

Ethernet frame structure and Ethernet's MAC (Medium Access Control) addresses.

Each Ethernet card has a unique 48-bit address that is assigned by the manufacturer. This is called the Hardware Address or the MAC address of the card. The hardware address is a unique property of the interface card, and it is usually expressed as colon-separated bytes of hexadecimal numbers, for example: 08:00:20:2b:0c:b0. The first 24 bits of this address identifies the manufacturer of the card, and the next 24 bits is unique to that card. The 24 bits that uniquely identify the manufacturer is assigned by IEEE. You can find a list at <http://standards.ieee.org/regauth/oui/oui.txt>

Usually, on a linux machine, `/sbin/ifconfig` command gives you the MAC addresses of Ethernet cards (among other things) that are active on that host. The same command can also be used to activate and de-activate Ethernet interfaces. Here's a sample output from `/sbin/ifconfig` :

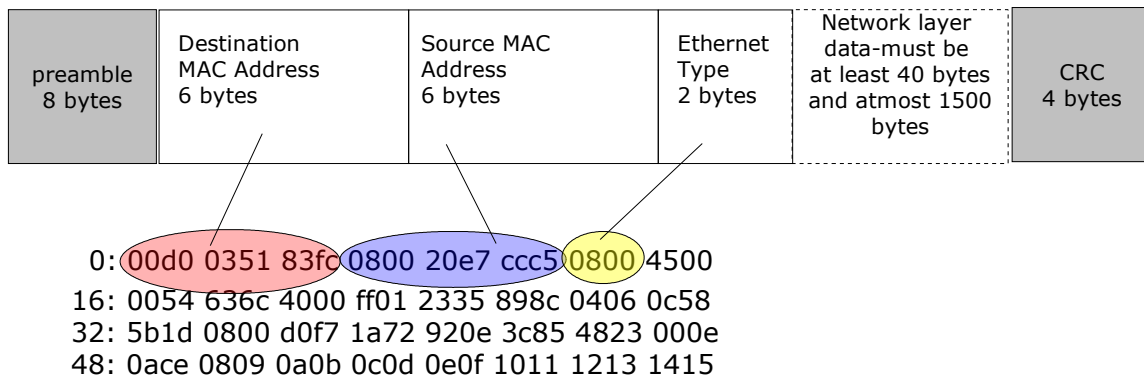
```
eth0  Link encap:Ethernet HWaddr 00:B0:D0:24:F2:1A
      inet addr:137.140.8.105 Bcast:137.140.8.255 Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:79360791 errors:0 dropped:0 overruns:65 frame:0
      TX packets:4256767 errors:0 dropped:0 overruns:0 carrier:0
      collisions:3566 txqueuelen:100
      RX bytes:2938524875 (2802.3 Mb) TX bytes:554138186 (528.4 Mb)
      Interrupt:5 Base address:0xe880

ppp0  Link encap:Point-to-Point Protocol
      inet addr:137.140.18.237 P-t-P:137.140.18.50 Mask:255.255.255.255
      UP POINTOPOINT RUNNING NOARP MULTICAST MTU:1500 Metric:1
      RX packets:2051 errors:0 dropped:0 overruns:0 frame:0
      TX packets:2571 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:3
      RX bytes:1019869 (995.9 KiB) TX bytes:269064 (262.7 KiB)

lo    Link encap:Local Loopback
      inet addr:127.0.0.1 Mask:255.0.0.0
      UP LOOPBACK RUNNING MTU:16436 Metric:1
      RX packets:147182 errors:0 dropped:0 overruns:0 frame:0
      TX packets:147182 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:81981556 (78.1 Mb) TX bytes:81981556 (78.1 Mb)
```

card manufactured
by Intel Corp.

The Ethernet frame usually consists of a *preamble* of 8 bytes at the beginning of the frame, used to “wake up” receivers and synchronize their clocks (See section 5.5.1 in the text). The preamble is followed by a 14-byte Ethernet header, then data from network layer, and ends with a 4-byte *CRC* (*cyclic redundancy check*), which is used for error detection (see section 5.2.3 of the text). The following figure shows a *tcpdump* output on a SUN machine with the Ethernet header:



In this example, 00:d0:03:51:83:fc is the MAC address (Ethernet address) of the network interface card (NIC) on the destination host, and 08:00:20:e7:cc:c5 is the MAC address of the NIC of the source host. *Note that the “destination” here means the immediate next-hop host, not the final destination of the IP datagram.* The link layer is only responsible for delivering data to the next-hop that is directly connected to the host; further sending from there is a routing decision to be made by the network layer of the next-hop host. As an IP datagram goes from one router to another, the Ethernet header on it changes, but the destination IP and source IP in the IP header does not change.

The type filed in the Ethernet header identifies the type of data the frame is carrying – 0x0800 means it is carrying an IP datagram.

Note that the *payload*, the network layer data in the frame, must be between 40 and 1500 bytes in length. If there is less than 40 bytes of data (as happens in ARP request/response), pad bytes of zeros are added. If there is more than 1500 bytes of data, the network layer need to fragment the data before passing it on to the Ethernet.