Datagrams & Frames

- The internet layer (IP)...
  - constructs a datagram;
  - determines the next hop;
  - hands the datagram to the network interface layer.
- The network interface layer...
  - binds the next-hop address to a hardware address;
  - prepares the datagram for transmission.

Datagrams & Frames

- BUT…

  …network hardware does not understand a datagram format or internet addressing (IP).
- So how is a datagram transmitted across a physical network which doesn't understand what a datagram is?
- Encapsulate!

Encapsulation

- The network interface layer <u>encapsulates</u> an entire datagram in the data area of a hardware frame.
  - The network hardware ignores the datagram format.
  - The hardware treats a frame containing a datagram like any other frame.
- The sender places a frame type in the frame header indicating that the frame contains a datagram.
- The receiver reads the frame type and knows the frame contains a datagram.

Encapsulation & Multiple Hops

- What happens when a datagram's journey has several hops through several physical networks?
- Each router in the path from the source to the destination does the following:
  - It unencapsulates the datagram from the incoming frame on one network.
  - It processes the datagram to determine the next hop.
  - It encapsulates the datagram in an outgoing frame on the next network.

Encapsulation & Multiple Hops

- Since different physical networks may use different technologies, the datagram may be encapsulated in frames with different formats.
- The datagram itself is (almost!) unchanged.
- Note: Hardware frame headers do NOT accumulate during a trip through an internet.

MTU

- Maximum Transmission Unit (MTU) - the largest amount of data that can be sent across a given network in a single packet.
- Each hardware technology specifies an MTU.
- Any datagram encapsulated in a hardware frame must be smaller than the MTU for that hardware technology.

MTU & Datagram Transmission

- Datagrams can be larger than most hardware MTUs.

  - IP datagram: 1 to 65,535 octets.

  - Ethernet MTU: 1500 octets.

  - Token ring MTU: 2048 or 4096 octets.

- Also, an internet may have heterogeneous networks with different MTUs.
- So what happens when a datagram arrives at a router and it's too large to be sent over the next network?
- Remark: One way to deal with this problem is to limit a datagram's size. Not feasible.

Fragmentation

- fragmentation - a technique used by IP to split a large datagram into smaller pieces called fragments.
- An IP router detects a datagram larger than a network's MTU and splits the datagram into fragments smaller than the outbound network's MTU.
- Each fragment becomes an independent datagram, and each fragment datagram includes all the IP datagram header fields.

Datagram Reassembly

- Once a datagram is fragmented, it stays fragmented.
- The final destination performs the reassembly. Reassembly is not done along the way.
- Two advantages to having the final destination do the reassembly:
  - It reduces the amount of information needed by routers.
  - It allows routers to change dynamically.

Fragment Loss

- IP does not guarantee datagram delivery. So fragments may be lost, too.
- If a fragment is lost, what happens to the original datagram?
  - The destination drops the entire original datagram.
- How does the destination identify lost fragments?
  - When the first fragment is received, the receiver starts a timer.
  - If all fragments arrive before the timer expires, the receiver cancels the timer and reassembles the datagram.

Fragment Loss

- If the timer expires before all fragments arrive, fragments are assumed lost. The entire datagram is discarded.

- Reasons for the all-or-nothing behavior:
  - There's no mechanism for a receiver to tell a sender which fragments arrived.
  - The sender does not know about any fragmentation which took place.
  - Even if the sender did retransmit a datagram, the datagram might take a different route and be fragmented differently.

Fragmenting a Fragment

- A fragment can encounter a subsequent network with an even smaller MTU.
- Thus, a subsequent router might further fragment a fragment.
- IP does not distinguish between fragments and subfragments. Subfragments look just like regular fragments.
- A subfragment's header includes its position in the <u>original</u> datagram. Consequently, there's no need to reassemble subfragments first. All fragments can be reassembled together.

Background to IPv6

- An effort to develop a successor to IPv4 was begun in the early 1990s by the IETF.
- The prime motivation was that the 32-bit address space was being used up.
- The two leaders of the IETF's Address Lifetime Expectations working group estimated that addresses would be exhausted in 2008 and 2018.
- In 1996, ARIN reported that:

  - all IPv4 Class A addresses were assigned;

  - 62% of Class B addresses were assigned;

  - 37% of Class C addresses were assigned.

Background to IPv6

- For details, see R. Hinden, “IP Next Generation (IP ng),” at **http://playground.sun.com/pub/ipng/html/ipng-main.html**

Main Changes in IPv6

- Expanded addressing capabilities.
  - The IP address is increased from 32 bits to 128 bits.
- A streamlined 40-byte header.
  - The base header is fixed at 40 bytes.
  - Optional information can be included in multiple headers.
- Flow labeling and priority.

Making the Move?

- Figuring out how to make the transition from IPv4 to IPv6 could be interesting.
- There's doubt about whether IPv6 will make significant inroads into the Internet in the near future or even at all.
  - Many North American ISPs said they don't plan to buy IPv6-enabled equipment.
  - They claim there's little customer demand for IPv6's capabilities when IPv4, with a few patches, works well enough.
  - More interest in IPv6 in Europe and Asia.

Making the Move?

- It's enormously difficult to change network-layer protocols.
  - Since the early 1990s, numerous network-layer protocols have been trumpeted as the next major revolution for the Internet, but most have had limited effect.
  - Introducing new protocols in the network layer is like replacing the foundation of a house. It is difficult to do without tearing down the whole house or at least temporarily relocating the house's residents.

Making the Move?

- On the other hand, the Internet has witnessed rapid development of new protocols in the application layer.
  - Examples include HTTP and the Web, audio and video streaming, chat.
  - Introducing new protocols in the application layer is like adding a new layer of paint to a house. It is easy to do. Also, if you choose an attractive color, others in the neighborhood will copy you.
- So we can expect to see changes in the Internet’s network layer, but much more slowly that the changes that will occur in the application layer.