Internet Architecture and Philosophy

- Conceptually, TCP/IP provides three sets of services to the user:

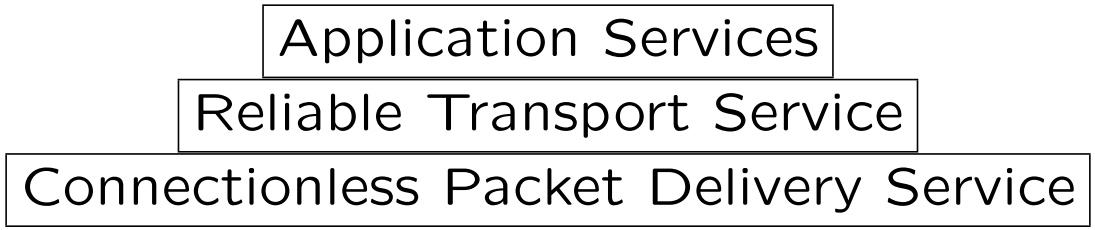

- The underlying service provided is "connectionless".
  - Internet "datagrams" are simply routed independantly between hosts
- The underlying service provided is "unreliable"
  - Any particular datagram might be lost, damaged, delayed, or re-ordered

Internet Protocol (IP)

- Basic unit of Internet transfer
- Embodies a connectionless packet delivery service
- Analogous to physical network packet
- Composed of
  - Header that contains source and destination
- Internet addresses, datagram type field, etc.
  - Data area that contains data being carried

IP Datagram Format

- An IP data consists of an IP header of (at least) 20 octets followed by the encapsulated IP data

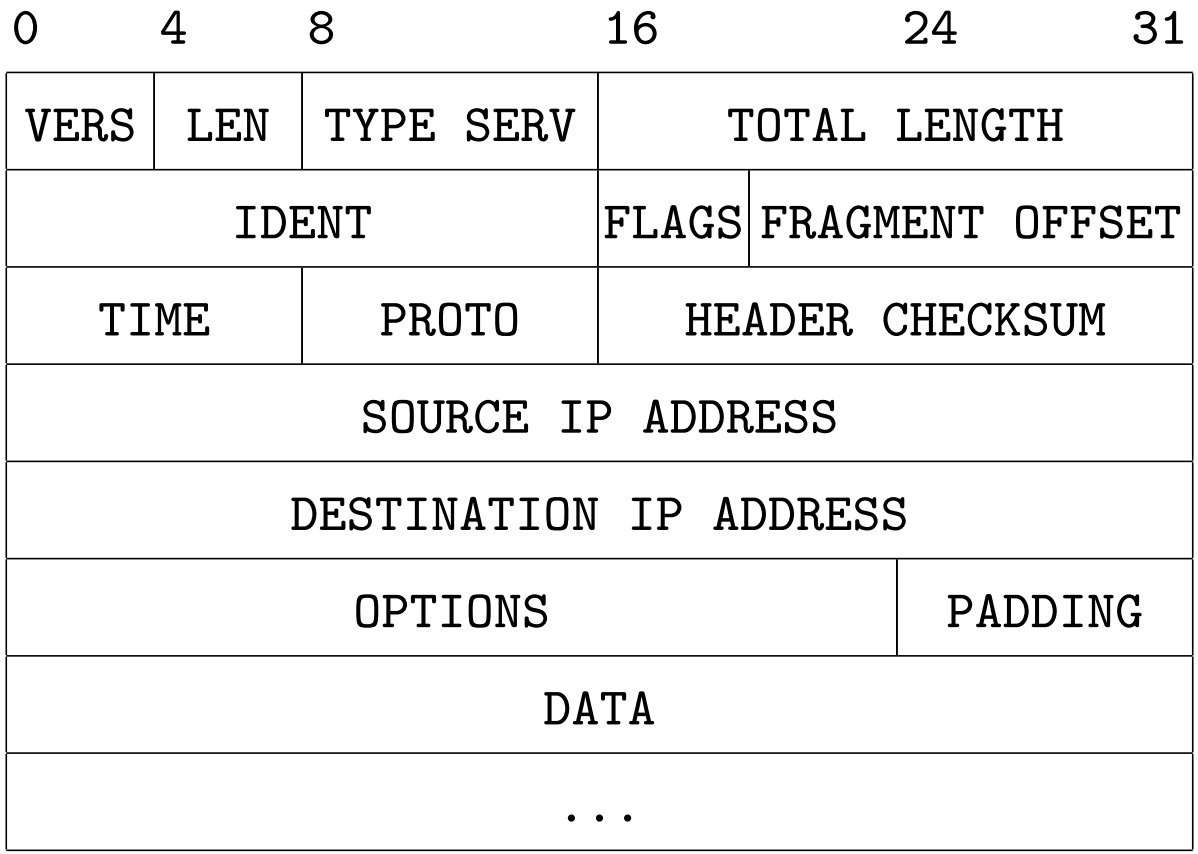

Encapsulation

- IP datagram travels in physical network frame
- Complete datagram is treated as data by the hardware
- TCP/IP defines standards for encapsulation on most network hardware technologies

