The IPv6 Compressed Routing Header (CRH)

Draft-bonica-6man-comp-rtg-hdr

R. Bonica, Juniper Networks
N. So, F. Xu, Reliance Jio
G. Chen, BAIDU
Y. Zhu, G. Yang, China Telecom
Y. Zhao, ByteDance

Network Programming -> Two SID Classes

Transport SIDs

- Steers packets to the terminal segment
- Processed at non-terminal segment endpoints (SL > 0)
- Example: END, END.X
- Relatively few of these
- Simple semantic
  - Carries relatively little information

Service SIDs

- Determines behavior at the terminal segment
- Processed at terminal segment endpoint only (SL = 0)
- Example: END.DX4, END.DX6
- Relatively many of these
- Rich semantic
  - Carry many bits of information

IPv6 -> Two Ways To Deliver Instructions To Downstream Nodes

Routing Extension Header

- Steers packets from ingress to egress
- Processed at non-terminal segment endpoints (SL > 0)
- Well-positioned to carry Transport SIDs

Destination Options Header

- Determines behavior at egress node
- Processed at terminal segment endpoint only (SL = 0)
- Well-positioned to carry Service SIDs

The Problem With Routing Headers

- Too Long
  - Typically 8 bytes of overhead (4 bytes are mandatory)
  - Typically, another 16 bytes per SID
  - Routing header with 3 SIDs is 56 bytes long
- Not ASIC Friendly
  - Processing long extension headers is computationally expensive
- Impose unreasonable bandwidth overhead
  - Short packets (>500) bytes are common on the Internet
  - Routing header with three SIDS may become common
  - > 10% Routing header overhead

Proposal

- Encode Transport SIDs in a new, compressed routing header
  - Draft-bonica-6man-comp-rtg-hdr
  - Topic of this talk
- Encode Service SIDs in new IPv6 Destination Options
  - Draft-bonica-6man-vpn-dest-opt
  - Topic of the next talk

Compressed Routing Header (CRH)

