HUB, BRIDGE, SWITCH, ROUTER: DEFINING THE FUNCTIONS OF THESE DIFFERENT DEVICES

Doug Jones
YAS Broadband Ventures

Abstract

Hubs, Bridges, Switches and Routers all play roles in creating LANs and IP networks. This paper discusses the networking characteristics these different devices.

Thirty years ago it was antenna design. Twenty years ago it was amplifier design. Ten years ago it was fiber optics. Now, it's the Internet. The basic building blocks of intranets and the Internet include the hub, bridge, switch, and router. These devices all have specific functions and roles to play when piecing together an IP network. The intent of this paper is to introduce each piece of equipment, describe what it does, and how it does it.

The first three devices; hub, bridge, and switch, are very closely related. They all operate on Ethernet frames. The fourth device, a router, operates on IP packets. The difference between a frame and a packet is semantic, having to do with which layer in the OSI protocol model a particular protocol operates. Regardless, both a frame and a packet are standardized sets of bits used to represent data.

This paper will begin with a brief description of the basics of Ethernet and how it works. This introduction is necessary to draw the distinctions between a hub, bridge, and switch. Next, the paper will briefly describe the Internetworking Protocol (IP) and how a router works.

Ethernet Background

Ethernet was invented by Bob Metcalfe in 1972 at the Xerox PARC (Palo Alto Research Center). This first Ethernet operated at about 3 Megabits per second (Mbps). In 1980, the first 10 Mbps Ethernet standard was published by DEC, Intel, and Xerox and is known as DIX Ethernet. At this point, the IEEE 802 committee, responsible for LAN/MAN standards development, took up the original DIX specification and used it as the basis for an IEEE standard. This work was completed in 1985 by the IEEE 802.3 subcommittee and is titled the *IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications*. Note that the word Ethernet does not appear in this title, however, both the IEEE 802.3 protocol and DIX are simply referred to as Ethernet.

The IEEE 802.3 standard is the official Ethernet standard, and it differs slightly from the DIX Ethernet standard, though the two are backward compatible. An Ethernet frame is shown in Figure 1. The difference lies in the Type/Length field. In DIX Ethernet, this field is used to identify the type of protocol carried in the payload. In IEEE 802.3, this field indicates the length of the payload. Though this one field is used for two purposes, it is possible to determine how a particular frame is encoded. When originally developed, Xerox made sure that none of the protocol identifiers had a value of less than 1500; hence, if the value is greater than 1500 it is a DIX frame. If the value is 1500 or less, then it is an IEEE 802.3 frame. In the case of an IEEE 802.3 frame, there are additional fields (LLC fields) in the payload that identify the type of protocol being carried in the payload.

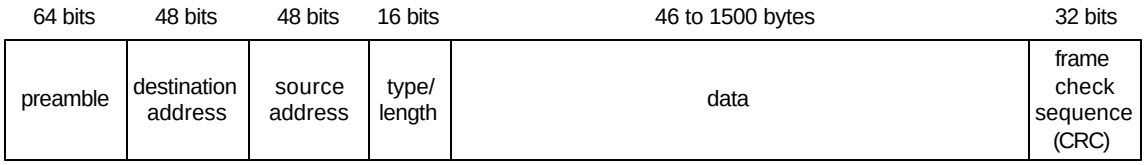

Figure 1

The source and destination addresses are both 48 bits in length, and are commonly referred to as "MAC addresses." The term MAC stands for Media Access Control, which is a generic term for a data link layer protocol, which Ethernet is. Note the payload is up to 1500 bytes. While an Ethernet frame can carry many types of payload, later on this paper will address the case when that payload is an IP packet.

The Ethernet Collision Domain

Ethernet networks originally operated over coaxial cable and multiple devices were connected to the same cable segment. Because multiple devices connected to this cable, there was the possibility of two devices transmitting at the same time. When two (or more) devices transmitted at the same time, those transmissions would collide and the data would be lost.

