WIRELESS ETHERNET (IEEE 803.11)

Wireless LANs form a very small percentage of LANs in operation today, but their use is growing rapidly. Wireless LANs transmit data through the air using radio or infrared transmission rather than through coaxial cable, twisted pair, or fiber optic cable. Until recently, there were few widely accepted standards for wireless LANs and as a result, equipment from different vendors could not be used in the same network. Over the past few years, however, several standards for wireless LANs have emerged, as have new terms: Wireless LAN (WLAN) and Local Area Wireless Network (LAW).

The IEEE 802.11 standard will likely become the dominant standard for WLANs. It is very similar to Ethernet, with a few differences. Most importantly, IEEE 802.11 systems are easily connected into Ethernet LANs and translate between IEEE 802.3 Ethernet and IEEE 802.11 wireless. For this reason, IEEE 802.11 is usually called wireless Ethernet, although its official name is Wireless LAN. IEEE 802.11 is rapidly evolving[1].

Topology

The logical and physical topologies of wireless Ethernet is the same as those of traditional Ethernet. It is a physical star, and a logical bus (see Figure 6-7). A central wireless Access Point (AP) is a radio transmitter that plays the same role as a hub in traditional Ethernet. All devices in the WLAN use the same radio frequencies, so the WLAN functions as a shared media LAN in the same manner as traditional Ethernet: computers must take turns using the one circuit. Because the system uses radio waves, the signal travels in all directions from the AP. The maximum range from the AP to the computers is determined by the amount of interference (e.g., concrete walls), but is typically 100-500 feet.

---- Figure 6-7 ---

The computers on the WLAN have a NIC inside the computer that is connected to an external transmitter that communicates with the AP (see Figure 6-8). The external transmitter transmits radio signals to a receiver that acts like a network hub and enables wireless computers to communicate with each other and with traditional wired networks.

---- Figure 6-8 updated figure 7-11 ---

[1] For more information, see the IEEE standards site at grouper.ieee.org/groups/802/11 and the Wireless Ethernet Compatibility Alliance at http://www.wirelessethernet.org.

Usually a set of APs are installed, so that there is complete wireless coverage in some area, enabling users to roam from AP to AP. When configured with a wireless network, a set of laptops or Palm-based devices becomes an effective way to provide a portable groupware configuration or to enable workers to walk through a facility and have constant network access at any point (e.g., warehouse, hospital, airport).

One potential problem is security. Because anyone within range of a WLAN can receive transmissions, eavesdropping a serious threat. IEEE 802.11 encrypts all transmissions using a 40-bit encryption scheme so that only those computers that have the key can decode and read the messages. However, as will be discussed in Chapter 12, a 40-bit key is not terribly good.

Media Access Control

Media access control in wireless Ethernet is Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), which is similar to the CSMA/CD used by traditional Ethernet. With CSMA/CA, computers listen before they transmit and if no one else is transmitting, they transmit. Detecting collisions is more difficult in radio transmission than in transmission over so wireless Ethernet attempts to avoid collisions to a greater extent than traditional Ethernet. CSMA/CA can use two approaches simultaneously to media access control.

Physical Carrier Sense Method

The first media access control method is the physical carrier sense method because it is based on the ability of computers to physically listen before they transmit. Each packet in CSMA/CA is sent using stop and wait ARQ (see Chapter 4). After the sender transmits one packet, it immediately stops and waits for an ACK from the receiver, before attempting to send another packet. When the receiver of a packet detects the end of the packet in a transmission, it waits a fraction of a second to make sure the sender has really stopped transmitting, and then immediately transmits an ACK (or a NAK). The original sender can then send another packet, stop and wait for an ACK, and so on.

While the sender and receiver are exchanging packets and ACKs, other computers may also want to transmit. So when the sender ends its transmission, why doesn't some other computer begin transmitting before the receiver can transmit an ACK? The answer is that the physical carrier sense method is designed so that the time the receiver waits after the transmission ends before sending an ACK is significantly less than the time a computer must listen to determine no one else is transmitting before initiating a new transmission. Thus the time interval between a transmission and the matching ACK is so short, that no other computer has the opportunity to begin transmitting.

