IPv6 (IP version 6) Essentials Ch2: The Structure of the IPv6 Protocol

Louis Chuang
Fu Jen Catholic University
EE ENCL

General Header Structure (1)

- The header structure of an IPv6 packet is specified in RFC 2460.
- The header has a fixed length of 40 bytes.
- The two fields for source and destination addresses each use 16 bytes (128 bits).
- In IPv6, five fields from the IPv4 header have been removed:
  - Header Length.
    - IPv4: is variable from 20~60-byte.
    - IPv6: is fixed 40-byte, & options are defined by Extension headers.
  - Identification.
  - Flags.
  - Fragment Offset.
    - The Identification field, Flags field, & Fragment Offset field handle fragmentation of a packet in the IPv4 header.
    - Fragmentation happens if a large packet has to be sent over a network that only supports smaller packet sizes.

General Header Structure (2)

- IPv4 router splits the packet into smaller slices and forwards multiple packets.
- The destination host collects the packets and reassembles them.
- If only one packet is missing or has an error, the whole transmission has to be retransmitted.
- It is very inefficient.
- In IPv6, a host learns the Path Maximum Transmission Unit (PMTU) size through a procedure called Path MTU Discovery.
- If an IPv6 host wants to fragment a packet, it will use an Extension header.
- IPv6 routers do not provide fragmentation.

- Header Checksum.
  - Header Checksum field was removed to improve processing speed.
  - Routers do not check and update checksums, its processing becomes much faster.

General Header Structure (3)

- Checksumming is done at the media access level (MAC layer), & the risk for undetected errors and misrouted packets is minimal.
- Moreover, there is also a checksum field at the transport layer (UDP and TCP).
- IP layer does not need the checksum anymore.

The Fields in the IPv6 Header (1)

- Fields in the IPv6 header.

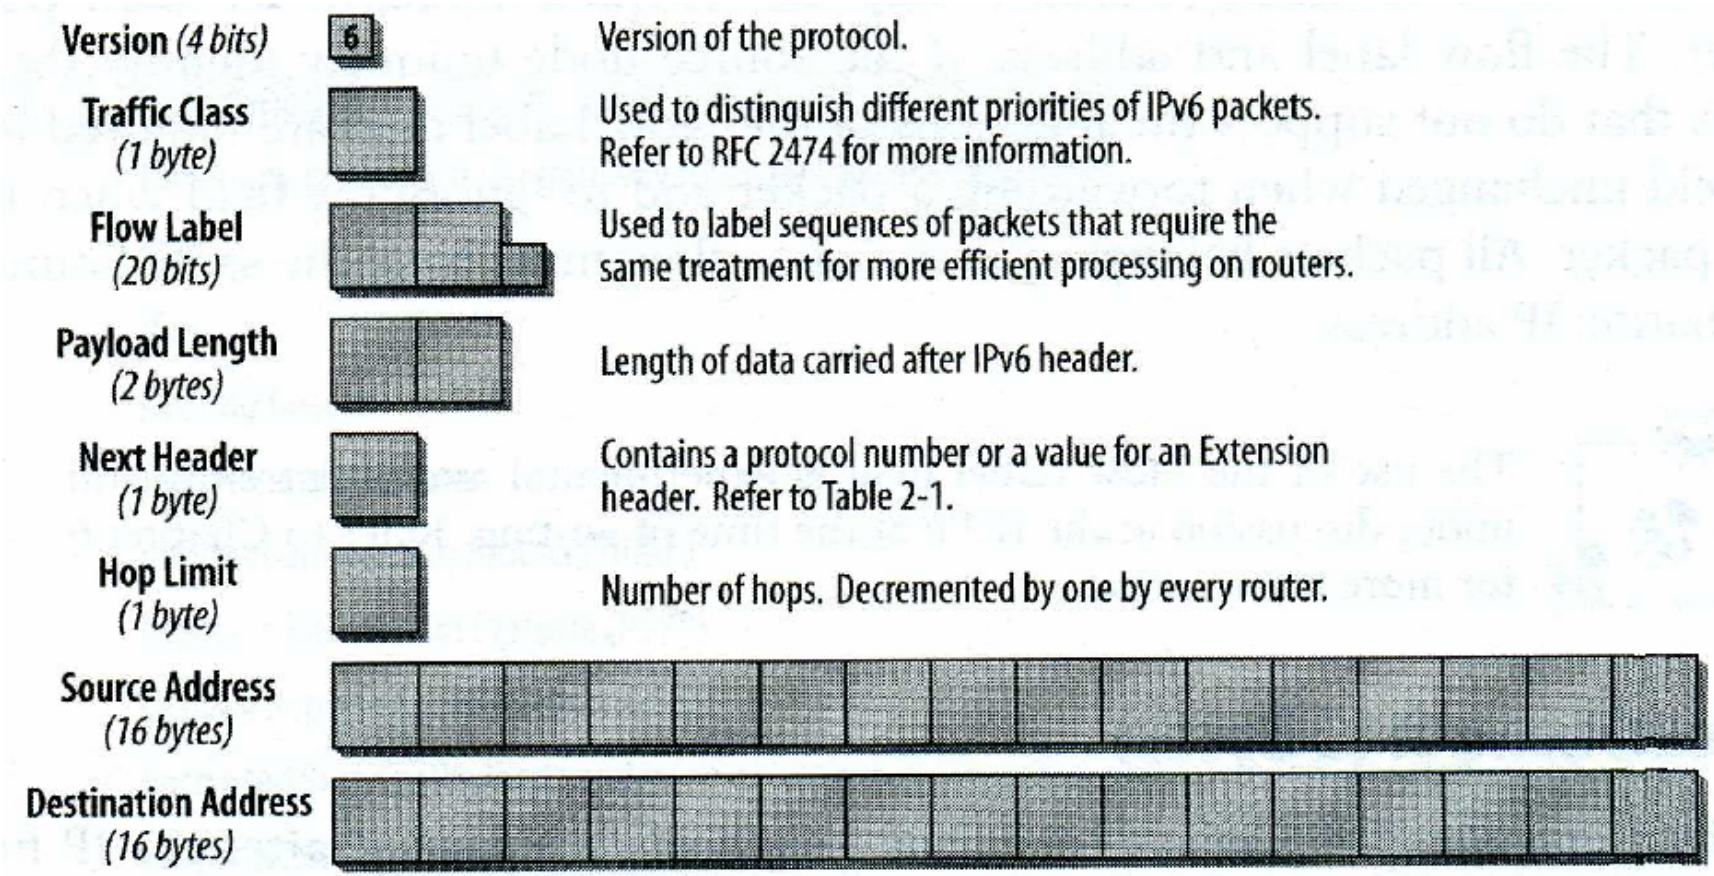

The Fields in the IPv6 Header (2)

- Version field (4 Bits):
  - Version 5 is for an experimental stream protocol (ST2, RFC 1819).
  - IPv6 is version 6.
- Traffic Class (1 Byte):
  - This field replaces the Type of Service (TOS) field in IPv4.
  - This field is for the real-time data and any other data that requires special process.
  - To identify different classes or priorities of IPv6 packets.
  - RFC 2474: "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers" explains how the Traffic Class field in IPv6 can be used.
- Flow Label (20 Bits):
  - This field distinguishes packets that require the same treatment.
  - It is helpful to the real-time traffics.

