BE Multicast using Multicast Routing Header

draft-chen-pim-be-mrh-00

Huaimo Chen, Donald E. Eastlake, Mike McBride (Futurewei)
Yanhe Fan (Casa Systems)
Gyan Mishra (Verizon)
Yisong Liu (China Mobile
Aijun Wang (China Telecom)
Xufeng Liu (IBM Corporation)
Lei Liu (Fujitsu)

IETF 114

Brief Description

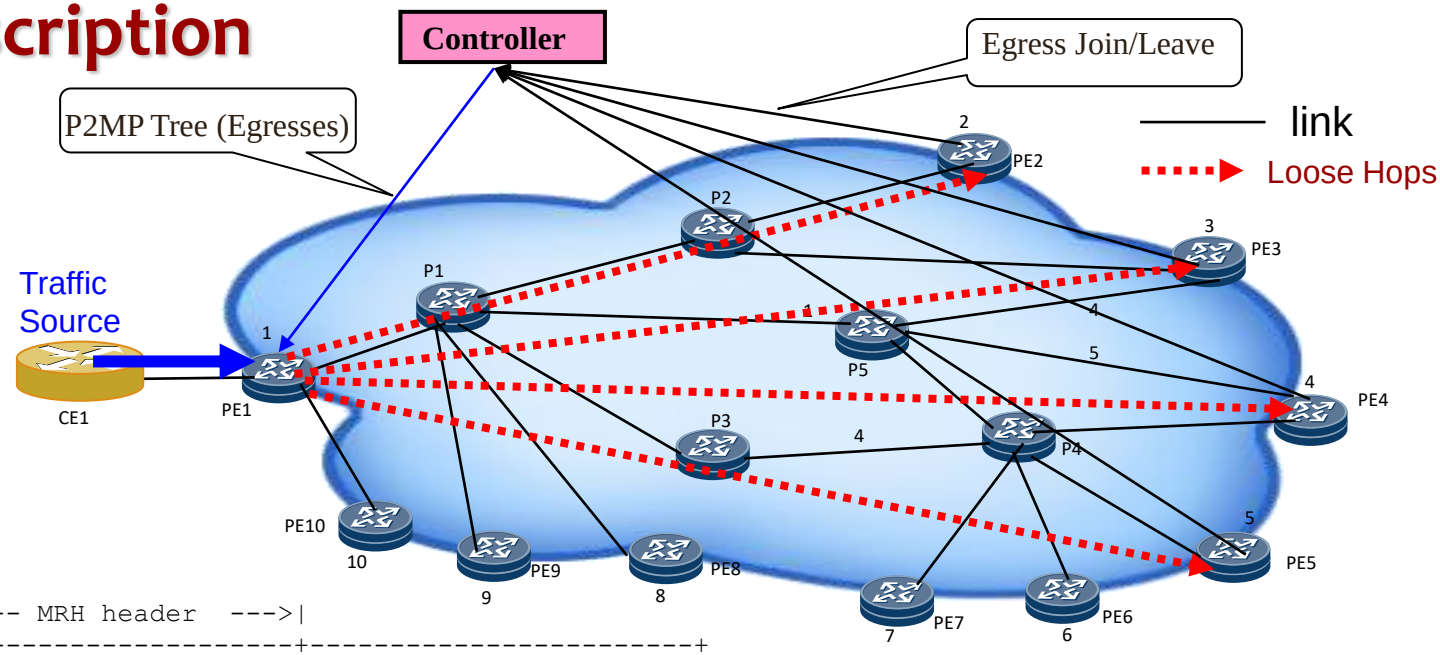

Figure 1. Tree from PE1 towards PE2 to PE5

```
|<--IPv6 header-->|<--- MRH header  --->|
+-----------------+---------------------+-----------------------+
|  Next Header =  | Next Header         | (an extension header) |
|60(DoH) or 43(RT)|                     | IP multicast datagram |
| SA=IPv6 Address |   Type = TBD (MRH)  |                       |
| DA=IPv6 Address |Data(Subtree/Egresses|                       |
+-----------------+---------------------+-----------------------+
                  |<------- MRH ------->|
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 | Next Header   |  Hdr Ext Len  |    Type = TBD |     . . .     |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+               +