Virtual Carrier Sense Method

The second media access control technique is called the virtual carrier sense method, because it does not rely on the physical media. The physical carrier sense method works well in traditional Ethernet, because every computer on the shared circuit receives every transmission on the shared circuit. However, in a wireless environment, this is not always true. A computer at the extreme edge of the range limit from the AP on one side may not receive transmissions from a computer on the extreme opposite edge of the AP's range limit. In Figure 6-7, both computers may be within range of the AP, but not be within range of each other. In this case, if one computer transmits, another computer on the opposite edge may not sense the other transmission, and transmit at the same time, causing a collision at the AP. This is called the hidden node problem, because the computer at the opposite edges of the WLAN are hidden from each other.

When the hidden node problem exists, the AP is the only device guaranteed to be able communicate with all computers on the WLAN. Therefore, the AP must manage the shared circuit using a controlled access technique, not the contention-based approach of traditional Ethernet (see Chapter 4). With this approach, any computer wishing to transmit first sends a request to transmit (RTS) to the AP which may or may not be heard by all computers. The RTS requests permission to transmit and to reserve the circuit for the sole use of the requesting computer for a specified time period. If no other computer is transmitting, the AP responds with a clear to transmit (CTS) specifying the amount of time for which the circuit is reserved for the requesting computer. All computers hear the CTS and remain silent for the specified time period.

The virtual carrier sense method is optional. It can be used always, never, or just for packets exceeding a certain size, as set by the LAN manager. From Chapter 4, you should remember that controlled access methods provide poorer performance in low traffic networks and better performance in high traffic networks, so unlike its traditional Ethernet cousin, a WLAN using this controlled access approach may provide a higher percentage of available capacity to the attached devices, perhaps as high as 90%.

Types of Wireless Ethernet

Two basic types of Wireless Ethernet have been defined: IEEE 803.11b, which is the most common today, and 802.11a which is the high speed alternative.

IEEE 803.11b

IEEE 803.11b in turns has two basic forms. Direct-sequence spread-spectrum (DSSS) systems transmit signals through a wide spectrum of radio frequencies simultaneously (in the 2.4 GHz band). The signal is divided into many different parts and sent on different frequencies simultaneously. Because several radio devices could be operating in these same frequency bands (not just wireless LANs, but

also cordless phones), devices add a special code to each bit transmitted that uniquely identifies the signal and enables the intended receiver to identify it.

Frequency-hopping spread-spectrum (FHSS) systems transmit signals through the same wide spectrum of radio frequencies, but use each frequency in turn. A short burst of data is sent on one frequency (usually less than half a second) and then the sender changes to another pseudorandom frequency and broadcasts another burst of data before changing to another frequency, and so on. The transmitter and receiver are synchronized so that they both know which frequencies will be used at which point. This approach minimizes jamming and eavesdropping because it is difficult for an outside listener to know what frequencies will be used next.

The FHSS version provides speeds of 1 Mbps and 2 Mbps. The DSSS version provides speeds of 1 Mbps, 2 Mbps, 5.5 Mbps, and 11 Mbps. A DSSS 20Mbps version is under development. Because greater distance from the computer to the AP can weaken the signal, making interference from microwave ovens, cordless phones, and baby monitors a major problem, both FSSS and DSSS have the ability to automatically seek changes in speeds. In good conditions at close range, for example, DSSS may provide 11 Mbps, but as the distance increases between the AP and the computer, or if interference increases, the transmission rate may back down to 1 Mbps.

It is important to remember that both versions are shared media implementations, meaning that all devices in the WLAN share the one logical circuit. So if the WLAN has 10 computers, and the speed is reduced to 1 Mbps due to inference, there may be noticeable response time delays. In this example, if we assume a 90% throughput rate (assuming controlled access), this would mean that each computer has about 90Kbps.