The Fields in the IPv6 Header (3)

- For a flow, routers can process packets belonging to the same flow more efficiently because they do not have to reprocess each packet's header.
- A flow is uniquely identified by the flow label & the source address.

- Payload Length (2 Bytes):
  - The calculation in IPv6 is different from the one in IPv4:
    - In IPv4: the length includes the IPv4 header & payload.
    - In IPv6: the length only for the payload, which follows the IPv6 header.
  - In IPv6, the Extension headers are considered part of the payload.
  - the Payload Length field has 2 bytes limits the maximum IPv6 packet payload size to 64 KB.
- Hop Limit (1 Byte):
  - This field is similar to the TTL (Time-to-Live) field in IPv4.

The Fields in the IPv6 Header (4)

- In IPv4: the TTL field contains a number of seconds. But most routers simply decremented this value by 1 at each hop.
- In IPv6: the value in this field (Hop Limit) represents a number of hops. Each router decrements this value by 1 at each hop.

- Source Address (16 Bytes) & Destination Address (16 Bytes):
  - Link local address.
  - Site local address.
  - Global address.
- Next Header (1 Byte):
  - In IPv4: this field is the Protocol Type field (upper layer protocol: TCP or UDP).
  - In IPv6: this field presents protocol number of the next header. If the next header is TCP or UDP, Next Header is 6 (TCP) or 17 (UDP).
  - If Extension headers are used in IPv6, this field presents the type of the next Extension header.

The Fields in the IPv6 Header (5)

- The Extension header is located between the IP header and the TCP or UDP header.

The Fields in the IPv6 Header (6)

- The IPv6 header in a trace file.

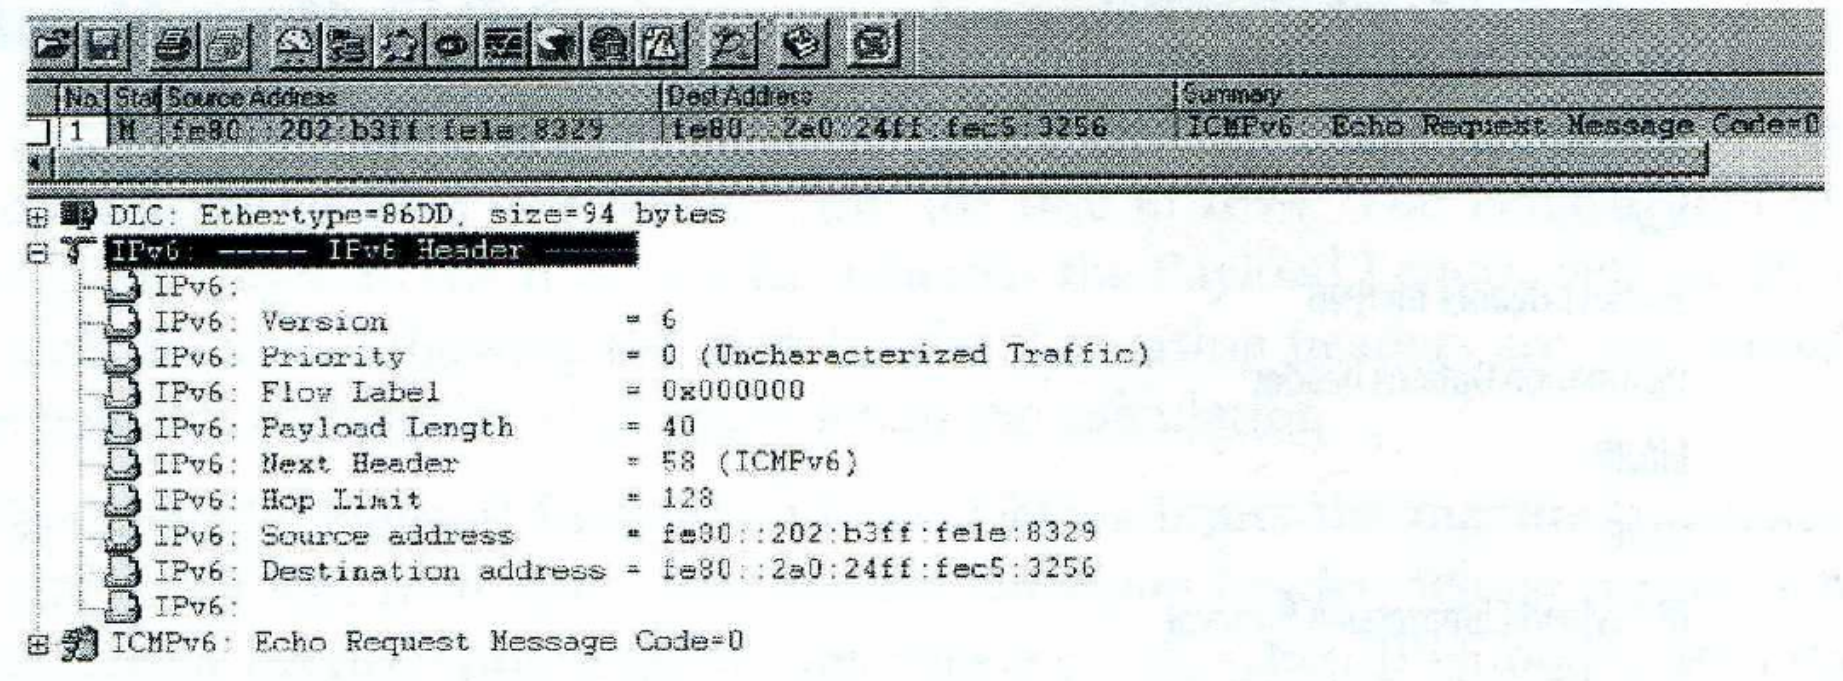

The Fields in the IPv6 Header (7)

Table 2-1. Values in the Next Header field

Value	Description
0	In an IPv4 header: reserved and not used In an IPv6 header: Hop-by-Hop Option Header following
1	Internet Control Message Protocol (ICMPv4)—IPv4 support
2	Internet Group Management Protocol (IGMPv4)—IPv4 support
4	IP in IP (encapsulation)
6	TCP
8	Exterior Gateway Protocol (EGP)
9	IGP - any private interior gateway (used by Cisco for their IGRP)
17	UDP
41	IPv6
43	Routing header
44	Fragmentation header
45	Interdomain Routing Protocol (IDRP)
46	Resource Reservation Protocol (RSVP)

The Fields in the IPv6 Header (8)

50	Encrypted Security Payload header
51	Authentication header
58	ICMPv6
59	No Next Header for IPv6
60	Destination Options header
88	EIGRP
89	OSPF
108	IP Payload Compression Protocol
115	Layer 2 Tunneling Protocol (L2TP)
132	Stream Control Transmission Protocol (SCTP)
134-254	Unassigned
255	Reserved

Extension Headers (1)

- Options in IPv4: such as Security Options, Source Routing Options, or Timestamping Options, etc. make the IPv4 header from a minimum of 20 bytes to 60 bytes. This capacity has rarely been used because it causes a performance hit.
- Options in IPv6: put them into extension headers.
- The current IPv6 specification (RFC 2460) defines 6 Extension headers:
  - Hop-by-Hop Options header (No. 0).
  - Routing header (No. 43).
  - Fragment header (No. 44).
  - Destination Options header (No. 60).
  - Authentication header (No. 51).
  - Encrypted Security Payload header (No. 50).
- Between the IPv6 header & the upper-layer protocol (TCP or UDP) header, it can put zero, one, or more than one Extension header.