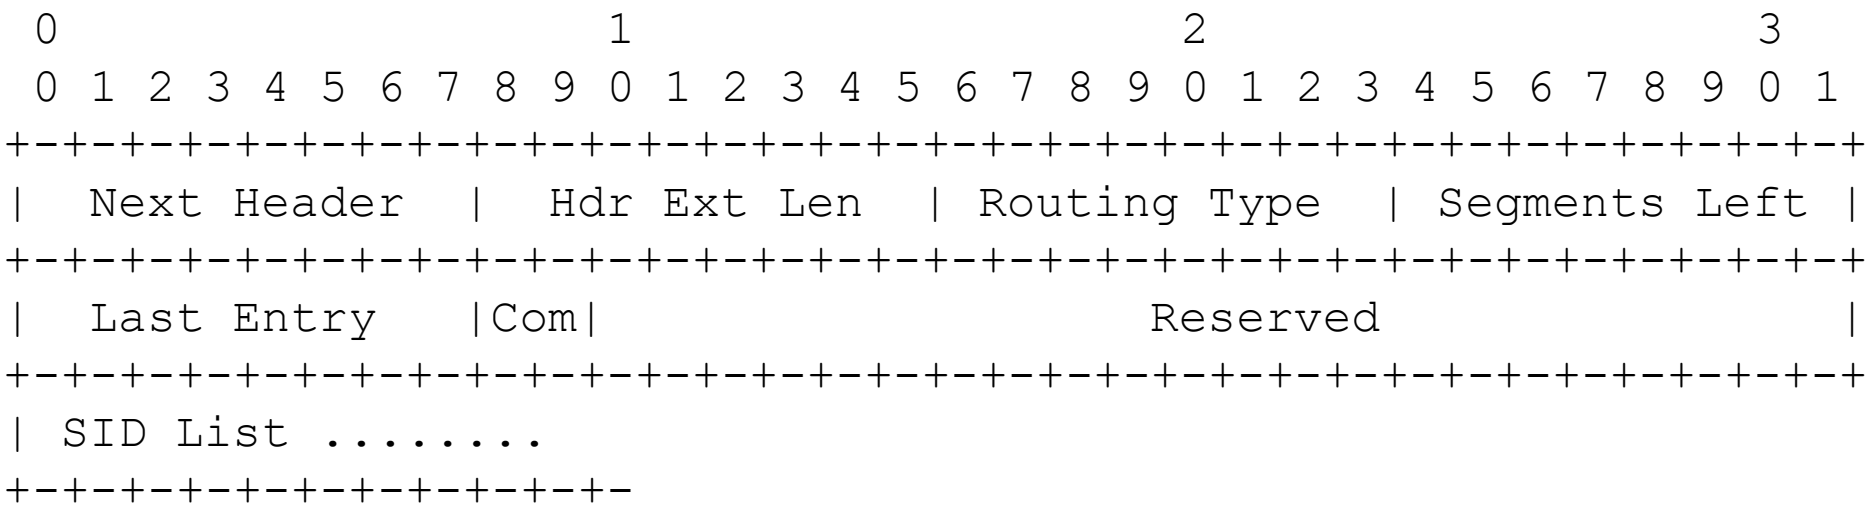

Compressed Routing Header (CRH)

- Initial fields are defined in RFC 8200 and common to all Routing headers
  - Next Header, Header Extension Length, Routing Type and Segments Left
- Last Entry is a pointer to the final entry
- Com field indicates whether SIDs are 8, 16, or 32 bits long
- Each SID maps to an IPv6 address
  - Either through a table lookup or an algorithm
- IPv6 address is copied to the Destination Address field of the IPv6 header

Compressed Routing Header (CRH): Com Equals Zero (8-bit SIDs)

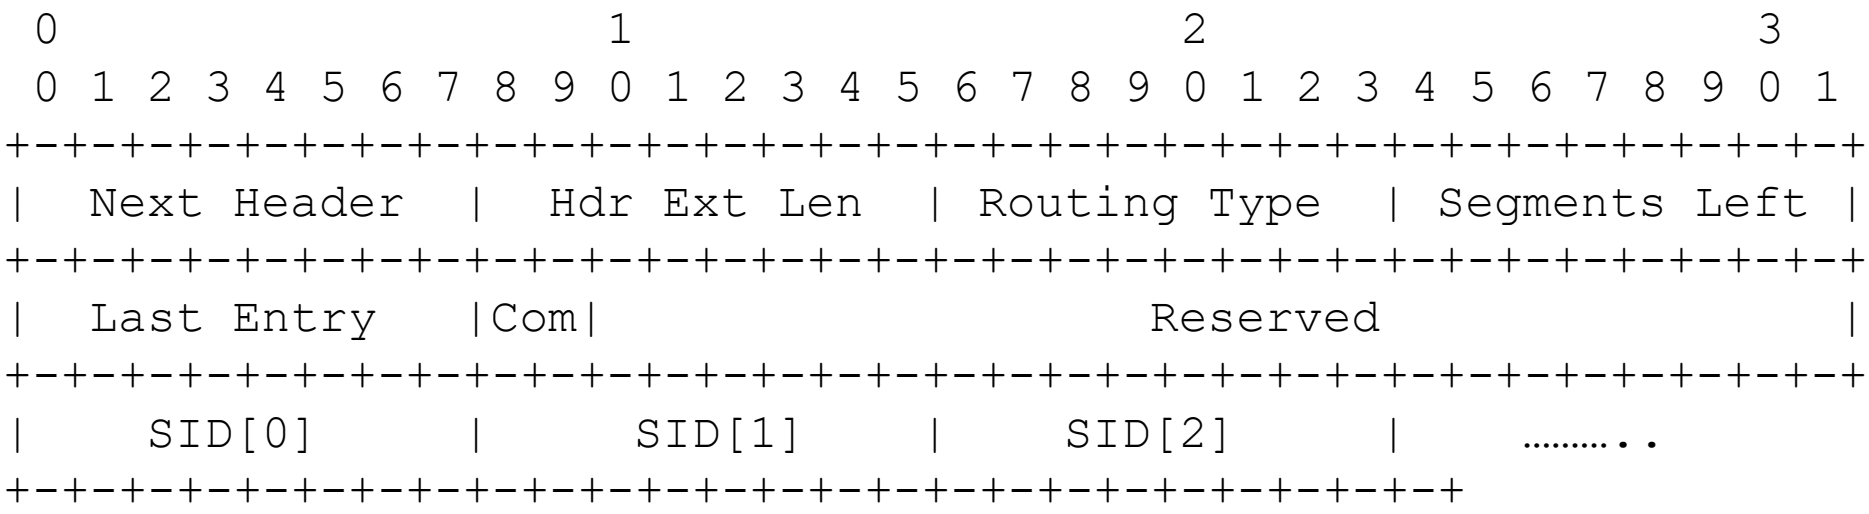

Compressed Routing Header (CRH): Com Equals One (16-bit SIDs)