IEEE 803.11a

As we write this, the IEEE 803.11a standard has not been completely defined, but it should be by the time you read this. IEEE 803.11a is expected to operate in the 5 GHz range, meaning that it is capable of much higher transmission speeds but also will likely be more susceptible to interference. The initial standard will likely provide a raw data rate of 54Mbps, but will probably only average 20 Mbps in practice.

OTHER WIRELESS TECHNOLOGIES

There are two other wireless technologies that may become more common: infrared technologies and Bluetooth.

Infrared Wireless LANs

In general, infrared wireless LANs are less flexible than the radio-based IEEE 802.11 WLANs because most require a direct line of sight between the transmitters and receivers in the same manner as your TV remote control. Transmitters and receivers are usually mounted in fixed positions to ensure the line of sight, so most infrared LANs are wireless only between the hubs. The NICs inside the computers are connected via traditional wires to a network hub which contains the transmitter (see Figure 6-9). The hubs use wireless transmission to communicate from hub to hub.

--- Figure 6-9 old 7-10 --

The primary advantage of a wireless LAN is the reduction in wiring. Infrared-based LANs are sometimes used for communication between buildings where installing underground cable would be expensive. In an old building where wiring is difficult and costs are extremely high, wireless LANs offer a low cost alternative by enabling communication without the installation of cables. Most infrared wireless systems provide 1-4 Mbps, but some provide 100 Mbps or more.

A new version of infrared, called diffuse infrared, operates without a direct line of sight by bouncing infrared light around a room. Most diffuse infrared systems have extremely short ranges (usually only 50-75 feet) and will only operate in the same room because the light cannot travel through walls.

Bluetooth

Bluetooth is strikingly different from the wireless LAN technologies discussed above. Its goal is to provide seamless networking of devices in a very small area (up to 30 feet, soon to increase to about 300 feet). Bluetooth is an unofficial standard but may soon be standardized as IEEE 802.15.

Bluetooth devices are small (about 1/3 of an inch square) and cheap (currently priced at about $30 but expected to quickly to drop to $4 or even lower. A Bluetooth network is called a piconet and consists of no more than 8 devices, but can be linked to other piconets to form a larger network. A typical application might be to connect a mouse to a computer, to connect a telephone headset to a base unit (up to 3 digital voice circuits can be in use simultaneously) , or to link your Palm handheld computer with your car so that your door unlocks and automatically opens as you approach.

Bluetooth provides a 1 Mbps shared circuit, but its data link protocol is inefficient, so it provides only a 780 Kbps throughput. It uses FHSS in the same crowded 2.4 Ghz range used by IEEE 803.11b, so interference with 803.11b devices can significantly reduce throughput in both the Bluetooth and 802.11b networks. It uses controlled access media access control, with one device, called a master, polling the other devices, called slaves.

BACKBONE NETWORK ARCHITECTURES

While there are an infinite number of ways in which network designers can build backbone networks, there are really only four fundamental architectures that can be combined in different ways. These four architectures are routed backbone (routers that move packets based on network layer addresses), bridged backbones (bridges that move packets based on data link layer addresses), collapsed backbones (switches that move packets based on data link layer addresses), and Virtual LANs (switches that move packets through LANs that are built virtually, not using physical location).

These four architectures are mixed and matched to build sets of backbone networks. Before we discuss these four architectures, we first must discuss the way in which network designers think about backbone designs and how to combine them; that is, the different layers of backbones that exist in most organizations today.

Backbone Architecture Layers

Network designers often think about three distinct technology layers [2] when they design backbone networks. The layer closest to the users is the access layer, the technology used in the LANs attached to the backbone network as described in the previous chapter (e.g., 100Base-T, switched 10Base-T, wireless Ethernet). See Figure 7-5 While the access layer is not part of the backbone network, the technologies used in the LANs (or access layer) can have major impacts on the design of the backbone.

