5.4 **Switched Local Area Networks**

In this section we give attention to switched local networks. Figure 5.15 **shows** a switched local network connecting three departments, two servers and a router with four switches.

Because these switches operate at the link layer, they switch link-layer frames (rather than network-layer datagrams), and therefore they don't recognize network-layer addresses, and don't use routing algorithms like RIP or OSPF to determine paths through the network of layer-2 switches. They use link-layer addresses to forward link layer frames through the network of switches.
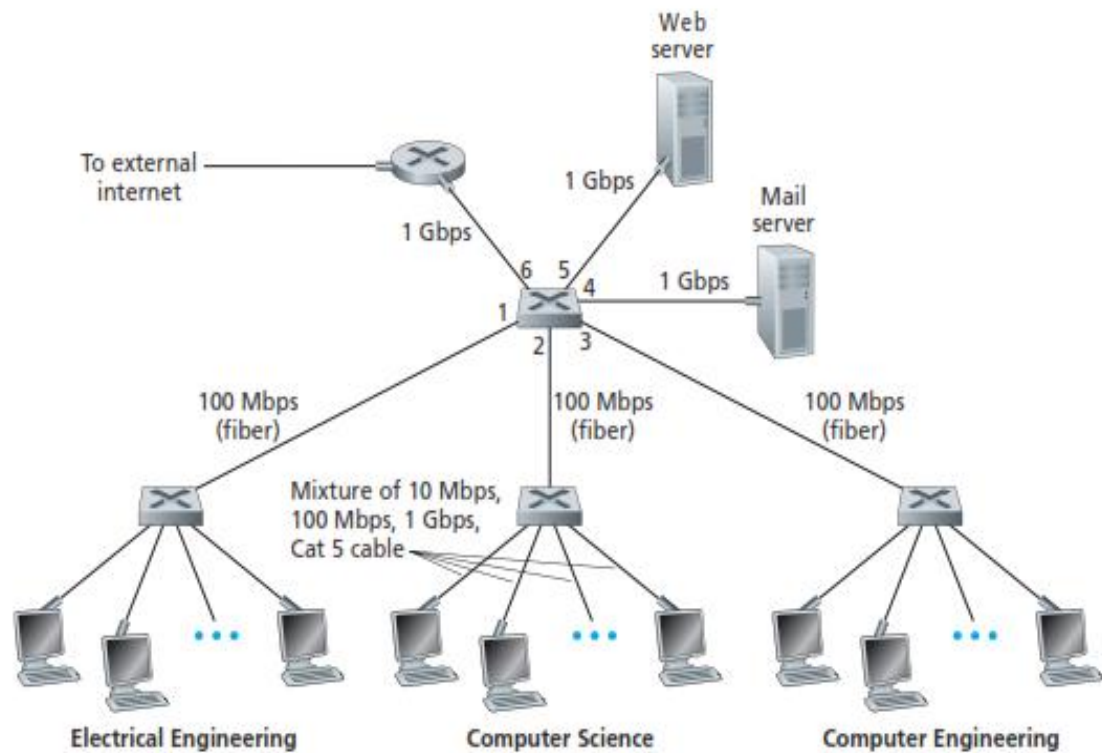
Figure 5.15 ♦ An institutional network connected together by four switches

How link-layer Switches do their function?

To answer this question we need to know:

MAC Addresses:

Each node has link-layer address, that is, each network interface has a link-layer address. A host or router with multiple network interfaces will thus have multiple link-layer addresses associated with it.

This is illustrated in Figure 5.16. A link-layer address is called a **LAN address**, a **Physical address,** or a **MAC address**.
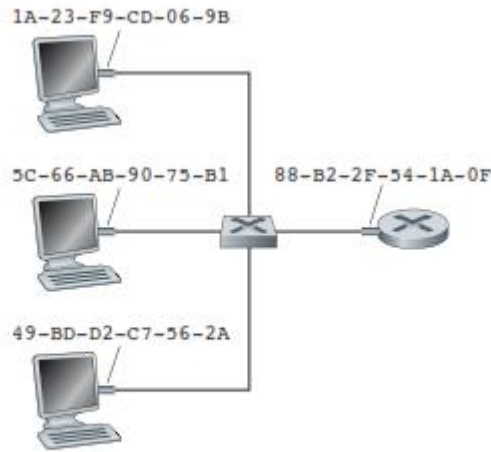
Figure 5.16 ♦ Each interface connected to a LAN has a unique MAC address

Because MAC address seems to be the most popular term, we'll henceforth refer to link-layer addresses as MAC addresses.

For most LANs (including Ethernet and 802.11 wireless LANs), the MAC address is **6 bytes long, giving 2^{48} possible MAC addresses**. As shown in Figure 5.16, these 6-byte addresses are **typically expressed in hexadecimal notation**, with each byte of the address expressed as a pair of hexadecimal numbers.

Flat and Hierarchical Structure

An adapter's **MAC address has a flat structure** (as opposed to a *hierarchical structure*) and doesn't change no matter where the adapter

goes. A laptop with an Ethernet interface always has the same MAC address, no matter where the computer goes. A smart phone with an 802.11 interface always has the same MAC address, no matter where the smart phone goes.