Extension Headers (2)

- The use of Extension headers.

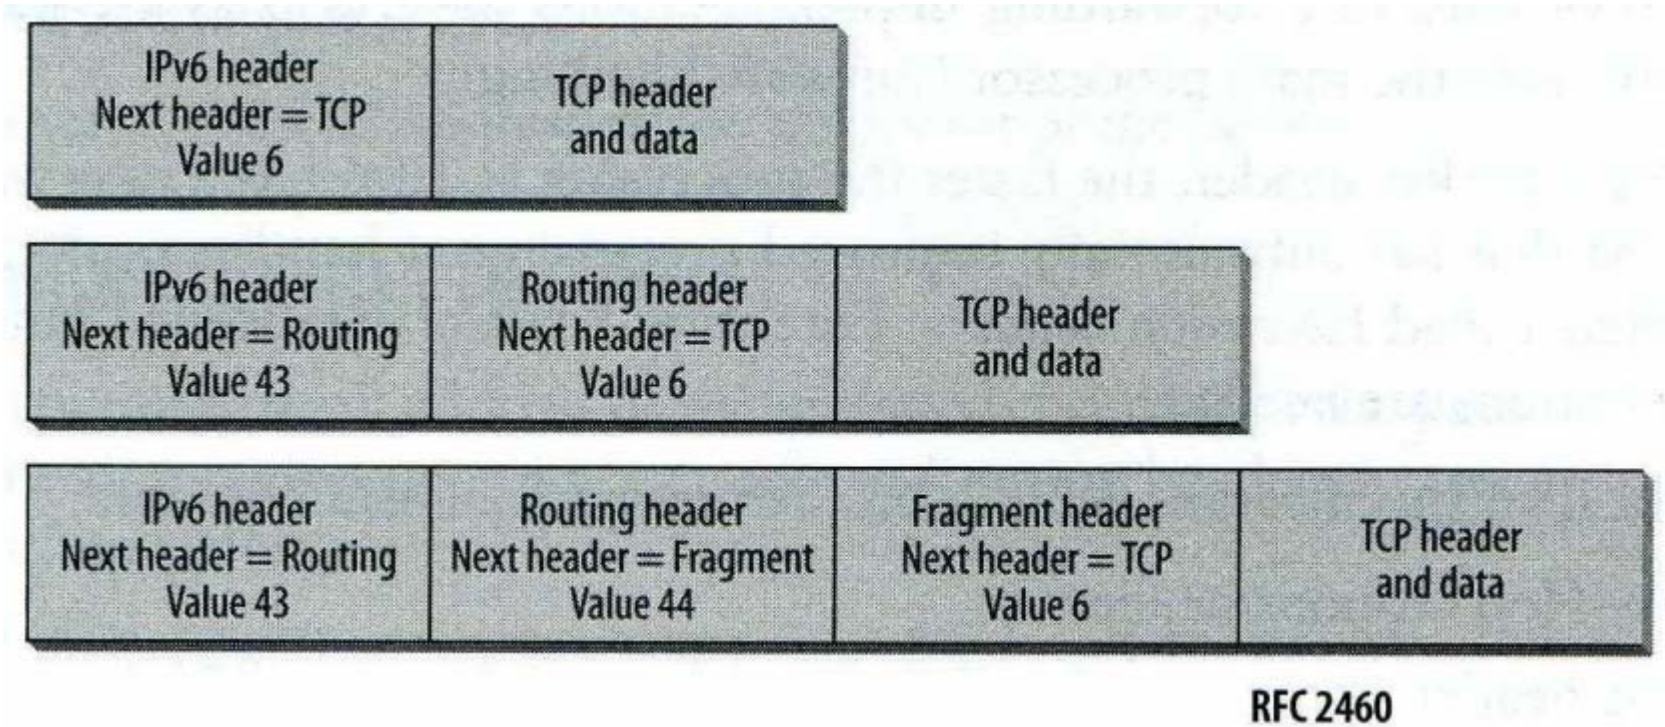

Extension Headers (3)

- Each Extension header is identified by the Next Header field in the preceding header.
- The Extension headers are examined or processed only by the Destination node specified in the IPv6 header.
  - An exception: if the Extension header is a Hop-by-Hop Options header, the information it carries must be examined and processed by every node along the path of the packet.
  - The Hop-by-Hop Options header must immediately follow the IPv6 header.
- Each Extension header is a multiple of 8 bytes long.
- If a processing node cannot identify the value in the Next Header field, it will discard the packet and send an ICMPv6 Parameter Problem message back to the source node.

Extension Headers (4)

- If more than one Extension header is put in a packet, the precedence of these headers should be obeyed (RFC 2460):
  - 1st: IPv6 header.
  - 2nd: Hop-by-Hop Options header.
  - 3rd: Destination Options header (is not the final destination node).
  - 4th: Routing header.
  - 5th: Fragment header.
  - 6th: Authentication header.
  - 7th: Encapsulating Security Payload header.
  - 8th: Destination Options header (is the final destination node).
  - 9th: Upper-Layer header (TCP or UDP).

Extension Headers (5)

- Hop-by-Hop Options Header (number 0):
  - In IPv4, the only way for a router to know whether it needs to examine a datagram is to parse upper layer data in all datagrams.
    - This slows down the routing process in a router.
  - In IPv6, a router processes the datagram only when the Hop-by-Hop Extension header is put in the datagram.
  - The following list describes each field:
    - Next Header (1 byte).
    - Header Extension Length (1 byte).
      - To identifies this header length in 8-byte units.
      - The length calculation does not include the first 8 bytes.
    - Options (variable size).
      - There can be one or more options.
      - The length of the options is variable.

Extension Headers (6)

- Format of the Hop-by-Hop Options header.