-- Figure 7-5 --

The distribution layer is the part of the backbone that connects the LANs together. This is the part of the backbone that contains the "TCP/IP gateways" described in Chapter 5. It usually runs throughout one building.

The core layer is the part of the backbone that connects the different backbone networks together, often from building to building. The core layer is technologies used in the campus network or the enterprise network. Some small organizations are not large enough to have a core layer; their backbone

2. Try not to be confuse the five basic layers in the network model (application layer, transport layer, and so on) with the layers of backbone technology we are describing here. They are different. We would have preferred to use a different work than "layer" to describe these, but unfortunately that is the term used in industry.

spans only the distribution layer. Other organizations are large enough that they have a core network at several locations that are in turn connected by WANs.

In the sections that follow, we describe the four basic BN architectures and discuss at which layer they are often used. We will focus on TCP/IP networks when comparing these four architectures. We assume that you are comfortable with the material on TCP/IP in Chapter 5; it you are not, you may want to go back and review the last section of the chapter entitled **TCP/IP Example** before you continue reading.

Routed Backbone

Routed backbones move packets along the backbone based on their network layer address (i.e., layer 3 address). The most common form of routed backbone uses a bus topology (e.g., using Ethernet 100Base-T). Routed backbones are sometimes called subnetted backbones or hierarchical backbones and are most commonly used to connect different buildings within the same campus network (i.e., at the core layer).

Figure 7-6 illustrates a routed backbone used at distribution layer (because it is simpler to explain how they work using the distribution layer than the core layer). A routed backbone is the basic backbone architecture we used to illustrate how TCP/IP worked in Chapter 5. There are a series of LANs (access layer) connected by routers or layer 3 switches to a single shared media backbone network. Each of the LANs are a separate subnet. Message traffic stays within each subnet unless it specifically needs to leave the subnet to travel elsewhere on the network, in which case the network layer address (e.g., TCP/IP) is used to move the packet.

---- Figure 7-6 old 10-11 --

Each LAN is usually a separate entity, relatively isolated from the rest of the network. There is no requirement that all LANs share the same data link layer. One LAN can use Ethernet, while another uses another technology. Each LAN can contain its own server designed to support the users on that LAN, but users can still easily access servers on other LANs over the backbone as needed.

Th primary advantage of the routed backbone is that it clearly segments each part of the network connected to the backbone. Each segment (usually a LAN or another backbone) has its own subnet addresses that can be managed by a different network manager. Each segment off the backbone also can use different data link layer technologies.

There are two primary disadvantages to routed backbones. First, the routers in the network impose time delays. Routing takes more time than bridging or switching, so routed networks can sometimes be slower.

Second, routed networks require a lot of management. Establishing separate subnet addresses for each LAN is time-consuming, and requires a large set of TCP/IP addresses. Any time a computer is moved from one LAN to another, it must be reconfigured (unless the network is using dynamic addressing which imposes costs of its own).

Bridged Backbone

Bridged backbones move packets along the backbone based on their data link layer address (i.e., layer 2 address). The most common form also uses a bus topology. They were common in the distribution layer, but their use is declining; few organizations install bridged networks because they have major performance problems as we shall shortly see. Bridged backbones are sometimes called flat backbones.

Figure 7-7 illustrates a distribution layer bridged backbone with a bus topology. This figure shows the same series of LANs as in Figure 7-6, but now the LANs are connected by bridges or layer 2 switches to the single shared media backbone network. As you can see, a bridged backbone looks very similar to a routed backbone. With a bridged backbone, however, the entire network (backbone and all connected network segments) are on the same subnet. All LANs are part of the same overall network and all must have the same data link layer protocol. This is in sharp contrast to the routed backbone in which the LANs are isolated and may be different.