In contrast, **IP addresses have a hierarchical structure** (that is, a network part and a host part), and a host's IP addresses needs to be changed when the host moves, i.e, changes the network to which it is attached.

How an Adapter Send a Frame?

1- When an adapter wants to send a frame to some destination adapter, **the sending adapter inserts the destination adapter's MAC address into the frame** and then sends the frame into the LAN.

2- The **Switch broadcast the frame onto all of its interfaces**.

3- When **an adapter receives a frame, it will check to see whether the destination MAC address in the frame matches its own MAC address**.

4- If there is a match, the adapter extracts the enclosed datagram and passes the datagram up the protocol stack. If there isn't a match, the adapter discards the frame, without passing the network-layer datagram up.

How an Adapter can send a Broadcast Frame?

In this case, the sending adapter inserts a special MAC **broadcast address** into the destination address field of the frame. For LANs that use 6-byte addresses (such as Ethernet and 802.11), the broadcast address is a string of 48 consecutive 1s (that is,

FF-FF-FF-FF-FF-FF in hexadecimal notation).

In the above answer the sending adapter inserts the destination adapter's MAC address.

The question is:

How the sending adapter obtain the MAC of the destination adapter?

The answer is: by the aid of an address resolution protocol.

5.4.2 Address Resolution Protocol (ARP)

Because there are both network-layer addresses (for example, Internet IP addresses) and link-layer addresses (that is, MAC addresses), there is a need to translate between them.

For the Internet, this is the job of the **Address Resolution Protocol (ARP)** [RFC 826]

Why we need such a protocol?

For a source node to send a datagram it must give its adapter not only the IP datagram but also the MAC address for destination.

The sending adapter will then construct a link layer frame containing the destination's MAC address and send the frame into the LAN.

Example: (Source and destination in the same subnet)

In Figure 5.17, each host and router has a single IP address and single MAC address.

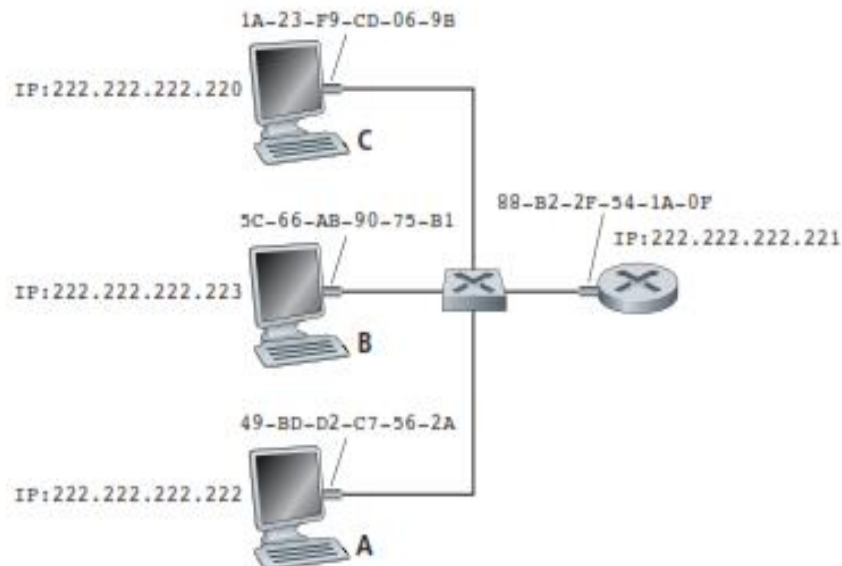As   usual,   IP   addresses   are   shown   in

Figure 5.17 ◆ Each interface on a LAN has an IP address and a MAC address

dotted-decimal notation and MAC addresses are shown in hexadecimal notation.

Now suppose that the host with IP address 222.222.222.220 wants to send an IP datagram to host 222.222.222.222.

To send a datagram, the source must give its adapter not only the IP datagram but also the MAC address for destination 222.222.222.222.

An ARP module in the sending host takes any IP address on the same LAN as input, and returns the corresponding MAC address.

In the example at hand, the sending host 222.222.222.220 provides its ARP module the IP address 222.222.222.222, and the ARP module returns the corresponding MAC address 49-BD-D2-C7-56-2A.

So we see that ARP resolves an IP address to a MAC address. In many ways it is analogous to DNS (studied in Section 2.5), which resolves host names to IP addresses.

However, one important difference between the two resolvers is **that DNS resolves host names for hosts anywhere in the Intern**et, **whereas ARP resolves IP addresses only for hosts and router interfaces on the same subnet.**

What is ARP Table?

1- Each host and router has an **ARP table** in its memory, which contains mappings of IP addresses to MAC addresses. Figure 5.18 shows what an ARP table in host 222.222.222.220 might look like.

2- The ARP table also contains a time-to-live (TTL) value, which **indicates when each mapping will be deleted from the table**. Note that a table does not necessarily contain an entry for every host and router on the subnet; some may have never been entered into the table, and others may have expired. (A typical expiration time for an entry is 20 minutes from when an entry is placed in an ARP table).