The IEEE 802.3 name for Ethernet, CSMA/CD, gives a hint at how Ethernet attempts to minimize collisions. Before transmitting on the cable, a device first "listens" on that cable (carrier sense) for existing transmissions. If the device senses the cable is busy, it does not transmit. If the device senses the cable is idle, it will transmit. The worse case scenario is when devices on opposite ends of a cable segment sense the cable is idle and begin transmitting at the same time. Due to propogation delay, this collision takes the longest time to detect.

Ethernet was designed to guarantee that a transmitting station could tell whether its transmission had failed due to a collision. In order to meet this guarantee, the length of cable had to be limited and the minimum frame size had to be specified. Given that the maximum cable length was 2.5 kilometers, given the speed of electricity on the cable, and given a transmission speed of 10 Mbps, it might take as long as 512 bit times to detect a collision, hence the minimum frame size of 64 bytes (64 bytes x 8 bits/byte = 512 bits).

As the number of devices on a cable segment increases, so does the traffic and the probability of a collision occurring. Collisions mean that data needs to be retransmitted and hence is an inefficient usage of the available bandwidth. One way to keep the number of collisions low was to keep the length of cable short.

What's an Ethernet Hub

An Ethernet hub is best described as an Ethernet repeater. That is, when a bit is received on a hub port, that bit is repeated onto every other port on the hub. On the one hand, this is a dumb, low cost device. On the other hand, a hub increases the size of the collision domain.

Each cable plugging into a hub is considered a "LAN segment." Each LAN segment is a collision domain. A hub, since it just repeats bits onto every segment, increases the size of the collision domain seen by all packets.

Consider the case of a 4-port hub. (There is nothing magical about 4 ports, this is just a standard configuration for commercially available hubs.) There are LAN segments connected to three of the 4 ports, call them segments 1, 2, and 3. A device connected to LAN segment 1 has data to transmit. This

device follows the CSMA/CD MAC protocol by "listening" on its cable segment and, determining it is idle, begins transmitting. The hub receives these bits and relays them onto LAN segments 2 and 3 without bothering to determining if those segments are idle. Hence, if other devices are transmitting on those segments, there are collisions.

Hubs are low cost devices intended for the quick interconnection of low traffic LAN segments. When used to connect high traffic LAN segments, hubs increase the probability of collisions, lost data, and retransmission on all of those segments. Cascading hubs should be discouraged as this just compounds the problem.

What's an Ethernet Bridge

A bridge is a lot like a hub, except this device respects the collision domain of the individual LAN segments connected to it.

Consider again the 4 port device, but this time a bridge. When a bridge receives a transmission, it buffers (stores in memory) the entire frame. Here is a first key differentiation from a hub. A bridge can determine what bits make up a frame, as opposed to simply repeating bits. Before the fame is sent onto other segments, the bridge "listens" on each of those segments for idle. A second key differentiator is that a bridge implements the CSMA/CD protocol on each of its ports, while a hub does not.

Therefore a bridge is a higher cost device. It has to contain both extra memory and the logic to perform the MAC protocol. However, a bridge does not increase the size of a collision domain like a hub does.

What's an Ethernet Switch

A switch is a lot like a bridge, in fact, another name for a switch is a "learning" bridge. As frames pass through the various ports of a switch, the switch learns which devices are connected to the LAN segments and stores their Ethernet MAC addresses in memory.

To learn how the switch operates, continue with the above example but place two devices on each LAN segment. As Figure 2 shows, segment 1 has devices A and B, segment 2 has devices C and D, and segment 3 has devices E and F.

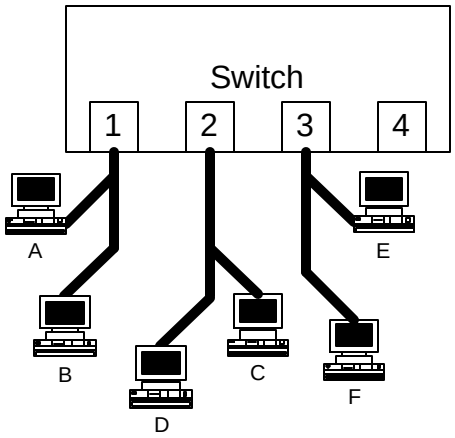

Figure 2

In this case, if device A sends a frame destined for device B, even though that frame is received at the switch, the switch will not forward the frame onto the other LAN segments because the switch knows that A and B are on the same LAN segment.