---- Figure 7-7 old 10-12 --

Bridged backbones have several distinct advantages and disadvantages compared to routed backbones. First, since bridges tend to be less expensive than routers, they are often cheaper. Second, they are usually simpler to install because the network manager does not need to worry about building many different subnets and assigning a whole variety of different subnet masks and addresses in each part of the network. However, since the backbone and all attached networks are considered part of the same subnet, it is more difficult to permit different individuals to manage different parts of the network (e.g., LANs); a change in one part of the network has the potential to significantly affect all other parts. Also, it is possible to run out of IP addresses if the entire network has many computers.

The single most major problem is network speed. Bridging is faster than routing, so one might expect the bridged backbone to be faster. For small networks, this is true. For large networks, it is not. Bridged backbone are slower than routed backbones. Since bridged backbone and all networks connected to them are part of the same subnet, broadcast messages (e.g., address requests) must be permitted to travel everywhere in the backbone. This means, for example, that a computer in one LAN attempting to find the data link layer address of a server in the same LAN will issue a broadcast message that will travel to every computer on every LAN attached to the backbone. (In contrast, on a routed backbone such messages would never leave the LAN in which they originated.)

There are many different types of broadcast messages other than address requests (e.g., a printer reporting it is out of paper, a server about to be shut down). These broadcast messages quickly use up network capacity in a large bridged network. The result is slower response times for the user. In a small network, the problems are not as great, because there are fewer computers to issue such broadcast messages.

Collapsed backbone

Collapsed backbones are probably the most common type of backbone network used in the distribution layer (i.e., within a building); most new building backbone networks designed today use collapsed backbones. They also are making their way into the core layer as the campus backbone, but routed backbones still remain common.

Collapsed backbone networks use a star topology with one device, usually a switch, at its center. Figure 7-8 shows a collapsed backbone connecting the same series of LANs. Here, the backbone circuit and set of routers or bridges is replaced by one switch and a set of circuits to each LAN. The collapsed backbone has more cable, but fewer devices. There is no backbone cable. The "backbone" exists only in the switch, which is why this is called a collapsed backbone.

---- Figure 7-8 old 10-13 --

There are two major advantages to collapsed backbones. First, performance is improved. With the traditional backbone network, the backbone circuit was shared among many LANs (eight LANs, in the case of Figure 7-8); each had to take turns sending messages. With the collapsed backbone, each connection into the switch is a separate point-to-point circuit. The switch enables simultaneous access, so that several LANs can send messages to other LANs at the same time. Throughput is increased significantly, often by 200 percent to 600 percent, depending upon the number of attached LANs and the traffic pattern.

Second, there are far fewer networking devices in the network. In Figure 7-8, one switch replaces eight routers. This reduces costs and greatly simplifies network management. All the key backbone devices are in the same physical location, and all traffic must flow through the switch. If something goes wrong or if new cabling is needed, it can all be done in one place.

Collapsed backbones often, but not always, have two important disadvantages that are the same as those for bridged networks. Because data link layer addresses are used to move packets, there is more broadcast traffic flowing through the network and it is harder to isolate and separately manage the individually attached LANs. Layer 3 switches can use the network layer address, so future collapsed backbones built with layer 3 will not suffer from this problem.

Collapsed backbones also have two relatively minor disadvantages. First, they use more cable, and the cable must be run longer distances, which often means that fiber optic cables must be used. Second, if the switch fails, so does the entire backbone network. However, if the reliability of the switch has the same reliability as the reliability of the routers in Figure 7-6, then there is less chance of an failure (because there are fewer devices to fail). For most organizations, these disadvantages are outweighed by benefits offered by collapsed backbones.

Rack-based Collapsed Backbones

Most organizations now use collapsed backbones in which all network devices for one part of the building are physically located in the same room, often in a rack of equipment. This form of collapsed backbone is shown graphically in Figure 7-9. This has the advantage of placing all network equipment in one place for easy maintenance and upgrade, but does require more cable. In most cases, the cost of the cable itself is only a small part of the overall cost to install the network, so the cost is greatly outweighed by the simplicity of maintenance and the flexibility it provides for future upgrades.

