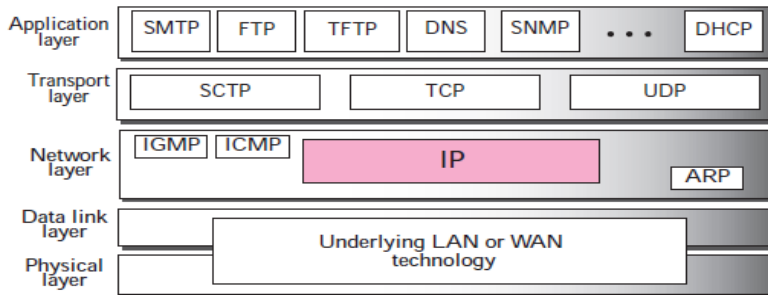




Internet Protocol Version 4 (IPv4)

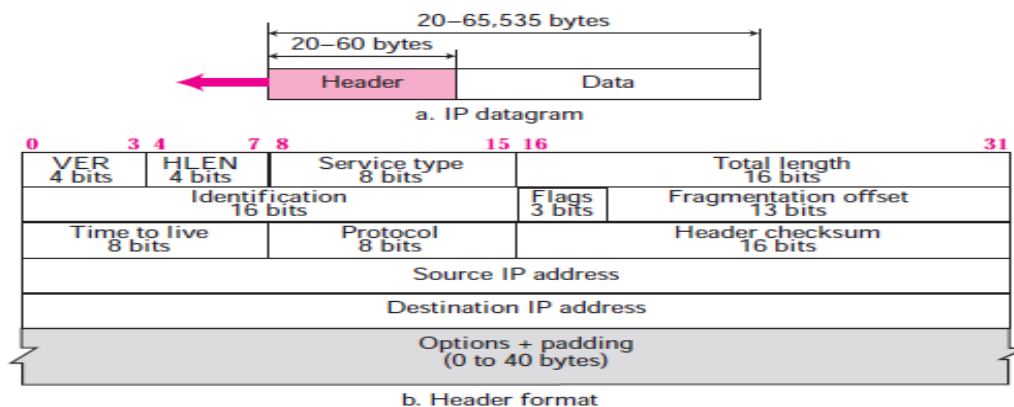
The **Internet Protocol (IP)** is the transmission mechanism used by the TCP/IP protocols at the network layer. The following Figure shows the position of IP in the suite.



- IP is an unreliable and connectionless datagram protocol (**a best-effort delivery service**).
- The term *best-effort* means that IP packets can be corrupted, lost, arrive out of order, or delayed and may create congestion for the network.
- If reliability is important, IP must be paired with a reliable protocol such as TCP.
- IP is also a connectionless protocol for a packet switching network that uses the datagram approach. This means that each datagram is handled independently, and each datagram can follow a different route to the destination.

DATAGRAMS

Packets in the network (internet) layer are called **datagrams**. A datagram is a variable-length packet consisting of two parts: header and data. The header is 20 to 60 bytes in length and contains information essential to routing and delivery. It is customary in TCP/IP to show the header in 4-byte sections. The following Figure shows the *IP datagram* format.



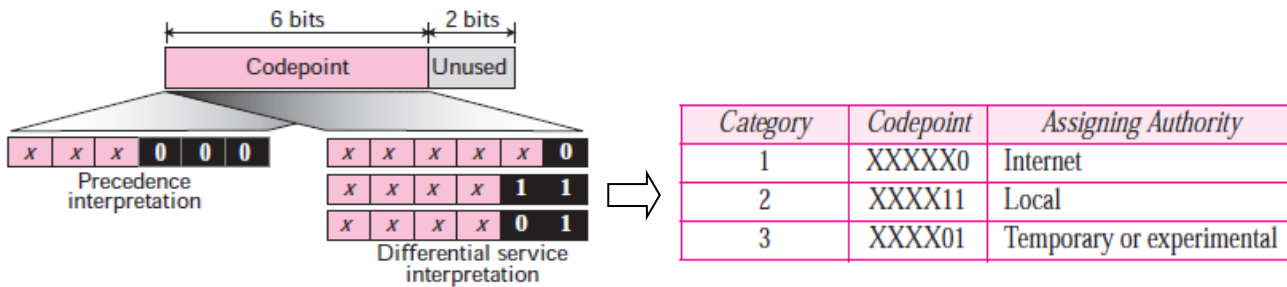
A brief description of each field is in order.