 |                                                               ~
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
 |           Tree/Sub-tree (i.e., egresses) encoded              |
 :                                                               :
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- Ingress (e.g., PE1) encapsulates the packet in a MRH with tree (i.e., egresses of tree)
- The packet is transmitted along the shortest IGP pathes to the egresses.
- Egress (e.g., PE2) decapsulates the packet in a MRH and sends it to next header process

Routing Header for BE Multicast

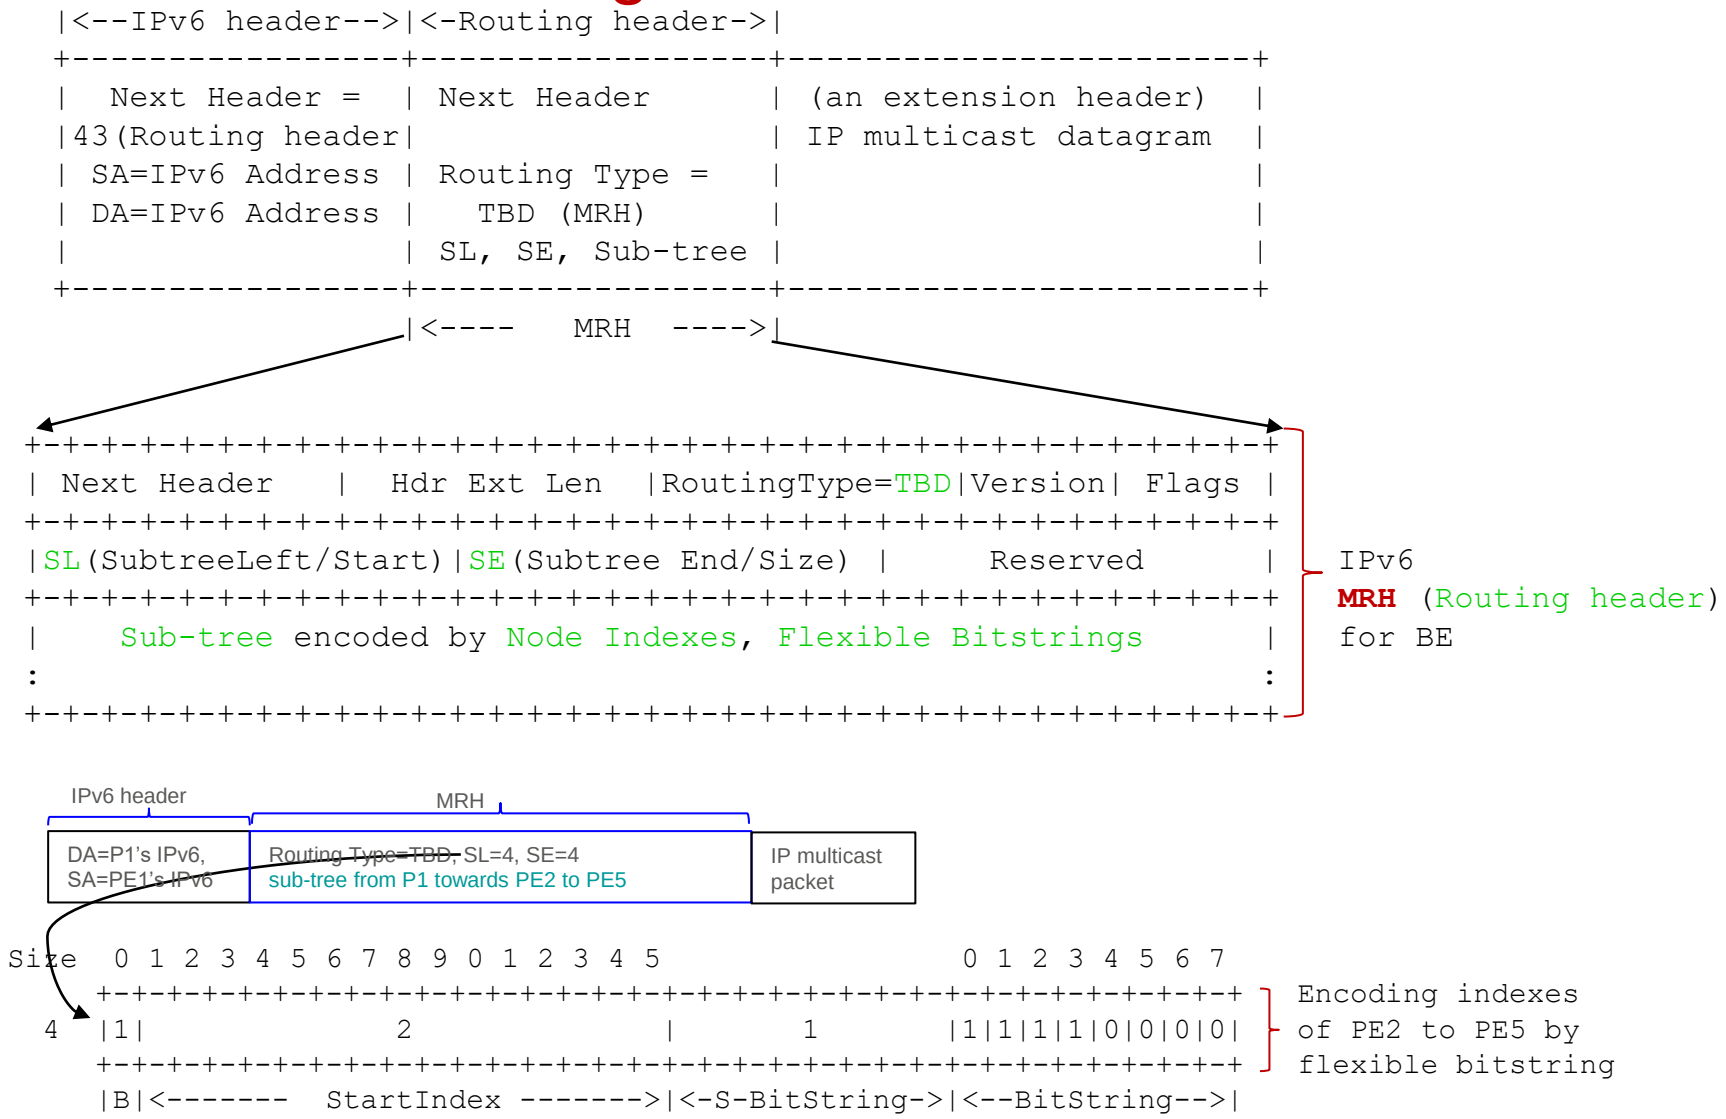

IPv6 packet with one bitstring sent to P1

Encoding by Flexible Bitstring and NodeIndex

A flexible bitstring has four fields:

1). B flag with value 1, 2). start index (StartIndex), 3). size of bitstring (S-BitString) in bytes and

