Network Layer

Deals with end-to-end communications.

The IP protocol does not guarantee delivery of data and is therefore unreliable and a best-effort type of protocol. All that is checked for errors is the header.

However, there is an error reporting mechanism called ICMP (Internet Control Message Protocol). However, it is only there to inform about problems that have occurred in the IP.

Fragmenting an IP Datagram

Because the Data Link layer often does not support frame sizes that are as large as the IP datagram, it must be sent in pieces.

The fields of the IP header involved in this are:

Identification

Identifies each fragment as part of a larger datagram. All fragments in a datagram are given the same identification number.

Flags

Three different bits:

1. **More** – Specifies whether there is another fragment in the datagram. This is almost always 1. The only times it is 0 is when the last fragment of the datagram is sent or when the datagram is not fragmented.

2. **Do not fragment** – There may be a datagram that we wouldn't want to fragment

3. **Unused**

Fragment Offset

The large data block that is in the datagram will be broken into smaller chunks that are about the same size, if that is possible.

The number of bytes in the fragment must be a number that is divisible by 8. The offset indicates the position of the data in the big datagram.

The offset number is the number of 8-byte chunks included in the current datagram.

Other Fields in the Header

Source and Destination Addresses

32 bits long in IPv4

The way you're informed that something did not reach its destination is through the ICMP message. These messages tell you like "the packet did not reach its destination." Or will explain that the time to live expired.

ICMP Ping

Used to request that

IP Address Classes

Original intent was to distribute IP addresses to companies based on need. Sometime around 1993, they realized that this classification was a mistake.

Class A (1-126)

MSB is always a 0
The next 7 bits are the network id.
The remaining 24 bits are the host id

This was intended for really big organizations. Since the network ID is only 7 bits, only 128 of these Class A addresses are available.

As of 2001, only 80 of these 128 were assigned.

Class B (128-191)

Start off with 10
The next 14 bits are the network id.
The remaining 16 bits are the host id.

Class C (192-223)

Start off with 110
The next 21 bits are for the network id
The remaining 8 bits are the host id.

Classless Inter-domain Routing (CIDR)

The network prefix is a variable number.

128.143.173.144 / 24

The 24 tells us that the first 24 bits are the network id. Whatever is left is the host id.

One reason they switched to this mechanism is that it removes the dependency to the rigid format dictated by the classes. It allows for flexible assignment based on the size of the company.