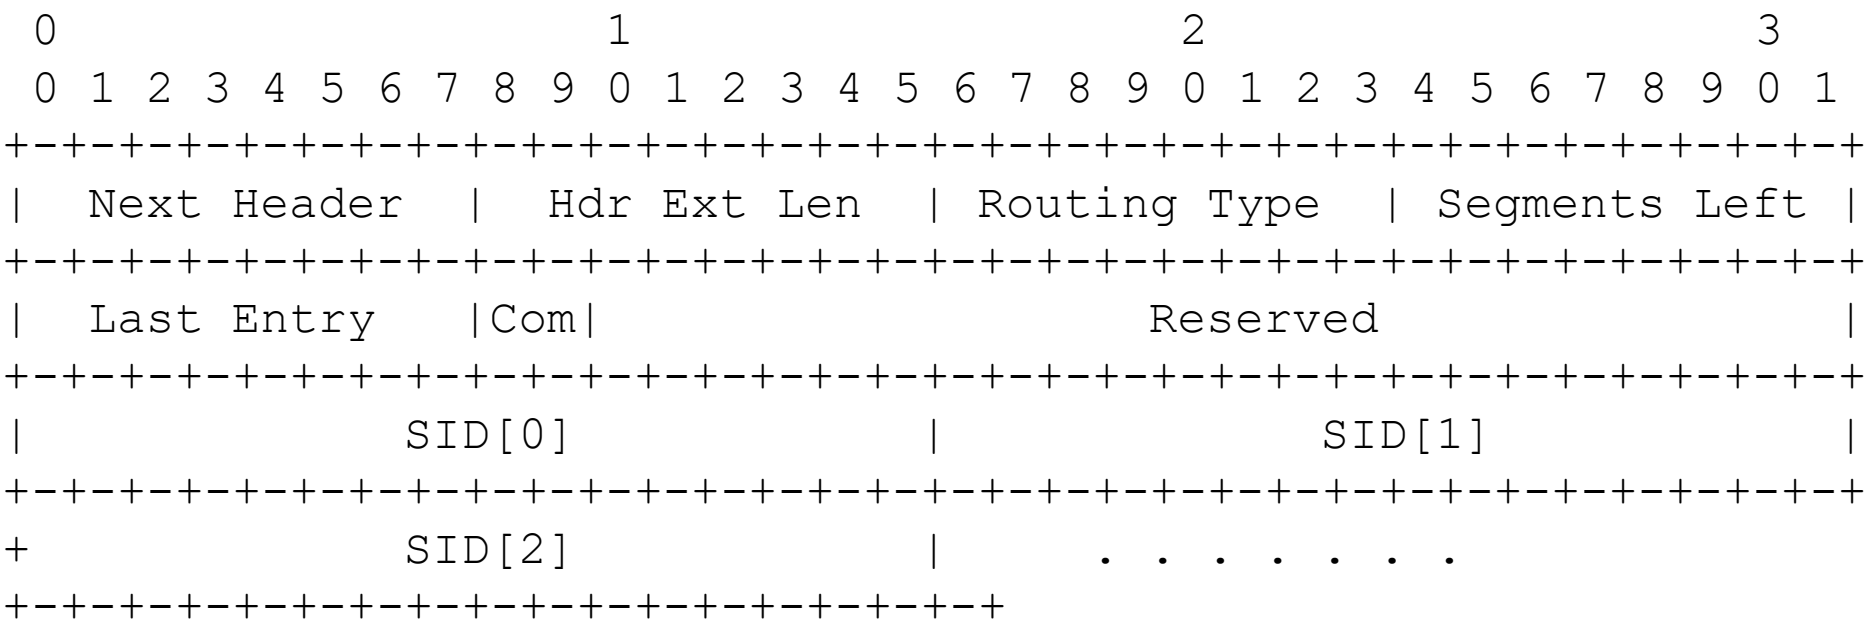

Compressed Routing Header (CRH): Com Equals Two (32-bit SIDs)

