Listed Requirements

- Optional fragmentation/ESP requirements can be met by one of two methods

1. IPv6 encapsulation with fragmentation/ESP, then treated as BIER payload
   - L2/tunnel header + BIER header + IPv6 header
   - IPv6 encapsulation and BIER are independent of each other
2. Generic fragmentation/ESP (no IP encapsulation), then treated as BIER payload
   - L2/tunnel header + BIER header + Generic fragmentation/ESP header
   - No IP encapsulation and its overhead
   - draft-zzhang-tsvwg-generic-transport-functions
     - Generically applicable to MPLS/BIER or any layer (even Ethernet if IEEE so desires)

BIERv6 Not Needed

- BIERin6, which is consistent with existing IPv4 solution, satisfies all requirements
  - No IPv6 overhead with BIERin6
    - Saves minimum 40 bytes compared to BIERv6
    - Use that 40 bytes for BitString – accommodates 320 more BFERs
- No need for BIERv6 as an additional solution
  - IPv6 encapsulation overhead with no obvious benefits
    - Why bother with BFIR->BFER IPv6 encapsulation
  - Complexities with handling BIER header encoding in DOH
    - Discussed separately
  - SRv6-style handling not needed but leads to complexities (e.g. MVPN procedures)
    - Hop-by-hop BIER forwarding does not need SRv6
    - BFIR/BFER do not need SRv6 based network programming

Suggested Next Steps

- Adopt BIERin6 as WG document
- Discuss BIERv6 further