Likewise, if E sends a frame to B, the switch will only transmit that frame on segment 1 and will not transmit that frame on segment 2. Like a bridge, a switch also implements the CSMA/CD protocol.

If the switch receives a frame and does not recognize the destination MAC address (i.e., broadcast frames, the first time a device transmits, frames destined to other places, etc.), that frame will be forwarded onto all other LAN segments.

More sophisticated switches allow the creation of Virtual LANs (VLANs) where a subset of ports can be grouped together into a logical "switch. For instance, a common commercial-grade switch configuration is to have 96 ports. A subset of these ports can be logically grouped together to create smaller switch domains.

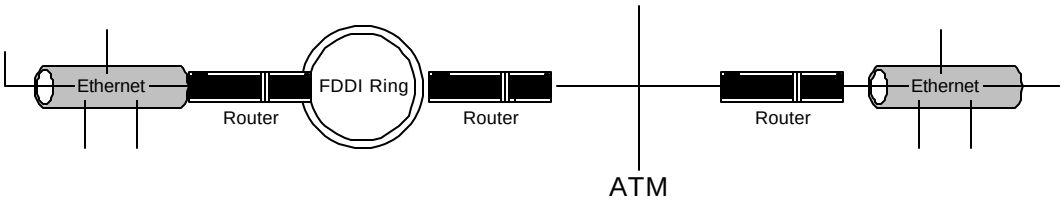

Figure 3

IP Background

Up to now, this paper has discussed Ethernet frames and the role of hubs, bridges, and switches in maintaining the most efficient collision domain on LAN segments. IP routers are an entirely different type of device. True, routers have Ethernet interfaces (10Base-T, 100Base-T, etc.), but IP has a different function than Ethernet.

Ethernet is a point-to-point protocol, and IP is an end-to-end protocol. While both use addresses, an Ethernet address only has meaning on a LAN segment whereas an IP address has global meaning.

As shown in Figure 3, routers can be connected with many different types of LANs. Ethernet, FDDI, and ATM are shown but there are many others. On each of these LANs, specific frame formats are used for carrying data. Just as the Ethernet frame in Figure 1 has a specific format, FDDI and ATM frames have their own formats and rules. All of these frames can carry IP packets. While on each LAN segment the frame format will differ, the IP packet carried in those frames will remain the same.

On each LAN segment, the router port will have a MAC address that is specific for the individual LAN type. As the IP packet goes from router to router (and traverses the LANs), the frames carrying that packet will change depending on the technology of the underlying LAN; however, the IP packet carried in the frame will not be changed.

What's an IP Router

A router has two main functions, IP routing and packet forwarding.

IP routing is the process whereby routers exchange route information so they know where to forward IP packets. Route information is exchanged by using a variety of protocols, including:

- RIP (Routing Information Protocol)
- OSPF (Open Shortest Path First)
- BGP (Border Gateway Protocol)
- Many others

Routing protocols are a study in and of themselves. Suffice it to say routers are continually exchanging route information because routes across the Internet are continually changing for a number of factors. For example, different paths through the Internet can be congested or out-of-service at any given time. Routing protocols identify these bottlenecks and provide routers with "up to the second" information to keep the traffic moving.