IP Address	MAC Address	TTL
222.222.222.221	88-B2-2F-54-1A-0F	13:45:00
222.222.222.223	5C-66-AB-90-75-B1	13:52:00

Figure 5.18 ♦ A possible ARP table in 222.222.222.220

How ARP works?

1- The sending host passes the destination IP address to its ARP.

2- The ARP looks the IP address in its Table.

3- If the sender's ARP table has an entry for the destination node, ARP returns the MAC address.

4- If the ARP table doesn't currently have an entry for the destination then ARP **protocol constructs a special packet called an ARP query packet.** (An ARP packet has several fields, including the sending

7

and receiving IP and MAC addresses. Both ARP query and response packets have the same format).

1- The purpose of the ARP query packet is to **query all the other hosts and routers on the subnet to determine the MAC address corresponding** to the IP address that is being resolved.

Ok, the question now what if the destination in another subnet?

Figure 5.19, shows a simple network consisting of two subnets interconnected by a router. In Figure 5.19., each host has exactly one IP address and one adapter. But, as discussed in Chapter 4, a router has an IP address for *each* of its interfaces. ==For each router interface there is also an ARP module (in the router) and an adapter.== Because the router in Figure 5.19 has two interfaces, it has two IP addresses, two ARP modules, and two adapters. Of course, each adapter in the network has its own MAC address.

Also note that Subnet 1 has the network address 111.111.111/24 and that Subnet 2 has the network address 222.222.222/24. Thus all of the interfaces connected to Subnet 1 have addresses of the form 111.111.111.xxx and all of the interfaces connected to Subnet 2 have addresses of the form 222.222.222.xxx.
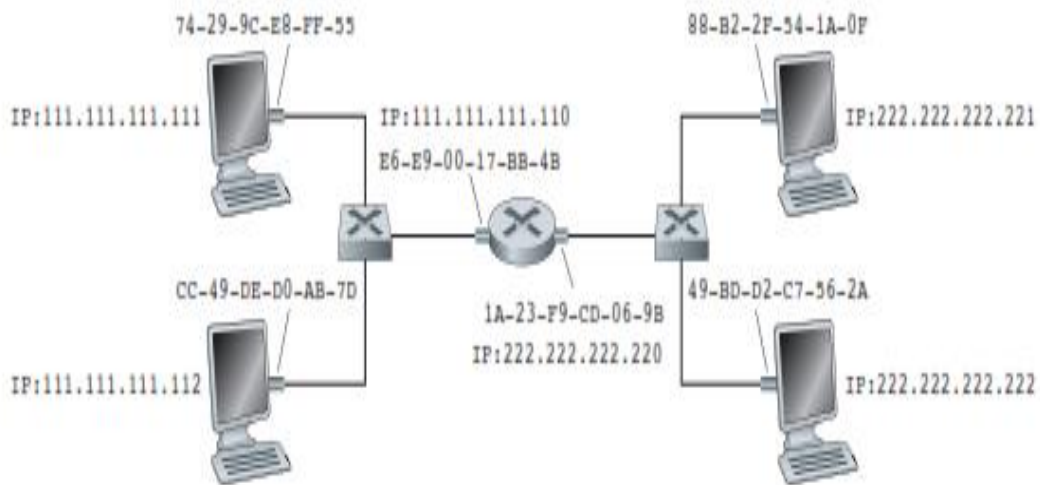
Figure 5.19 ◆ Two subnets interconnected by a router

Suppose a host on Subnet 1 would send a datagram to a host on Subnet 2.

Specifically, suppose that host 111.111.111.111 wants to send an IP datagram to a host 222.222.222.222. The sending host passes the datagram to its adapter, as usual. But the sending host must also indicate to its adapter an appropriate destination MAC address.

What MAC address should the adapter use?

If we look carefully at Figure 5.19, we see that in order for a datagram to go from 111.111.111.111 to a host on Subnet 2, the datagram must first be sent to the router interface 111.111.111.110, which is the IP address of the first-hop router on the path to the final destination. Thus, the appropriate MAC address for the frame is the address of the adapter for router interface 111.111.111.110, namely,

E6-E9-00-17-BB-4B.

The sending host knows the IP address of the interface to its subnet from the routing table.

How does the sending host acquire the MAC address for 111.111.111.110?

By using ARP, of course! Once the sending adapter has this MAC address, it creates a frame (containing the datagram addressed to 222.222.222.222) and sends the frame into Subnet 1.

The router adapter on Subnet 1 sees that the link-layer frame is addressed to it, and therefore passes the frame to the network layer of the router.

The router now has to determine the correct interface on which the datagram is to be forwarded. **As discussed in Chapter 4, this is done by consulting a forwarding table in the router**. The forwarding table tells the router that the datagram is to be forwarded via router interface 222.222.222.220.

This interface then passes the datagram to its adapter, which encapsulates the datagram in a new frame and sends the frame into Subnet 2.

This time, the destination MAC address of the frame is indeed the MAC address of the ultimate destination. And how does the router obtain this destination MAC address? From ARP, of course!