- **Version (VER):** This 4-bit field defines the version of the IP protocol. If the machine is using some other version of IP, the datagram is discarded rather than interpreted incorrectly.
- **Header length (HLEN):** This 4-bit field defines the total length of the datagram header in 4-byte words. This field is needed because the length of the header is variable (between 20 and 60 bytes).



When there are no options, the **header length** is 20 bytes, and the value of this field is 5 (5* 4= 20).
When the option field is at its maximum size, the value of this field is 15 (15*4=60).

- **Service type:** In the original design of IP header, this field was referred to as type of service (**TOS**), which defined how the datagram should be handled. Part of the field was used to define the **precedence** of the datagram; the rest defined the type of service (**low delay, high throughput**, and so on). IETF has changed the interpretation of this 8-bit field. This field now defines a set of differentiated services. The new interpretation is shown in the following Figure.




In this interpretation, the first 6 bits make up the **codepoint subfield** and the last 2 bits are not used. The codepoint subfield can be used in two different ways.

- When the 3 right-most bits are 0s, the 3 left-most bits are interpreted the same as the precedence bits in the service type interpretation. In other words, it is compatible with the old interpretation. The precedence defines the eight-level priority of the datagram (0 to 7) in issues such as congestion. If a router is congested and needs to discard some datagrams, those datagrams with lowest precedence are discarded first.
 - When the 3 right-most bits are not all 0s, the 6 bits define 56 (64-8) services based on the priority assignment by the Internet or local authorities according to the above Table. The first category contains 24 service types; the second and the third each contain 16.
- **Total length:** This is a 16-bit field that defines the total length (header plus data) of the IP datagram in bytes.

$$\text{Length of data} = \text{total length} - \text{header length}$$

The total length field defines the total length of the datagram including the header.

Largest size of datagram is 2^{16} (65.536) bytes. Max header size is 60 bytes. Maximum Transfer Unit (**MTU**) of Data link layer may restrict this size.




Q: why we need this field anyway?

- **Identification.** This field is used in fragmentation.
- **Flags.** This field is used in fragmentation.
- **Fragmentation offset.** This field is used in fragmentation (discussed in the next section).



- **Time to live.** A datagram has a limited lifetime in its travel through an internet. this field is mostly used to control the maximum number of hops (routers) visited by the datagram. When a source host sends the datagram, it stores a number in this field. This value is approximately two times the maximum number of routes between any two hosts. Each router that processes the datagram decrements this number by one. If this value, after being decremented, is zero, the router discards the datagram.
 - This field is needed because routing tables in the Internet can become corrupted. This field limits the lifetime of a datagram.
 - Another use of this field is to intentionally limit the journey of the packet. For example, if the source wants to confine the packet to the local network, it can store 1 in this field. When the packet arrives at the first router, this value is decremented to 0, and the datagram is discarded.
- **Protocol.** This 8-bit field defines the higher-level protocol that uses the services of the IP layer. An IP datagram can encapsulate data from several higher level protocols such as TCP, UDP, ICMP, and IGMP. This field specifies the final destination protocol to which the IP datagram should be delivered.



Value	Protocol	Value	Protocol
1	ICMP	17	UDP
2	ICMP	89	OSPF
6	TCP		

- **Checksum.** Checksum only covers the header and not data
- **Source address:** This 32-bit field defines the IP address of the source.
- **Destination address:** This 32-bit field defines the IP address of the destination.

Example 1:

An IP packet has arrived with the first 8 bits as shown: 01000010

The receiver discards the packet. Why?

Solution:

There is an error in this packet. The 4 left-most bits (0100) show the version, which is correct. The next 4 bits (0010) show the wrong header length ($2 * 4 = 8$). The minimum number of bytes in the header must be 20. The packet has been corrupted in transmission.

Example 2:

In an IP packet, the value of HLEN is 1000 in binary. How many bytes of options are being carried by this packet?

Solution:

The HLEN value is 8, which means the total number of bytes in the header is $8 * 4$ or 32 bytes. The first 20 bytes are the base header, the next 12 bytes are the options.



Example 3

In an IP packet, the value of HLEN is 5_{16} and the value of the total length field is 0028_{16} . How many bytes of data are being carried by this packet?