---- Figure 7-9 --

The room containing the rack of equipment is sometimes called the main distribution facility (MDF) or central distribution facility (CDF). See Figure 7-10. The cables from all computers and devices in the area served by the MDF (often hundreds of cables) are run into the MDF room. Once in the run they are connected into the various devices. The devices in the rack are connected among themselves using very short cables called patch cables.

---- Figure 7-10 photo of racks --

With rack-based equipment, it becomes simple to move computers from one LAN to another. In the traditional routed backbone design as shown in Figure 7-6, for example, all the computers in the same

general physical location are connected to the same hub and thus share the capacity of the hub. While this often works well, it can cause problems if many of the computers on the hub are high traffic computers. For example, in Figure 7-6, if all the busy computers on the network are located in the upper left area of the figure, the hub in this area may become a severe bottleneck.

With a MDF, all cables run into the MDF. If one hub becomes overloaded, it is straightforward to unplug the cables from several high-demand computers from the overloaded hub and plug them into one or more less-busy hubs. This effectively spreads the traffic around the network more efficiently and means that network capacity is no longer tied to the physical location of the computers; computers in the same physical area can be connected into very different network segments.

Chassis-based Collapsed Backbones

Sometimes a chassis switch is used instead of a rack. A chassis switch enables users to plug modules directly into the switch. Each module is a certain type of network device. One module might be a 16-port 10Base-T hub, another might be a router, while another might be an 4-port 100Base-T switch, and so on. The switch is designed to hold a certain number of modules and has a certain internal capacity, so that all the modules can be active at one time. For example, a switch with five 10Base-T hubs, two 10Base-T switches (with 8 ports each), a 100Base-T switch (with 4 ports) and a 100Base-T router would have to have an internal switching capacity of at least 710 Mbps (5 x 10Mbps + 2 x 8 x 10Mbps + 4 x 100 Mbps + 100 Mbps = 710 Mbps).

The key advantage of chassis switches is their flexibility. It becomes simple to add new modules with additional ports as the LAN grows, and to upgrade the switch to use new technologies. For example, if you want to add gigabit Ethernet or ATM (discussed below) you simply lay the cable and insert the appropriate module into the switch.

-- Figure 7-11 --

Virtual LAN

For many years, the design of local area networks remained relatively constant. However, in recent years, the introduction of high speed switches has begun to change the way we think about local area networks. Switches offer the opportunity to design radically new types of LANs. Most large organizations today have traditional LANs, but many are considering the virtual LAN (VLAN), a new type of LAN/BN architecture made possible by intelligent, high speed switches.

VLANs are networks in which computers are assigned to LAN segments by software, rather than by hardware. In the section above, we described how in rack-based collapsed backbone networks, a computer could be moved from one hub to another by unplugging its cable and plugging it into a different hub. VLANs provide the same capability via software so that the network manager does not have to unplug and replug physical cables to move computers from one segment to another.

VLANs are often faster and provide greater opportunities to manage the flow of traffic on the LAN and BN than the traditional LAN and routed BN architecture. However, VLANs are significantly more complex so they usually are used only for large networks. There are two basic approaches to designing VLANs: single switch VLANs and multi-switch VLANs.

Single Switch VLAN

A single switch VLAN means that the VLAN operates only inside one switch. The computers on the VLAN are connected into the one switch and assigned by software into different VLANs (see Figure 7-12). The network manager uses special software to assign the dozens or even hundreds of computers attached to the switch to different VLAN segments. The VLAN segments function in the same way as physical LAN segments; the computers in the same VLAN act as though they are connected to the same physical switch or hub. For example, broadcast messages sent by computers in a VLAN segment are sent only to the computers on the same VLAN. VLANs can be designed so that they act as though computers are connected via hubs (i.e., several computers share a given capacity and must take turns using it) or via workgroup switches (i.e., all computers in the VLAN can transmit simultaneously). While switched circuits are preferred to the shared circuits of hubs, buying VLAN switches with the capacity to provided a complete set of switched circuits for hundreds of computers is more expensive than those that permit shared circuits.