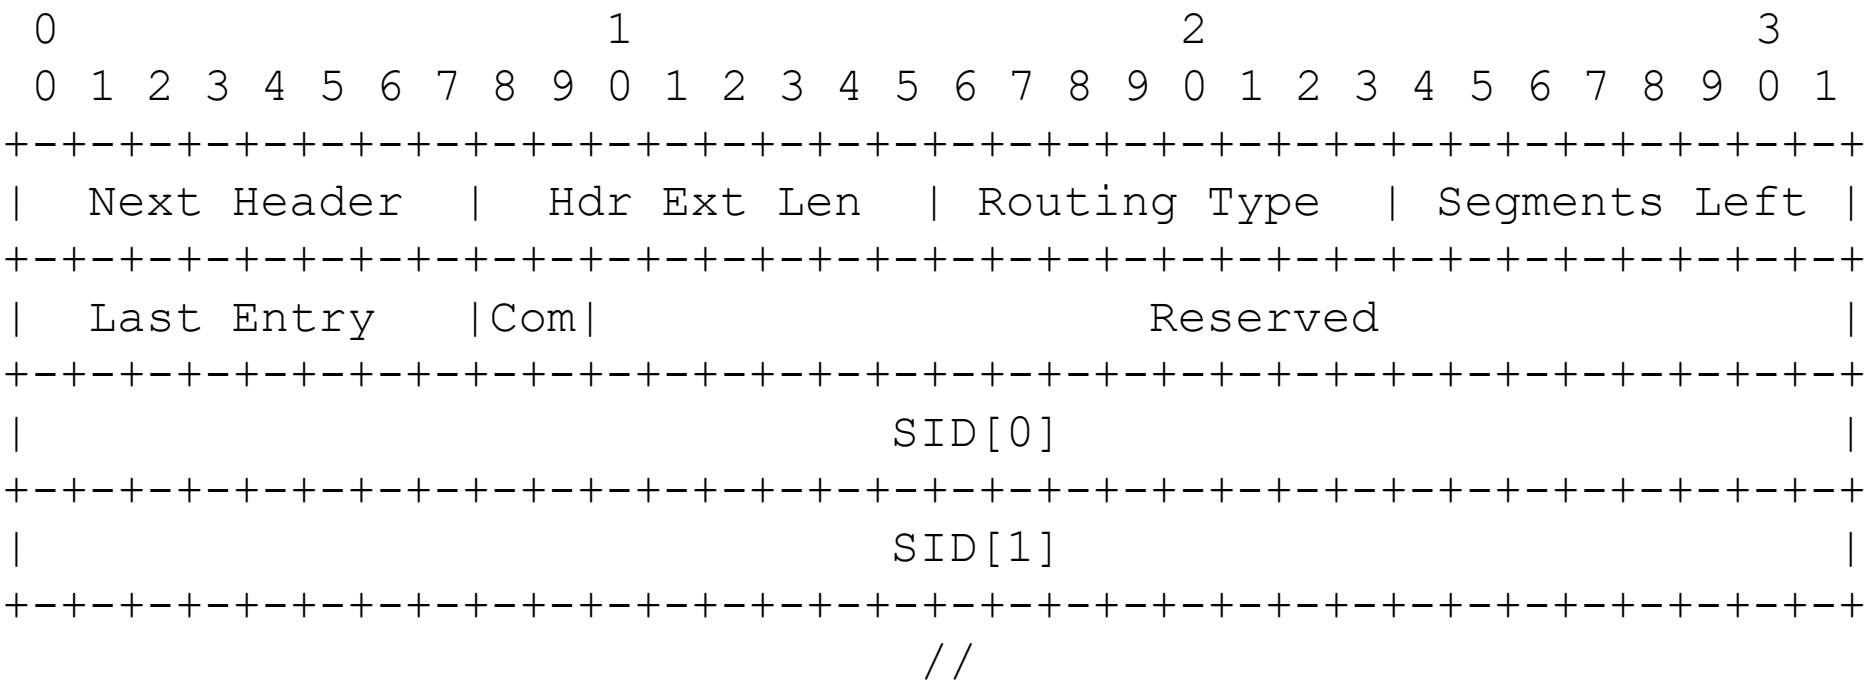

Compliance Benefits

- IPv6 address semantics remain unchanged
  - An IPv6 address represents a nodes interface to a subnetwork [RFC 4291]
- No need for transit routers to insert or delete extension headers
- No need for two instances of the CRH in the same packet
- Extension headers are processed in the order that they appear in the packet
  - No need to backtrack when Segments Left is equal to zero

Status

- Operators are expressing interest
- Prototypes under development
  - Forwarding plane
  - ISIS Extensions to support SID Advertisement

Next Steps

- Wide review in SPRING and 6man WGs
- Call for adoption in 6man WG