Solution:

The HLEN value is 5, which means the total number of bytes in the header is $5 * 4$ or 20 bytes (no options). The total length is 40 bytes, which means the packet is carrying 20 bytes of data ($40 - 20$).

Example 4:

An IP packet has arrived with the first few hexadecimal digits as shown: $45000028000100000102 \dots$

How many hops can this packet travel before being dropped? The data belong to what upper layer protocol?

Solution:


To find the time-to-live field, we skip 8 bytes (16 hexadecimal digits). The time-to-live field is the ninth byte, which is 01. This means the packet can travel only one hop. The protocol field is the next byte (02), which means that the upper layer protocol is IGMP.

Fragmentation

- A datagram can travel through different networks. Each router decapsulates the IP datagram from the frame it receives, processes it, and then encapsulates it in another frame.
- The format and size of the sent frame depend on the protocol used by the physical network through which the frame is going to travel. For example, if a router connects a LAN to a WAN, it receives a frame in the LAN format and sends a frame in the WAN format.

Maximum Transfer Unit (MTU)

- Each data link layer protocol has its own frame format in most protocols. The total size of the datagram must be less than the maximum size of the data field, which is defined by the restrictions imposed by the hardware and software used in the network.
- In order to make the IP protocol independent of the physical network, the designers decided to make the maximum length of the IP datagram equal to 65,535 bytes. This makes transmission more efficient if we use a protocol with an MTU of this size.



The value for the Ethernet LAN is 1500 bytes, for FDDI LAN is 4352 bytes, and for PPP is 296 bytes.


- **Only data in a datagram is fragmented.**
- The source usually does not fragment the IP packet.
- When a datagram is fragmented, each fragment has its own header with most of the fields repeated, but some changed. A datagram can be fragmented several times before it reaches the final destination.



- A datagram can be fragmented by the source host or any router in the path. The reassembly of the datagram, however, is done only by the destination host because each fragment becomes an independent datagram.
- When a datagram is fragmented, required parts of the header must be copied by all fragments. The host or router that fragments a datagram must change the values of three fields: **flags, fragmentation offset, and total length**. The rest of the fields must be copied. Of course, the value of the checksum must be recalculated regardless of fragmentation.