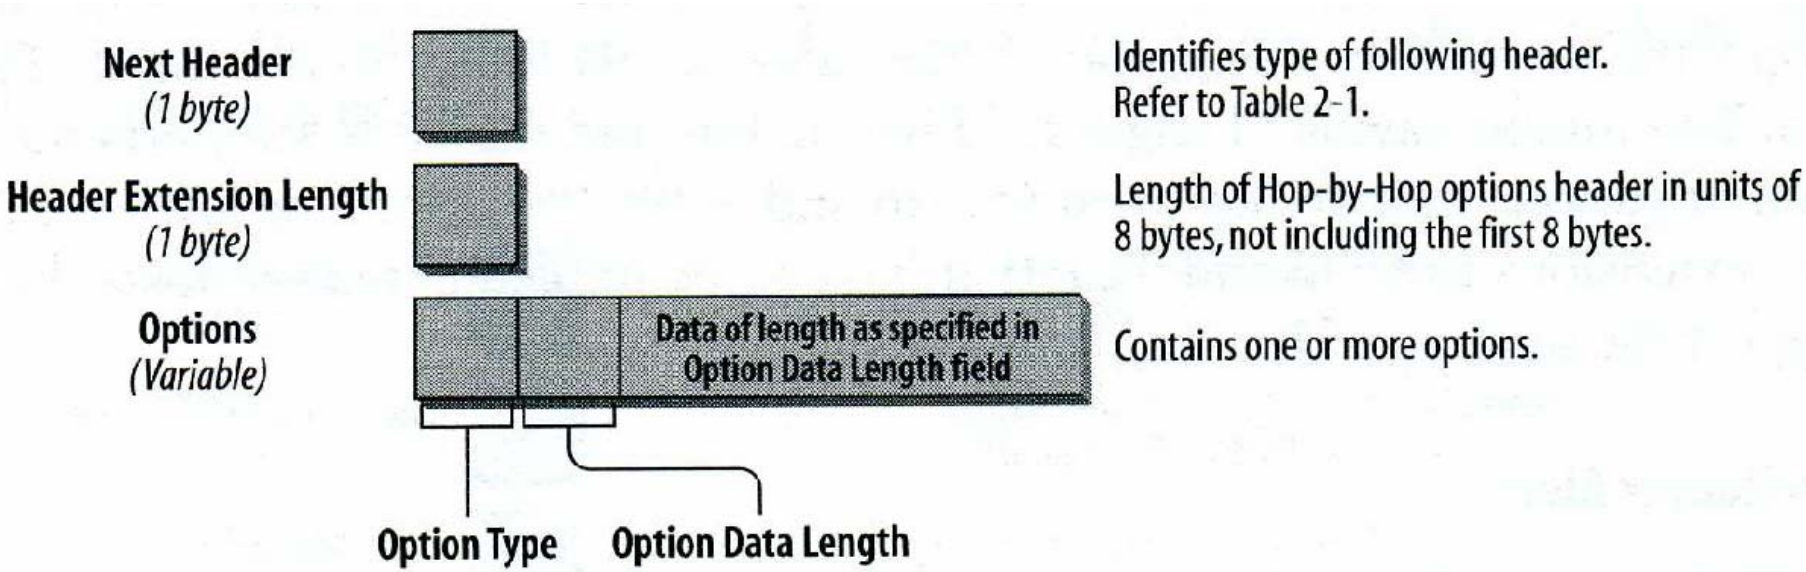

Extension Headers (7)

- The first byte is the Option Type Field.
  - The value of the first two bits in the Option Type Field contains the information about how this option must be treated when the processing node does not recognize the option.
  - 00: skip and continue processing.
  - 01: discard the packet.
  - 10: discard the packet and send ICMP Parameter Problem message (with Code 2) to the packet's source node, pointing to the unrecognized option type.
  - 11: discard the packet and send ICMP Parameter Problem message (with Code 2) to the packet's source address, but the destination must be not a multicast address.

Extension Headers (8)

- Routing Header (number 43):
  - It is used to give a list of one or more intermediate nodes, which should be visited on the packet's path to its destination.
  - In IPv4, this is called the Loose Source & Record Route option (in option of IPv4 header).
  - The following list describes each field:
    - Next Header (1 byte).
    - Header Extension Length (1 byte).
      - To identifies this header length in 8-byte units.
      - The length calculation does not include the first 8 bytes.
    - Routing Type (1 byte).
      - RFC 2460 describes Routing Type zero.
    - Segments Left (1 byte).
      - This field identifies how many nodes are left to be visited before the packet reaches its final destination.

Extension Headers (9)