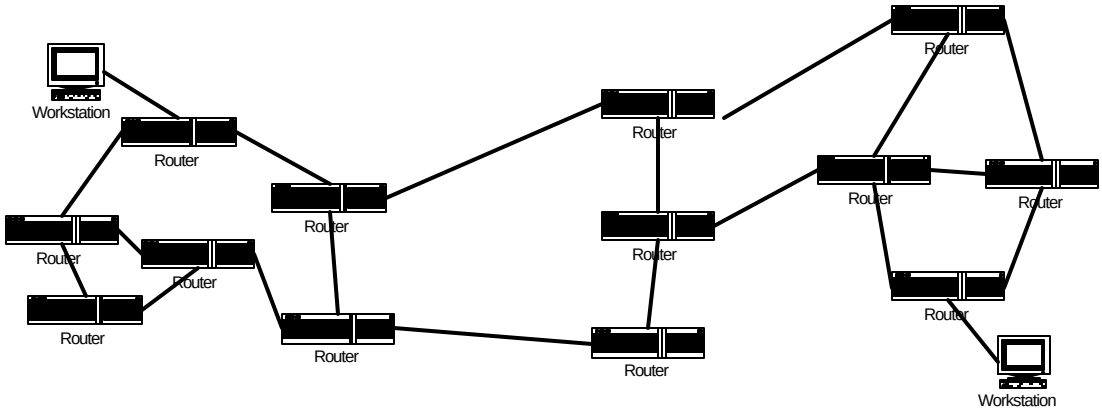

Figure 4

Consider the case of a single large email message traversing the Internet. This message is too big to fit into one IP packet and therefore must be fragmented into many IP packets. In going from one email account to another, the individual packets of this email message may take different paths across the Internet. As these packets are received at the other end, even out of order, higher layer protocols are responsible for reassembling them into the original email (in this specific case, that protocol is TCP, or the Transmission Control Protocol).

These self-healing and redundant properties make the Internet very robust. The original Internet was created under the guidance of the U.S. Military in the 1960s and one of the design criteria was to remain operational even in the event of massive destruction of parts of the network (i.e., nuclear attack).

The routing protocols exchange route tables, essentially lists of information that identify paths through the Internet to particular destinations. When an IP packet arrives at a router, the IP address in that packet is used to index into the route table and the route information is used to decide how to forward that packet. While Figure 3 was an example of a simple connection between routers, a more realistic network design is given in Figure 4.

In Figure 4, there are over a dozen possible paths to get data between the two workstations. The routing protocol calculates the most efficient path between the two workstations and the routers in between forward packets along that path. These decisions are made on a hop-by-hop basis. Based on its route table, each router along the path makes an autonomous decision about the best path to forward the packet. In the case of one router (or LAN segment) going out of service, the routing protocol quickly (on the order of seconds) recalculates the new most efficient path for packets to take and propogates this information among the affected routers. While some packets may be dropped during the reconfiguration, the network is self-healing and communications continue.

The difference between a MAC address and an IP address

As written earlier, Ethernet is a point-to-point protocol and IP is an end-to-end protocol. This means that as a frame traverses LAN segments, its source and destination addresses will be changed, however, the IP addresses in the IP packet carried by that frame will remain the same.