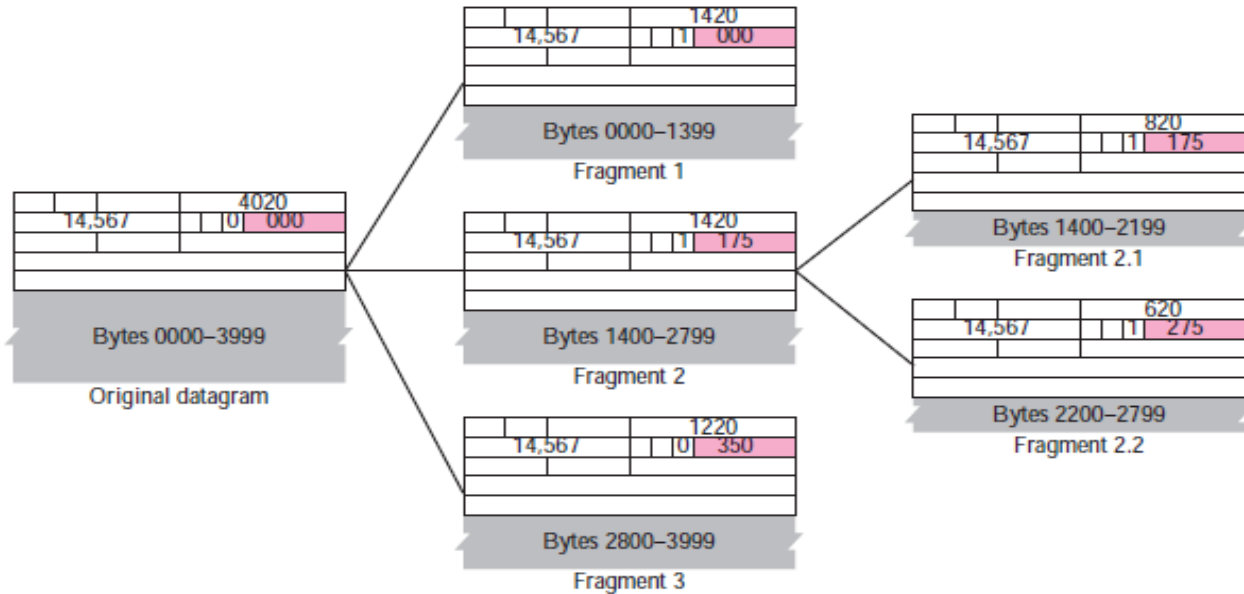
Fields Related to Fragmentation

- The fields that are related to fragmentation and reassembly of an IP datagram are the **identification, flags, and fragmentation offset fields**.
 - **Identification:** This 16-bit field identifies a datagram originating from the source host. The combination of the identification and source IP address must uniquely define a datagram as it leaves the source host. To guarantee uniqueness, the IP protocol uses a counter to label the datagrams. When the IP protocol sends a datagram, it copies the current value of the counter to the identification field and increments the counter by one. When a datagram is fragmented, the value in the identification field is copied into all fragments. The identification number helps the destination in reassembling the datagram.
 - **Flags:** This is a three-bit field. The first bit is reserved (not used). The **second bit** is called the **do not fragment** bit. **If** its value is **1**, the machine must not fragment the datagram. If it cannot pass the datagram through any available physical network, it discards the datagram and sends an ICMP error message to the source host. **If** its value is **0**, the datagram can be fragmented if necessary. The **third bit** is called the **more fragment** bit. **If** its value is **1**, it means the datagram is not the last fragment. **If** its value is **0**, it means this is the last or only fragment.
 - **Fragmentation offset:** This 13-bit field shows the relative position of this fragment with respect to the whole datagram. It is the offset of the data in the original datagram **measured in units of 8** bytes. The following Figure shows a datagram with a data size of 4000 bytes fragmented into three fragments.





Remember that the value of the offset is measured in units of **8 bytes**. This is done because the length of the offset field is only **13 bits** long and cannot represent a sequence of bytes greater than **8191**. This forces hosts or routers that fragment datagrams to choose the size of each fragment so that the first byte number is divisible by 8.



- It is obvious that even if each fragment follows a different path and arrives out of order, the final destination host can reassemble the original datagram from the fragments received (if none of them is lost) using the following strategy:
 1. The first fragment has an offset field value of zero.
 2. Divide the length of the first fragment by 8. The second fragment has an offset value equal to that result.
 3. Divide the total length of the first and second fragment by 8. The third fragment has an offset value equal to that result.
 4. Continue the process. The last fragment has a *more bit* value of 0.

Example 5:

A packet has arrived with an **M** bit value of **0**. Is this the first fragment, the last fragment, or a middle fragment? **Do we know if the packet was fragmented?**

Solution

If the **M** bit is **0**, it means that there are no more fragments; the fragment is the last one. However, we cannot say if the original packet was fragmented or not. A nonfragmented packet is considered the last fragment.

Example 6:

A packet has arrived with an **M** bit value of **1**. Is this the first fragment, the last fragment, or a middle fragment? **Do we know if the packet was fragmented?**

Solution



If the **M** bit is **1**, it means that there is at least one more fragment. This fragment can be the first one or a middle one, but not the last one. We don't know if it is the first one or a middle one; we need more information (the value of the fragmentation offset). See also the next example.

Example 7:

A packet has arrived with an **M** bit value of **1** and a **fragmentation offset** value of **zero**. **Is this the first fragment, the last fragment, or a middle fragment?**

Solution

Because the **M** bit is **1**, it is either the first fragment or a middle one. Because the offset value is 0, it is the first fragment.

Example 8:

A packet has arrived in which the **offset value is 100**. **What is the number of the first byte? Do we know the number of the last byte?**

Solution

To find the number of the first byte, **we multiply the offset value by 8**. This means that the first byte number is 800. We cannot determine the number of the last byte unless we know the length of the data.

Example 9:

A packet has arrived in which the **offset value is 100**, the value of **HLEN is 5** and the value of **the total length field is 100**. **What is the number of the first byte and the last byte?**

Solution

The **first byte number** is $100 \times 8 = 800$. The **total length** is **100 bytes** and the **header length** is **20 bytes** (5×4), which means that there are 80 bytes in this datagram. If the first byte number is 800, the last byte number must be 879.