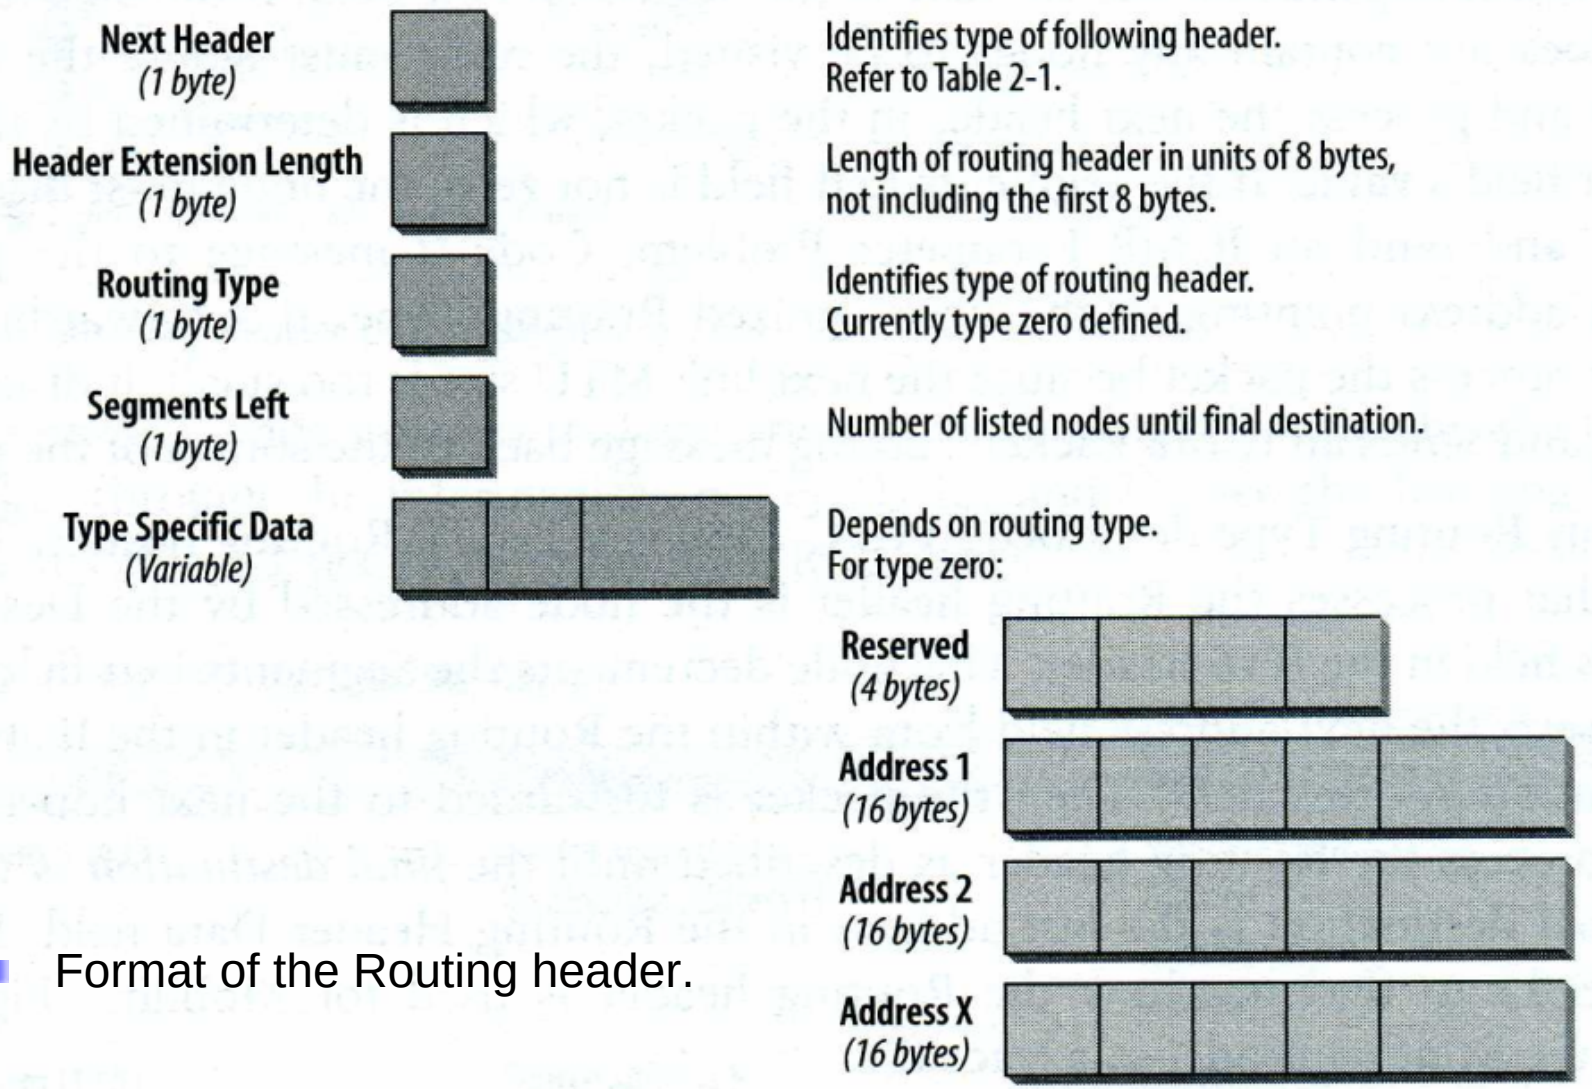

- Format of the Routing header.

Extension Headers (10)

- If a processing node cannot identify a Routing Type value, then it will process the content of the Segments Left field.
  - If the Segments Left field does not contain any nodes to be visited: then the node ignores the Routing header and process the next extension header.
  - If the Segments Left field is not zero: the node must discard the packet and send an ICMP Parameter Problem message (with Code 0) to the source node.
- If a forwarding node cannot process the packet because the next link MTU size is too small, then it will discard the packet and sends an ICMP Packet Too Big message back to the source node.

❖ Type-Specific Data (Variable-length).

- The length of this field depends on the Routing Type.
- It records the visited intermediate nodes & destination node.

Extension Headers (11)

- In Routing Type Zero (RFC 2460):
  - The first visited node (D1) is the node addressed by the Destination address field in the IPv6 header.
  - This node D1 decrements the Segments Left field by 1, & exchange the next address in the Routing header with the Destination address in the IPv6 header.
  - Then the packet is forwarded to the next hop that will again process the Routing header as the above, until the final destination is reached.
  - The final destination is the last address in the Routing Header Data field.

Extension Headers (12)

- Routing header in a trace file.

