So far …

- Ethernet (IEEE 802.3):
  - Good for routing within local area network (LAN).
  - **Difficult for truly global routing**, every switch everywhere would need to store all MAC addresses – (we really need some kind of address hierarchy).
  - **Unreliable**:
    - No automatic retransmission on error.
    - No acknowledgements – sender doesn't know if receiver got the data.

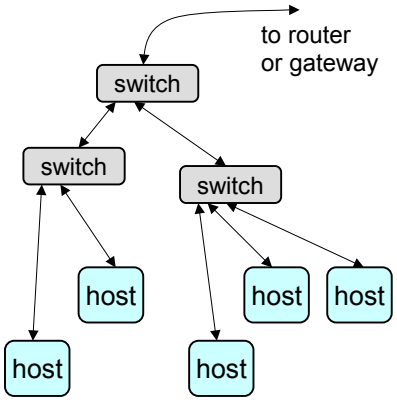

TCP/IP

A suite of protocols for global host addressing and reliable transmission on the internet.

- TCP/IP is an example of a layered protocol: each layer builds upon the layer below it, adding new functionality.
- The protocol stack is the collection of protocol that make up the suite:

protocol for transferring files / delivering mail	P2
protocol for routing and reliability	P1
protocol for sending and receiving data using specific hardware	P0

- Each protocol layer *encapsulates* the layer above it:

packet format:

P0 header	P1 header	P2 header	data

- Stacks are modular, so they can easily change when a new hardware model is adapted or needs of applications change. (Replace one module).

TCP/IP

- TCP/IP is a 4-layer protocol:

Application layer:	FTP, SMTP, HTTP
Transport layer:	TCP, UDP
Network layer:	IP
Link Layer:	IEEE 802.x, PPP, SLIP

- Link level examples:
  - IEEE 802.3 for Ethernet, 802.5 for token-ring, 802.11 for wireless,
  - Used with dial-up modems: Serial line IP (SLIP), Point-to-Point protocol (PPP).

IP (Internet Protocol)

Extends the idea of host address from MAC to a hierarchical "soft" address. All hosts take on an IP address.

- *The job of IP is to enable data to be transmitted between networks* (adds very little in the context of a LAN over what is possible with MAC addresses).
- Features of IP:
  - *Connectionless* – no concept of a job or session. Every packet treated individually.
  - *In-order delivery not ensured*.
  - *Unreliable* protocol.

The link layer (Ethernet) needs to know the unique address (MAC) of the specific place to next deliver the message. TCP/IP suite include ARP (address resolution protocol) to map from IP address to MAC address. Protocol works by broadcasting a request on the network – if a host sees its IP address, it replies with its MAC. If the IP is outside this subnet, then the router (connecting out) will reply).

IP Packets

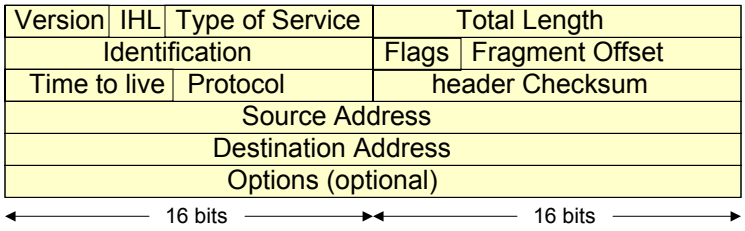

- *Protocol* field: says which high-level protocol sent the data – used by destination to pass packet to right protocol module.
- TTL (time to live): Initialized by the sender (usually 64) then decr. by 1 by every router the packet passes through. When reaches 0, the packet is discarded and the sender is notified with the Internet control message protocol (ICMP). This keeps packets from getting stuck in loops. (Also, used by traceroute).
- *Internet Addressing*: every host directly connected to the internet has a unique address (issued by InterNIC).
- Internet addresses are 32-bits long written as 4-Bytes separated by periods. Range:

1.0.0.1 to 223.255.255.255

Internet Addresses

- The 4-Byte address (x.x.x.x) is really in to parts: network ID and host ID. The dividing line between the two is not constant. Instead IP address are split into three classes:

Class A: 1st Byte in range 1-126, remaining 3-Bytes used for uniqe host address (126 class A networks each with up to 16M hosts).