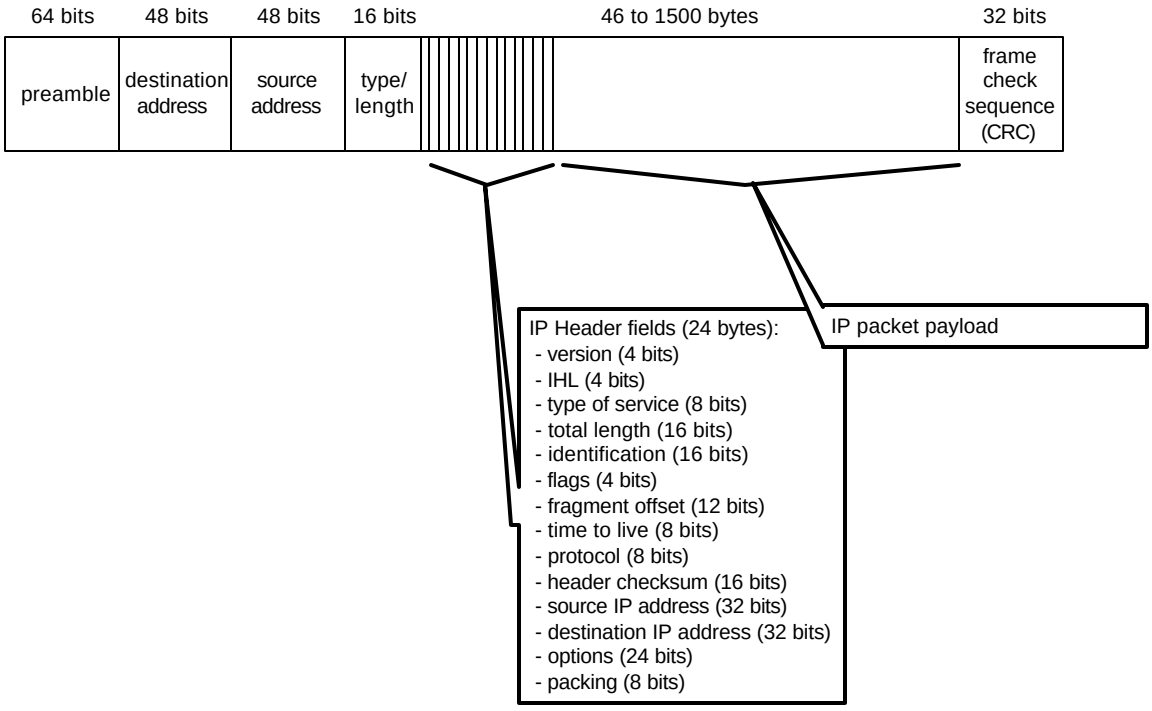

Figure 5

Figure 5 shows a redrawing of Figure 1, except the payload of the Ethernet frame contains an IP packet. Consider the case when this Ethernet frame (carrying the IP packet) is forwarded from one workstation to another (as diagramed in Figure 4).

The scenario is as follows. An application running on the workstation creates the IP packet and sends it to its Ethernet LAN interface for transmission. The IP packet has the source IP address set to this workstation and the destination IP address set to the other workstation. The Ethernet LAN interface on the workstation encapsulates the IP packet in an Ethernet frame, and sets the source MAC address equal to that of the workstation and the destination MAC address to that of the first router. The workstation then puts this frame on the Ethernet link according to the CSMA/CD protocol rules. The router "sees" the frame on the LAN segment with destination MAC address set to its own, hence this frame is destined for that router.

Although other devices on the LAN segment see the frame, they would discard this frame because the destination MAC address is for the router. (Note, while Ethernet MAC addresses only have significance on a particular Ethernet segment, each Ethernet MAC address is unique in that it should never be duplicated in another device.)

Now that the router has the Ethernet frame, that router strips off the Ethernet bits and uses just the IP packet (mainly the destination IP address) to make the next-hop packet forwarding decision. When that packet exits the router, the router re-

encapsulates the IP packet in a new Ethernet frame (or whatever protocol is appropriate for that LAN segment, as proposed in Figure 3) with a source MAC address of the router port and the destination MAC address of the next hop router. Note that the IP addresses in the IP packet have not been changed. In this manner, the IP packet makes its way, router-to-router, LAN segment-to-LAN segment, between the two workstations. When finally delivered to the destination workstation, the Ethernet frame will contain the MAC address of the last-hop router in the source address field and the MAC address of the workstation in the destination address field.

Summary

Ethernet and IP are separate protocols that do different things. However, IP packets are routinely encapsulated in Ethernet frames for transmission on LAN segments.

Hubs, bridges, and switches operate in the Ethernet domain and Routers operate in the IP domain. Hubs, bridges, and switches connect Ethernet segments. Hubs are low cost, but increase the collision domain a particular frame is exposed to. Switches have additional cost for memory and logic, but they create more efficient LAN segments.

Routers perform routing and forwarding. Routing involves exchanging route information with neighboring routers on the network. The route information is used to forward individual IP packets. Routing is very dynamic and is a main reason why the Internet is such a robust network.

References:

1. Spurgeon, Charles E.; Ethernet - The Definitive Guide, O'Reilly & Associates, 2000.
2. IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, IEEE 802 Committee, 1998.
3. Perlman, Radia; Interconnections - Bridges and Routers, Addison Wesley, 1992.
4. Huitema, Christian; Routing in the Internet, Prentice-Hall, Inc., 1995.

Author Contact:

Doug Jones, Chief Architect
YAS Broadband Ventures
voice: (303) 661-3823
doug@yas.com