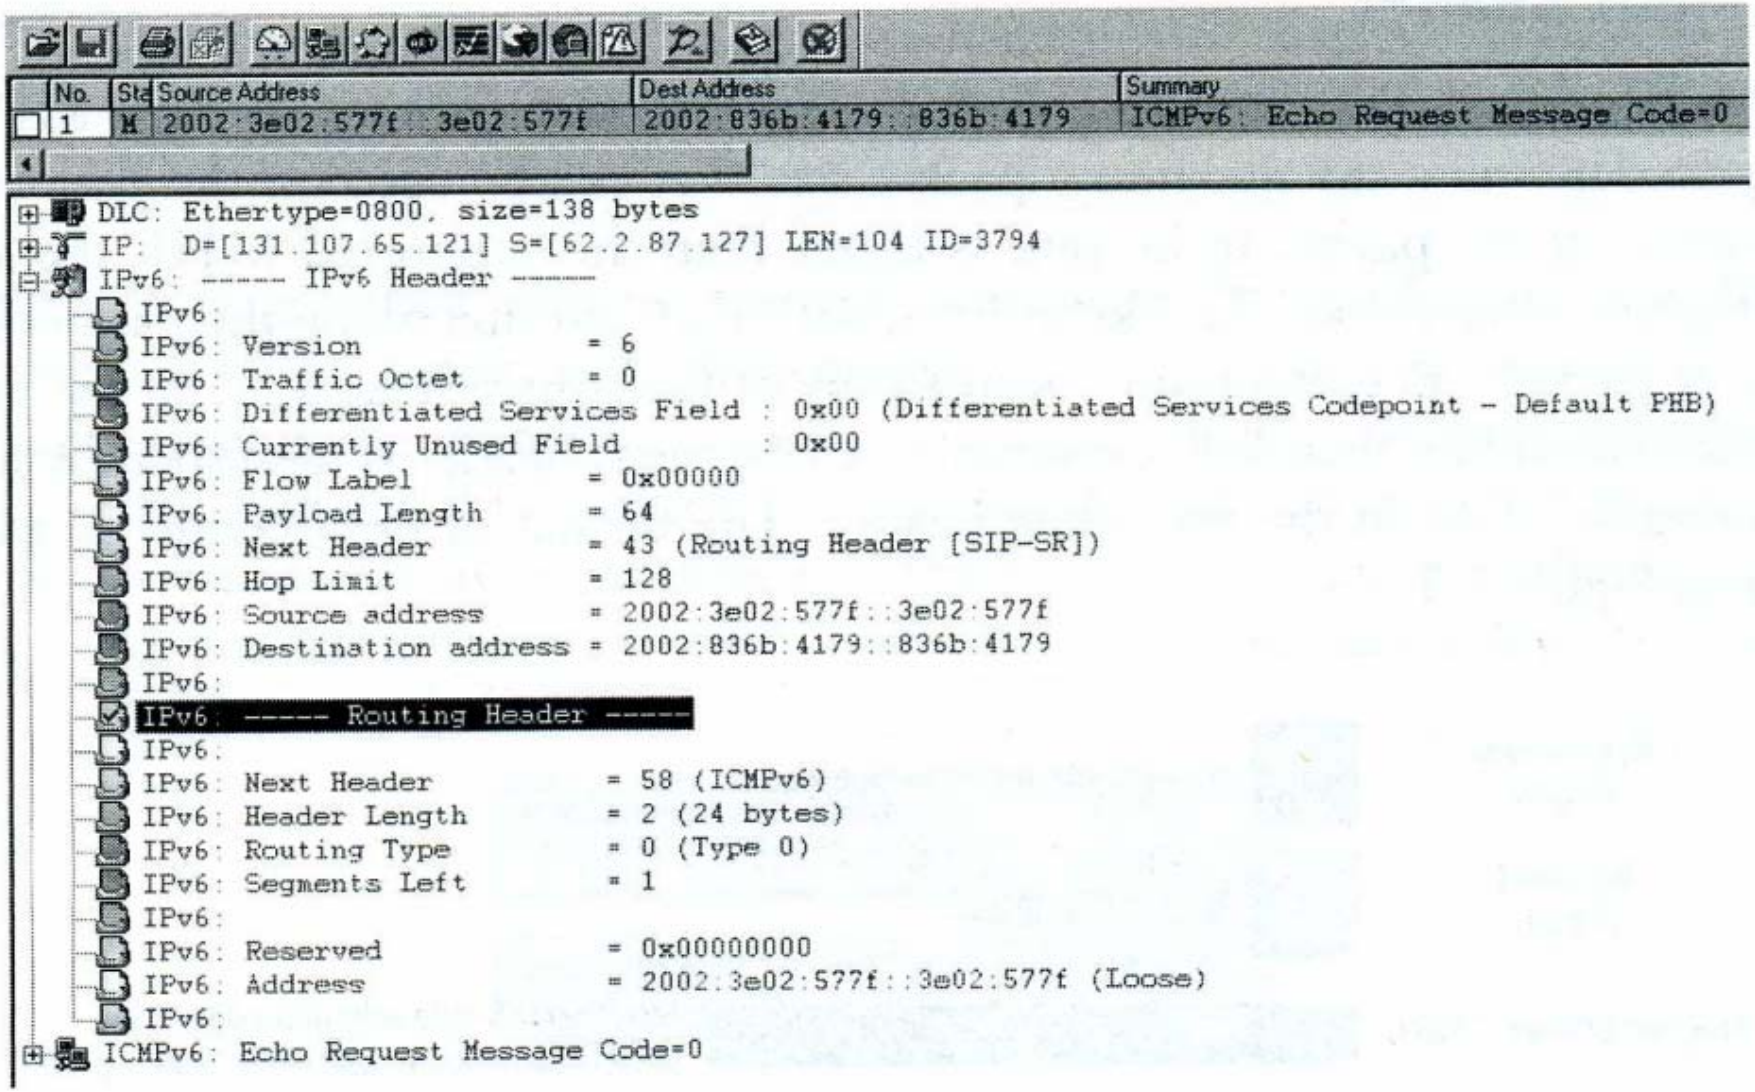

No.	Sta	Source Address	Dest Address	Summary
1	M	2002:3e02:577f::3e02:577f	2002:836b:4179::836b:4179	ICMPv6: Echo Request Message Code=0

```
DLC: Ethertype=0800, size=138 bytes
IP:  D=[131.107.65.121] S=[62.2.87.127] LEN=104 ID=3794
IPv6: ------ IPv6 Header ------
IPv6:
IPv6: Version                = 6
IPv6: Traffic Octet          = 0
IPv6: Differentiated Services Field : 0x00 (Differentiated Services Codepoint - Default PHB)
IPv6: Currently Unused Field       : 0x00
IPv6: Flow Label             = 0x00000
IPv6: Payload Length         = 64
IPv6: Next Header            = 43 (Routing Header [SIP-SR])
IPv6: Hop Limit              = 128
IPv6: Source address         = 2002:3e02:577f::3e02:577f
IPv6: Destination address    = 2002:836b:4179::836b:4179
IPv6:
IPv6: ----- Routing Header -----
IPv6:
IPv6: Next Header            = 58 (ICMPv6)
IPv6: Header Length          = 2 (24 bytes)
IPv6: Routing Type           = 0 (Type 0)
IPv6: Segments Left          = 1
IPv6:
IPv6: Reserved               = 0x00000000
IPv6: Address                = 2002:3e02:577f::3e02:577f (Loose)
IPv6:
ICMPv6: Echo Request Message Code=0
```

Extension Headers (13)

- Processing the routing header.

	IPv6 header	Routing header
Packet from S to I1	Source address S Destination address I1	Segments Left 3 Address (1) = I2 Address (2) = I3 Address (3) = D
Packet from I1 to I2	Source address S Destination address I2	Segments Left 2 Address (1) = I1 Address (2) = I3 Address (3) = D
Packet from I2 to I3	Source address S Destination address I3	Segments Left = 1 Address (1) = I1 Address (2) = I2 Address (3) = D
Packet from I3 to D	Source address S Destination address D	Segments Left = 0 Address (1) = I1 Address (2) = I2 Address (3) = I3

Extension Headers (14)

- Fragment Header (number 44):
  - In IPv4, a big packet can be fragmented by a router, but it does not be done in IPv6 router.
  - An IPv6 host, which wants to send a packet to an IPv6 destination, uses Path MTU (PMTU) discovery to determine the maximum packet size that can be used on the path to that destination.
  - If the packet to be sent is larger than the supported PMTU, the source host fragments the packet.
  - The following list describes each field:
    - Next Header (1 byte).
    - Reserved (1 byte).
    - Fragment Offset (13 bits).
      - The offset is in 8-byte units of the data in this packet (because the packet is fragmented).
      - The offset refers to the start of the data in this packet.

Extension Headers (15)

- Format of the Fragment header.