Class B: 1st 2-Bytes in range 128.0 – 191.255, last 2-Bytes for host IDs. (16,000 nets with up to 16,000 hosts each).

Class C: 1st 3-Bytes in range 224.0.0 to 239.255.255. (2M nets with up to 254 hosts each).

- The IP address can be further divided to obtain a *subnet ID*. You decide how the subnet ID is arrived at by defining a 32-bit subnet mask. Its gets ANDed with the IP address to obtain subnet address. Example:
  subnet_mask = 255.255.255.0
  ID = 128.124.14.5,
  128.124 identifies the class B network,
  128.124.14 the subnet, and
  5 the host on the subnet.

Internet Addresses – special meanings

- Network ID of 0 in an address means "this network" – so for local communication only the host ID need be specified.
- Host ID of 0 means "this host".
- Network ID of 127 means loopback interface – also means "this host". Packets don't make it to the network. Allows different applications running on the same host to communicate.
- 224.x.x.x to 239.x.x.x are Class D addresses, used for multi-casting.
- 240.x.x.x to 247.x.x.x reserved for experimental purposes.
- Three sets of addresses are reserved for private address space – networks of computers that do not need to be addressed from the Internet.

  Class A: 10.x.x.x

  Class B: 172.16.x.x to 172.31.x.x

  Class C: 192.168.0.x to 192.168.255.x

  If you have equipment which uses IP addresses that have not been allocated by InterNIC then the addresses should be within these ranges and your router configured to not pass packets from/to these addresses.

We are running out of the address available in the 32-bit address space for IP address! IP version 6 (IPv6) has a new addressing system and is supposed to fix this.

IP Routing

- Local routing is done according to the specifics of the LANs own protocol.
- Routing to outside networks is done through *routers (these are either hosts with multiple NICs and special routing software, or special router hardware*.)
  - Each host on the LAN is assigned a default router, used to connect it to outside.

- A router examines every packet and compares the destination address with a table of address.
  1. If it finds an exact match, it forwards the packet to the address associated with that entry in the table.
  2. If the router doesn't find a match, it runs through to the table looking for a match just on the network ID. If a match is found, the packet is sent on to the address associated with that entry.
  3. If no match, the router sends it to the default, next-hop router, if present.
  4. If no default router present, the router sends an ICMP "host unreachable" messsage back to the sender.

- Routers build up their tables in multiple ways:
  - Static – read from a file on startup.
  - Dynamically, by broadcasting ICMP router solicitation messages to which other routers respond.
  - Other protocols are used to discover the shortest path to a location.
  - Routers are updated periodically in response to traffic conditions and availability of a route.

Transport Layer

Two most popular transport protocols are TCP and UDP.

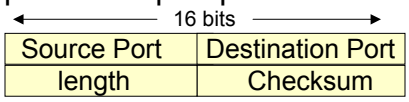

UDP Header

- UDP – User Datagram Protocol
  - Port numbers represent a software port.
  - They identify which protocol module sent (or is to receive) the data.
  - Standard port numbers exist:

    Telnet: port 23, Simple Mail Transfer Protocol: port 25
  - UDP and TCP use the port numbers to determine which application layer protocol should receive the data.
  - UDP isn't reliable, but appropriate for many application s like real-time audio and video (where if data is lost it is better to do without it than to send it again.) Also, gets used for online games.

TCP – Transmission Control Protocol

- Transport layer protocol used by most internet applications: FTP, HTTP, Telnet, …

- Connection-oriented: 2 hosts, one a client, and the other a server must establish a connection before any data can be transferred between them (SYN/ACK handshake). Once done the connection must be closed (FIN flag).
- TCP sends data using IP in blocks called segments.
- TCP includes mechanisms for ensuring data which arrives out of sequence is put back into the order it was sent.
- TCP implements flow-control, so a sender app. Cannot overwhelm a receiver app with data.
- TCP provides reliability: When data is received correctly, TCP sends an acknowledgement back to the sender. If the sender doesn't receive an ack within a certain period, the data is recent. For efficiency, the sender will usually send multiple segments without waiting for acks. It keeps track of what segments have or have not been acked – keeping a copy of those that have not, in case they need to be resent.
- ACKs are piggy-backed on data segments for efficiency.