4). bitstring (BitString), where each bit with value1 indicates a node index equal to StartIndex plus the bit number. Note that the bit number is counted from right to left and from 0.

For example, the indexes of egresses PE2 to PE5 (i.e., PE2, PE3, PE4 and PE5) are encoded by a flexible bitstring (suppose their indexes are 2, 3, 4 and 5 respectively):

B = 1, StartIndex = 2, S-BitString = 1, BitString = 0b11110000 indicating four node indexes 2, 3, 4 and 5.

BitString's first bit (bit 0) with value 1 indicates the first node index 2 equal to 2 + 0; the BitString's second bit (bit 1) with value 1 indicates the second node index 3 equal to 2 + 1, and so on.

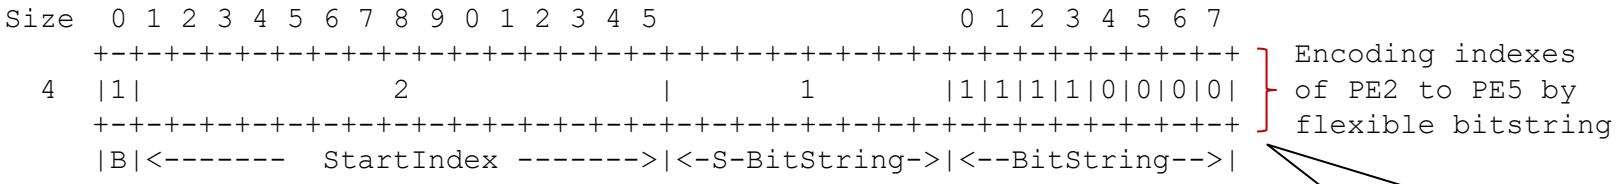

Indexes of PE2 to PE5 by string:
4 bytes, operations on 8 bits

A NodeIndex field with B = 0 represents a node index directly.

For example, the indexes of egresses PE2 to PE5 are represented by NodeIndex.

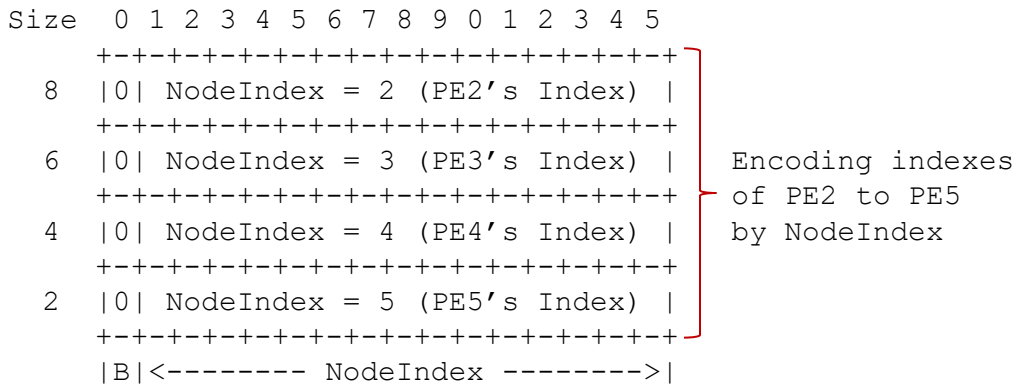

Indexes of PE2 to PE5 by NodeIndex:
8 bytes, operations on 4 Node Indexes

Using bitstring is more efficient than using NodeIndex