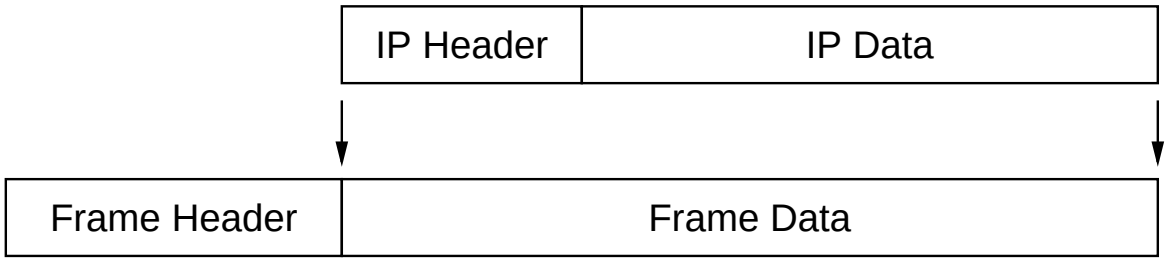

- Complete IP datagram is treated as data in physical network frame
- Encapsulation occurs in network interface software as the last step before the datagram is transmitted

Datagram Encapsulated In An Ethernet Frame

```
02 07 01 00 27 ba 08 00 2b 0d 44 a7 08 00 45 00
00 54 82 68 00 00 ff 01 35 21 84 eb 01 01 84 eb
01 02 08 00 73 0b d4 6d 00 00 04 3b 8c 28 28 20
0d 00 08 09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15
16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25
26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35
36 37
```

- IP header follows 14 octet Ethernet frame header and contains 20 octets
- IP source: `132.235.1.1 (84eb0101)`
- IP destination: `132.235.1.2 (84eb0102)`
- IP type: `01` (ICMP)

Network MTU

- Each network hardware technology imposes a fixed limit on the maximum size of a packet
- Size limit called *Maximum Transmission Unit* (MTU)
- Encapsulated datagram must be less than network MTU
- Possible solutions
  - Force datagram to be less than smallest possible MTU
    - Inefficient
    - Difficult to know minimum MTU
      - (See RFC 1191 – Path MTU Discovery)
  - Choose initial datagram size that seems appropriate and handle problems later
  - (IP uses the latter)

Datagram Fragmentation

- Needed when datagram larger than network MTU over which it must travel
- Performed by gateways
- Divides datagram into several, smaller datagrams called fragments
- Each fragment routed as independent datagram
- Final destination reassembles fragments
  - What's the other alternative?
  - Which way is "right"?

Datagram Fragmentation Details

- Each fragment is a datagram
- Gateway replicates initial datagram header for all fragments
- Offset field in header gives offset in original datagram for data in this fragment
- Two necessary bits in fragment flags field
  - "More Fragments" bit
    - Can infer "this is a fragment"
    - Can infer "last fragment"
  - Additional bit set in header to indicate "don't fragment"