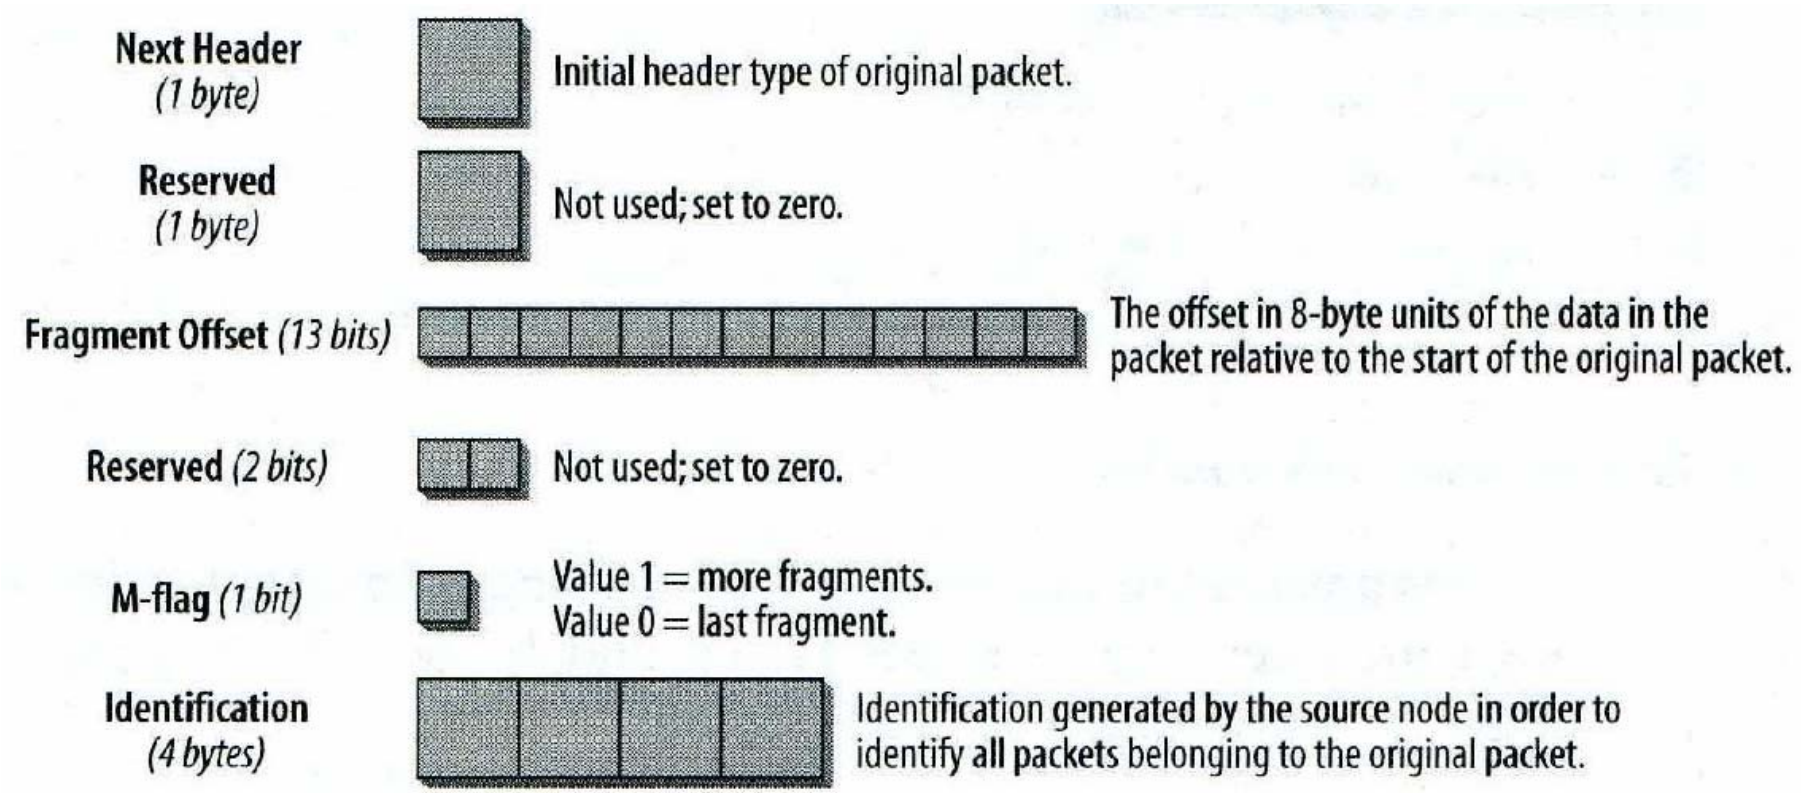

Extension Headers (16)

- Reserved (2 byte).
- M-Flag (1 bit).
  - Value 1: indicates more fragments.
  - Value 0: indicates the last fragment.
- Identification (4 Bytes).
  - Sequence number, which is generated by the source host in order to identify all packets belonging to the original big packet (be fragmented).

- The initial unfragmented packet is the original packet. It has an unfragmentable part:
  - The IPv6 header.
  - Some Extension headers: they must be processed by nodes along the path (i.e., Hop-by-Hop Options).
- The fragmentable part of the original packet consists of:
  - Some Extension headers: they need only to be processed by the final destination.
  - The Upper-Layer headers, & payload data.

Extension Headers (17)

- Fragmentation with IPv6.

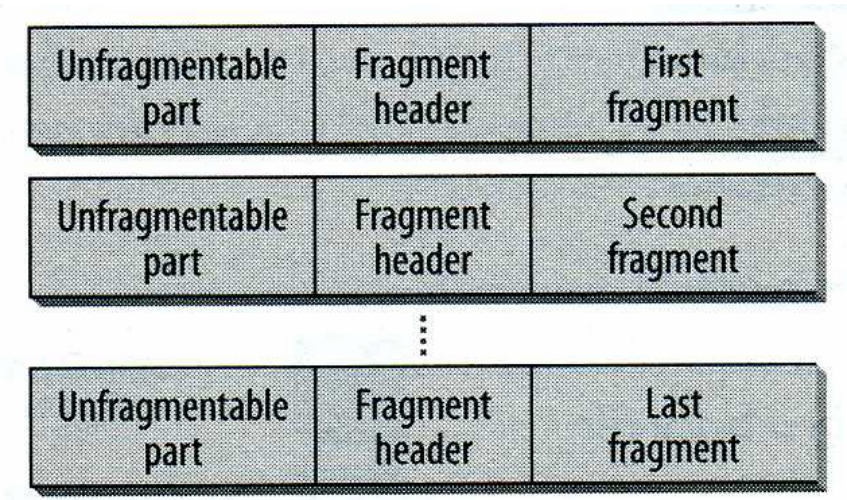

Extension Headers (18)

- The unfragmentable part of the original packet appears in every fragmented packet,
  - Then the following is the Fragmentation header, & the fragmentable part.
- The destination node collects all the fragments & reassembles them.
  - The fragments must have the same Source and Destination addresses & the same identification value in order to be reassembled.
- If all fragments do not arrive at the destination node within 60 seconds after the first fragment, the destination will discard all fragments.
  - If the destination node has received the first fragment (offset = 0), then it will send back an ICMPv6 Fragment Reassembly Time Exceeded message to the source node.

Extension Headers (19)

- Fragment header in a trace file.

No.	Sta	Source Address	Dest Address	Summary
1	M	fe80::202:b3ff:fe1e:8329	fe80::2a0:24ff:fec5:3256	ICMPv6: Echo Request Message Code=0
2		fe80::202:b3ff:fe1e:8329	fe80::2a0:24ff:fec5:3256	IPv6: Continuation of frame 1;656 Bytes of data IPv6: Flow=0x000000
3		fe80::2a0:24ff:fec5:3256	fe80::202:b3ff:fe1e:8329	ICMPv6: Echo Reply Message Code=0
4		fe80::2a0:24ff:fec5:3256	fe80::202:b3ff:fe1e:8329	IPv6: Continuation of frame 3;656 Bytes of data IPv6: Flow=0x000000

```
DLC: Ethertype=86DD, size=1510 bytes
IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version                = 6
IPv6: Traffic Octet          = 0
IPv6: Differentiated Services Field : 0x00 (Differentiated Services Codepoint - Default PHB)
IPv6: Currently Unused Field       : 0x00
IPv6: Flow Label             = 0x00000
IPv6: Payload Length         = 1456
IPv6: Next Header            = 44 (Fragment Header [SIP-FRAG])
IPv6: Hop Limit              = 128
IPv6: Source address         = fe80::202:b3ff:fe1e:8329
IPv6: Destination address    = fe80::2a0:24ff:fec5:3256
IPv6:
IPv6: ----- Fragment Header -----
IPv6:
IPv6: Next Header     = 58 (ICMPv6)
IPv6: Reserved        = 0
IPv6: Offset          = 0000
IPv6: .... 00  = Reserved = 0
IPv6: .... ...1 = More fragments
IPv6: Identification = 1
IPv6:
ICMPv6: Echo Request Message Code=0
```

Extension Headers (20)

- The last packet in the fragment set.

