Best-Effort Delivery

- Because IP is designed to operate over all types of network hardware, the underlying hardware may misbehave.
- As a result, IP datagrams may be lost, duplicated, delayed, delivered out of order, or delivered with corrupted data.
- Higher layers of protocol software are required to handle each of these errors.

IP Encapsulation

- A datagram is encapsulated in a frame for transmission across a physical network.
- The destination address in the frame is the address of the next hop to which the datagram should be sent.
- The address is obtained by translating the IP address of the next hop to an equivalent hardware address.

** An IP datagram encapsulated in a hardware frame. The entire datagram resides in the frame data area. In practice, the frame format used with some technologies includes a frame trailer as well as a frame header.*

Internet Transmission

- When a datagram arrives in a network frame, the receiver extracts the datagram from the frame data area and discards the frame header.

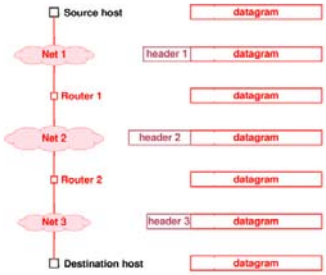

An IP datagram as it appears at each step during a trip across an internet. Whenever it travels across a physical network, the datagram is encapsulated in a frame appropriate to the network.

Maximum Transmission Unit

- A datagram cannot be larger than the MTU of a network over which it is sent.

An example of a router that connects two networks with different MTU values. A frame that travels across Network *1* can contain *1500* octets of data, while a frame that travels across network *2* can contain at most *1000* octets of data.

MTU

- When a router receives a datagram that is larger than the MTU of the network over which it is to be sent, the router divides the datagram into smaller pieces called fragments.
- Each fragment uses the IP datagram format, but carries only part of the data.

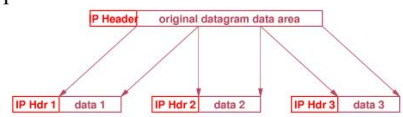

An IP datagram divided into three fragments. Each fragment carries some data from the original datagram, and has an IP header similar to the original datagram.

Fragmentation Example

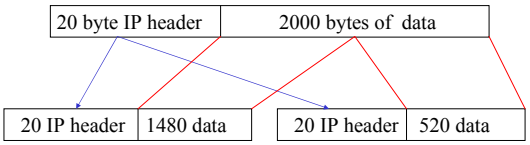

Reassembly

- The process of creating a copy of the original datagram from fragments.

** An example internet in which hosts can generate Datagrams that require fragmentation. Once a datagram has been fragmented, the fragments are forwarded to the final destination, which reassembles them.*

Fragmentation by Intermediates

- Any router between the source and the destination may need to fragment a IP datagram.
- Datagrams are only reassembled at the final destination.

Identification/Flags/Fragment Offset

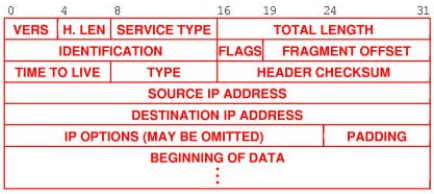

IDENTIFICATION – used to identify all of the fragments of a packet.
FRAGMENT OFFSET – tells a receiver how to order fragments within a datagram.
FLAGS – Indicates the last of a series of fragments.

Fragment Loss

- IP does not guarantee datagram delivery.
- To avoid exhausting memory, IP specifies a maximum time to hold fragments.
- The result of the IP reassembly timer is all or nothing. Either all fragments are received and assembled or completely discarded.

Fragmentation Minimization

Consider sending 1700 bytes

The first network will create 1500 and 200 byte packets. The second network will divide the 1500 byte packet into 1000 and 500 byte packets creating 3 packets in total.

Starting with a 1000 byte packet and a 700 byte packet will produce only 2 packets with no intermediate fragmentation.

Determining Minimum Fragment Size

- An option in the IP header prevents IP from fragmenting a packet.
- An ICMP message is sent back to the source whenever a packet is discarded because it could not be sent due to MTU limits.
- The sender can adjust the packet size to meet all MTU limits.