The Essentials of Computer Org. and Arch. 3rd edition Null and Lobur
ISBN: 978-1-4496-0006-8

THE IP VERSION 4 DATAGRAM HEADER

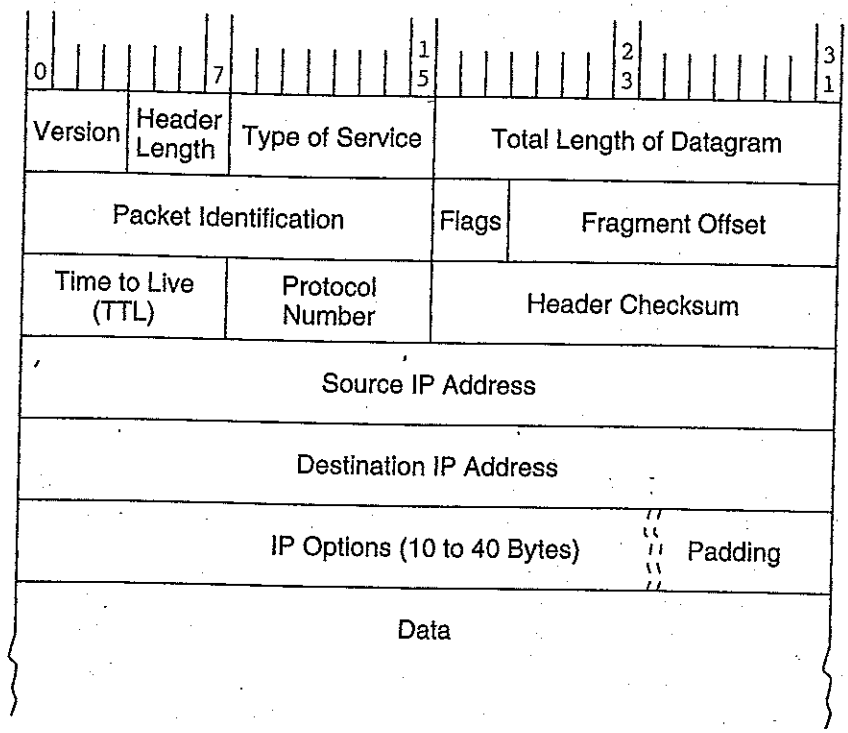

Each IPv4 datagram must contain at least 40 bytes, which include a 24-byte header as shown above. The horizontal rows represent 32-bit words. Upon inspection of the figure, you can see, for example, that the Type of Service field occupies bits 8 through 15, while the Packet Identification field occupies bits 32 through 47 of the header. The Padding field shown as the last field of the header ensures that the data that follows the header starts on an even 32-bit boundary. The Padding always contains zeroes. The other fields in the IPv4 header are:

- Version—Specifies the IP protocol version being used. The version number tells all of the hardware along the way the length of the datagram and what content to expect in its header fields. For IPv4, this field is always 0100 (because $0100_2 = 4_{10}$).
- Header Length—Gives the length of the header in 32-bit words. The size of the IP header is variable, depending on the value of the IP Options fields, but the minimum value for a correct header is 5.
- Type of Service—Controls the priority that the datagram is given by intermediate nodes. Values can range from "routine" (000) to "critical" (101). Network control datagrams are indicated with 110 and 111.
- Total Length—Gives the length of the entire IP datagram in bytes. As you can see by the layout above, 2 bytes are reserved for this purpose. Hence, the largest allowable IP datagram is $2^{16} - 1$, or 65,535.

- Packet ID—Each datagram is assigned a serial number as it is placed on the network. The combination of Host ID and Packet ID uniquely identifies each IP datagram in existence at any time in the world.
- Flags—Specify whether the datagram may be fragmented (broken into smaller datagrams) by intermediate nodes. IP networks must be able to handle datagrams of at least 576 bytes. Most IP networks can deal with packets that are about 8KB long. With the "Don't Fragment" bit set, an 8KB datagram will not be routed over a network that says it can handle only 2KB packets, for example.
- Fragment Offset—Indicates the location of a fragment within a certain datagram. That is, it tells which part of the datagram the fragment came from.
- Time to Live (TTL)—TTL was originally intended to measure the number of seconds for which the datagram would remain valid. Should a datagram get caught in a routing loop, the TTL would (theoretically) expire before the datagram could contribute to a congestion problem. In practice, the TTL field is decremented each time it passes through an intermediate network node, so this field does not really measure the number of seconds that a packet lives, but the number of hops it is allowed before it reaches its destination.
- Protocol Number—Indicates which higher-layer protocol is sending the data that follows the header. Some of the important values for this field are:

  0 = Reserved

  1 = Internet Control Message Protocol (ICMP)

  6 = Transmission Control Protocol (TCP)

  17 = User Datagram Protocol (UDP)

  TCP is described in Section 12.5.3.

- Header Checksum—This field is calculated by first calculating the one's complement sum of all 16-bit words in the header, and then taking the one's complement of this sum, with the checksum field itself originally set to all zeroes. The one's complement sum is the arithmetic sum of two of the words with the (seventeenth) carry bit added to the lowest bit position of the sum. (See Section 2.4.2.) For example, 11110011 + 10011010 = 110001101 = 10001110 using one's complement arithmetic. What this means is that if we have an IP datagram of the form shown to the right, each w_i is a 16-bit word in the IP datagram. The complete checksum would be computed over two 16-bit words at a time: $w_1 + w_2 = S_1$; $S_1 + w_3 = S_2$; . . . $S_k + w_{k-2} = S_{k+1}$.

w_1	w_2
w_3	w_4
. . .	. . .
. . .	. . .
w_{n-1}	w_n

- Source and Destination Addresses—Tell where the datagram is going. We have much more to say about these 32-bit fields in Section 12.5.2.
- IP Options—Provides diagnostic information and routing controls. IP Options are, well, optional.