Example Of Fragmentation

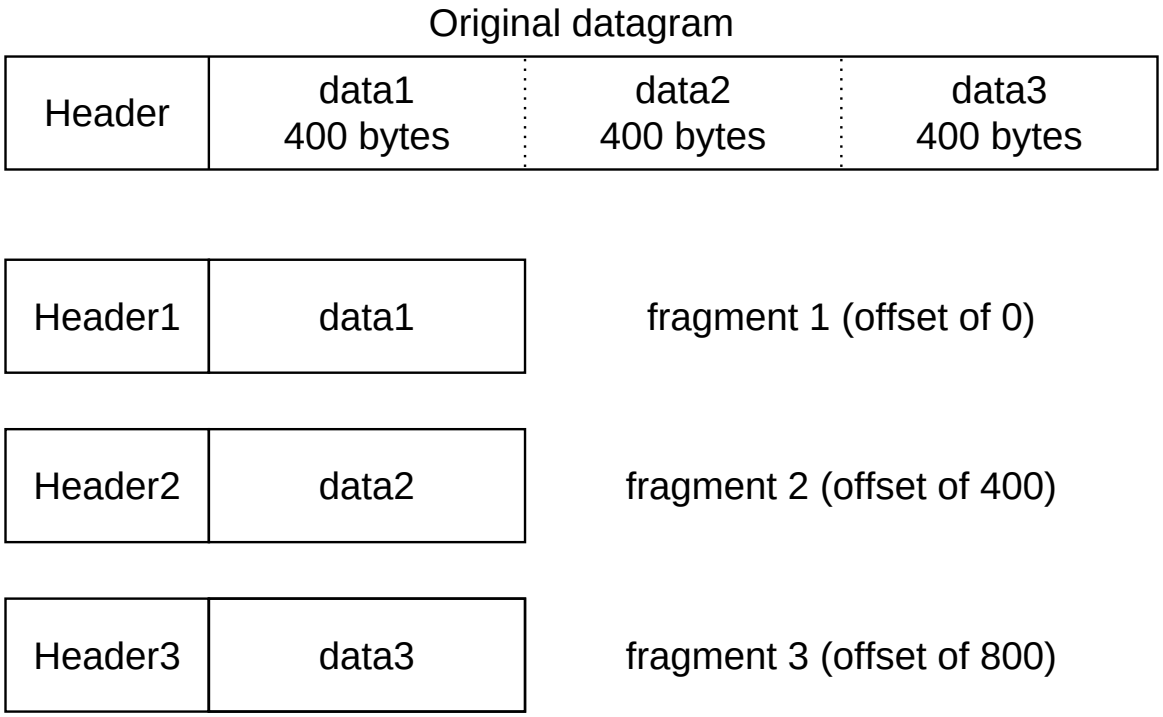

- Offset specifies where data belongs in original datagram
- Offset actually stored as multiples of 8 octets
- More fragments bit turned OFF in header of fragment #3

IP Options

- Option field is broken into 3 sub-fields

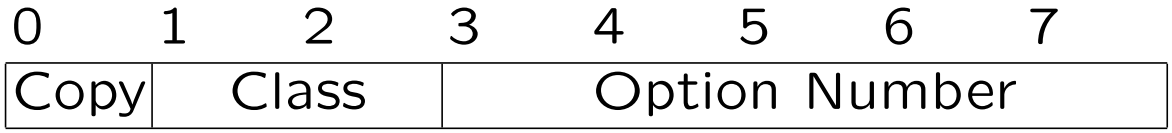

- Copy
  - 0 means copy only into first fragment
  - 1 means copy into all fragments
- Class
  - 0 – Datagram or network control
  - 2 – Debugging and measurement
  - 1,3 – Reserved

IP Options (continued)

- A few of the defined IP options:

Option Class	Option Number	Length	Description
0	0	-	End of option list. Used if options do not end at end of header
0	1	-	No operation (used for padding)
0	3	var	Loose source routing. Used to route a datagram along a specified path
0	7	var	Record route. Used to trace a datagram's route
0	9	var	Strict Source Routing. Used to route a datagram along a specified path
2	4	var	Internet timestamp. Record timestamps along route

Record Route Option

- Original sender enters "record route" option and leaves enough space for the expected number of gateways
- Each gateway along a datagram's path adds its IP address into the header

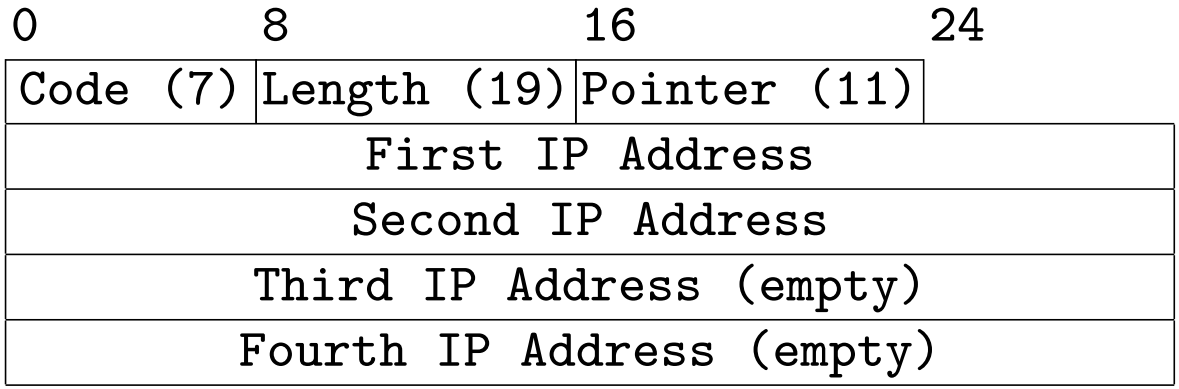

- Length - total length of the option, including entries not filled (plus first 3 octets)
- Pointer - offset within the option of the next available slot

Source Route Options

- IP provides two forms of restricted routing, *strict* and *loose* source routing
  - With *strict* routing, each gateway must be directly connected to the next gateway on the list

0	8	16	24
Code (9)	Length (19)	Pointer (11)	
First Hop IP Address			
Second Hop IP Address			
Third Hop IP Address			
Fourth Hop IP Address			

- Length - total length of the option (plus first 3 octets)
- Pointer - offset within the option of the next address to use