---- Figure 7-12 --

We should also note that it is possible to have just one computer in a given VLAN. In this case, that computer has a dedicated connection and does not need to share the network capacity with any other computer. This is commonly done for servers.

There are four ways in which computers attached to VLAN switches can be assigned to the specific virtual LANs inside them. The first approach, used by port-based VLANs (also called Layer 1 VLANs), uses the physical layer port number on the front of the VLAN switch to assign computers to VLAN segments. Each computer is physically cabled into a specific port on the VLAN switch. The network manager uses special software provided by the switch manufacturer to instruct the switch which ports are assigned to which VLAN. This means that the network manager must know which computer is connected to which port.

The second approach, used by MAC-based VLANs (also called Layer 2 VLANs), uses the data link layer address to form the VLANs. The network manager uses special software to instruct the switch which incoming data link layer addresses are assigned to which VLAN segment. The advantage of a layer 2 VLAN is that they are simpler to manage when computers are moved. If a computer is moved in a layer 1 VLAN, then the network manager must reconfigure the switch to keep that computer in the same VLAN because the computer has moved from one port to another. With a layer 2 VLAN, no reconfiguration is needed. Although the computer may have moved from one port to another, it is the permanently assigned data link layer address that is used to determine which VLAN the computer is on.

The third approach, used by IP-based VLANs (also called Layer 3-VLANs), uses the network layer address to form the VLANs. As before, the network administrator uses special software to instruct the switch which network layer addresses are assigned to which VLAN. Layer 3 VLANs reduce the time spent reconfiguring the network when computers move in the same way as layer 2 VLANs. Layer 3 VLANs tend to be a bit slower at processing each message than layer 2 VLANs because processing layer 3 protocols is slightly slower than processing layer 2 protocols.

The fourth approach, used by application-based VLANs (also called policy-based VLANs or Layer 4 VLANs), uses the type of application indicated by the port number in the TCP packet in combination with the network layer addresses to form the VLAN groups. As before, the network administrator uses special software to instruct the switch which types of packets from which addresses are assigned to which VLAN. This process is very complex because the network manager must decide on a variety of different factors in forming the VLANs. The advantage is a very precise allocation of network capacity. Now VLANs can be formed to allocate a certain amount of network capacity for Web browsing to certain individuals, so much to Web browsing for others, so much to transaction processing, and so on. In this way, the network manager can restrict the amount of network capacity used by potentially less productive applications (e.g., Web surfing) and thus provide much better allocation of resources.

Multi-switch VLAN

A multi-switch VLAN works the same way as a single switch VLAN, except that now several switches are used to build the VLANs (see Figure 7-13). In this case, the switches must be able to send packets among themselves in a way that identifies the VLAN to which the packet belongs. There are two approaches to this.

---- Figure 7-13 --

The first approach is to use a proprietary protocol that encapsulates the packet (i.e., a protocol that is not standard, but instead is used only by specific companies). In this case, when a packet needs to go

from one VLAN switch to another VLAN switch, the first switch puts a new VLAN packet around the outside of the Ethernet packet. The VLAN packet contains the VLAN information and is used to move the packet from switch to switch within the VLAN network.. When the packet arrives at the final destination switch, the VLAN packet is stripped off and the unchanged Ethernet packet inside is sent to the destination computer.

The other approach is to modify the Ethernet packet itself to carry the VLAN information. IEEE 802.1q is an emerging standard that inserts 16-bytes of VLAN information into the normal IEEE 802.3 Ethernet packet. In this case, when a packet needs to go from one VLAN switch to another VLAN switch, the first switch replaces the incoming Ethernet packet with an 802.1q packet that contains all the information in the original 802.3 Ethernet packet, plus 16-bytes of VLAN information. The additional VLAN information is used to move the packet from switch to switch within the VLAN network. When the packet arrives at the final destination switch, the IEEE 802.1q packet is stripped off and replaced with a new Ethernet packet that is identical to the one with which it entered the VLAN and is sent to the destination computer.