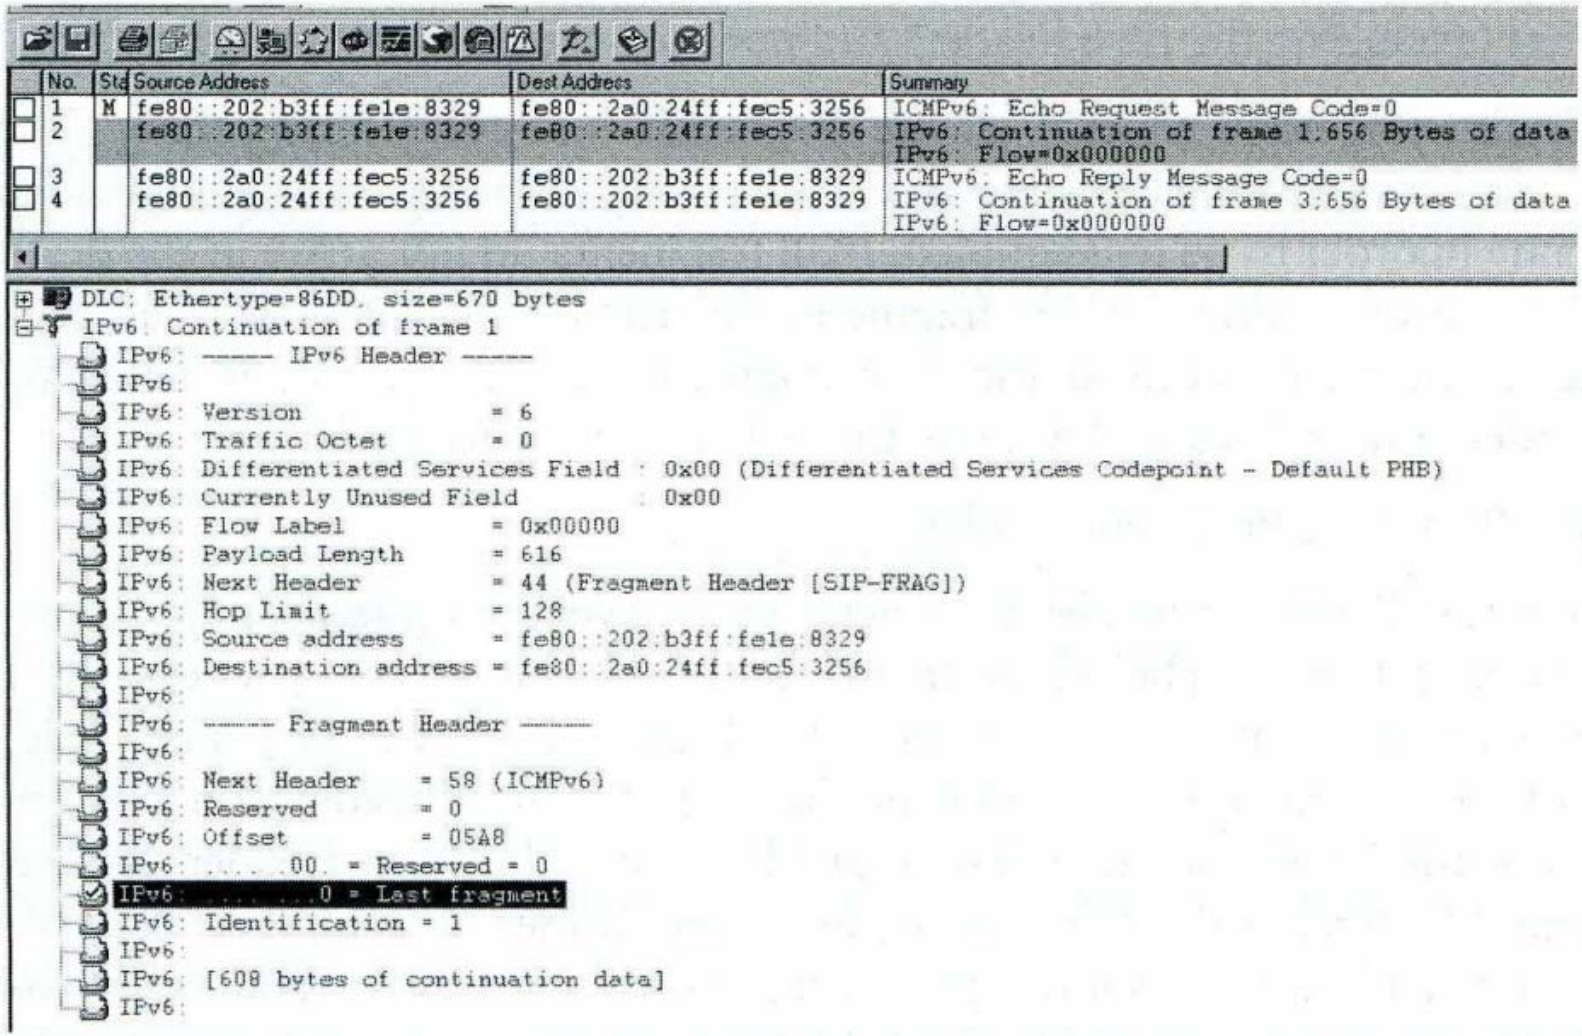

No.	Sta	Source Address	Dest Address	Summary
1	M	fe80::202:b3ff:fe1e:8329	fe80::2a0:24ff:fec5:3256	ICMPv6: Echo Request Message Code=0
2		fe80::202:b3ff:fe1e:8329	fe80::2a0:24ff:fec5:3256	IPv6: Continuation of frame 1;656 Bytes of data IPv6: Flow=0x000000
3		fe80::2a0:24ff:fec5:3256	fe80::202:b3ff:fe1e:8329	ICMPv6: Echo Reply Message Code=0
4		fe80::2a0:24ff:fec5:3256	fe80::202:b3ff:fe1e:8329	IPv6: Continuation of frame 3;656 Bytes of data IPv6: Flow=0x000000

```
DLC: Ethertype=86DD, size=670 bytes
IPv6: Continuation of frame 1
IPv6: ----- IPv6 Header -----
IPv6:
IPv6: Version              = 6
IPv6: Traffic Octet        = 0
IPv6: Differentiated Services Field : 0x00 (Differentiated Services Codepoint - Default PHB)
IPv6: Currently Unused Field        : 0x00
IPv6: Flow Label           = 0x00000
IPv6: Payload Length       = 616
IPv6: Next Header          = 44 (Fragment Header [SIP-FRAG])
IPv6: Hop Limit            = 128
IPv6: Source address       = fe80::202:b3ff:fe1e:8329
IPv6: Destination address  = fe80::2a0:24ff:fec5:3256
IPv6:
IPv6: ------- Fragment Header -------
IPv6:
IPv6: Next Header     = 58 (ICMPv6)
IPv6: Reserved        = 0
IPv6: Offset          = 05A8
IPv6: .... .00. = Reserved = 0
IPv6: .... ...0 = Last fragment
IPv6: Identification = 1
IPv6:
IPv6: [608 bytes of continuation data]
IPv6:
```

Extension Headers (21)

- Destination Options Header (number 60):
  - A Destination Options header carries the information, which is examined by the destination node only.
  - The following list describes each field:
    - Next Header (1 byte).
    - Header Extension Length (1 byte).
      - To identifies this header length in 8-byte units.
      - The length calculation does not include the first 8 bytes.
    - Options (variable size).
      - There can be one or more options.

Extension Headers (22)

- Format of the Destination Options header.