Network Basics

In order to attain a better understanding of wireless technologies, it's important to grasp computer-networking basics. Wireless networking is based on the same concepts as those of wired networks, so a review of the underlying technologies used in wired networks will help clarify some of the performance and security issues relating to wireless ones.

A Layered Approach

One of the basic principles in modern computing involves dividing the overall operation into layers. Each layer performs a general function in a specific way. Often, there are competing approaches in dealing with a given task or sub-task, so compartmentalizing each function in layers allows for changing out one approach for another without having to rework the entire system.

Networks especially benefit from a layered design. The classic conceptual network model was expressed by the Open Systems Interconnect (OSI) initiative. Although OSI did not catch on as a real-world network scheme, its conceptual model helps users understand the networking technologies in use today. The terminology of the OSI reference model finds common use in networking literature and documentation.

Ethernet

In order for a network to function, a set of media access rules must exist to govern the flow of data. Operating at Layer 2 (see Figure 1), these rules define how data will be packaged as well the basic rules of transmission. Ethernet is the Layer-2 protocol most often used in local area networks. While wireless networks use a slightly different set of media access rules, Ethernet serves as a good example of a low-level network protocol.

IFLANET—OSI for Libraries: Standards to Services (1992)
www.ifla.org/VI/5/reports/rep3/312.htm

Addressing–One of the basic requirements of Ethernet is that all devices have a unique address. Called the MAC (media access control) address, this 48-bit binary number–typically expressed as 12 digits in hexadecimal–must be unique worldwide. Each Ethernet equipment manufacturer is assigned a code used in the prefix of the MAC address and is responsible for assigning unique addresses for every individual device it manufactures, burning that number into its firmware. My laptop, for example, has the address 00-0f-66-06-b3-68.

Packets–In order for information to travel on a network, it must be broken down into structured packets. A long information stream will be chopped up into smaller portions, each of which is bracketed by structured fields at the beginning and end. The packet headers function much like an envelope in the postal system, carrying the MAC address of the originating device, the address of the destination, packet length data, as well as CRC (cyclic redundancy code) checksums that can be used to validate contents.

CSMA/CD–Ethernet follows a set of media access rules called Carrier Sense Multiple Access with Collision Detection, or CSMA/CD. The formal definitions for Ethernet have been established by the IEEE (Institute for Electrical and Electronics Engineers) 802.3 committee and apply to all its versions.

One can compare CSMA/CD to the old-fashioned party-line telephone. Remember how the neighborhood telephone system worked in our grandparents' time? Many houses would share the same phone line, making it possible for each house to listen in on all the conversations. Each household had a distinctive ring. While you were only supposed to pick up the phone when you heard your distinct ring, eavesdropping was common.

Ethernet follows the same rules:

- **Carrier Sense**–Devices must indicate the network is active before transmission. It's like listening for the dial tone.
- **Multiple Access**–Lots of conversations can take place on the line at the same time.
- **Collision Detection**–If two devices transmit at the same time, a collision occurs and all the data packets involved are lost. With Ethernet, it's okay to transmit at any time, but the network must detect collisions. When a collision happens, the devices involved will wait a random number of milliseconds and retransmit. Making the interval random avoids causing a new collision, as would happen if they waited a fixed period. While collisions involve some inefficiency, it turns out it's better to incur a certain number of collisions than to implement some scheme in which devices wait for a quiet time.

It's interesting to compare Ethernet, a free-wheeling West Coast network protocol (invented in 1973 by Bob Metcalf at PARC, the Xerox Palo Alto Research Center) with IBM's Token Ring, which is an East Coast, more deterministic protocol. Both perform the same task–operating on OSI Layer 2–but the Ethernet strategy for moving data involves putting as many packets on the media as quickly as possible. Ethernet performs error-checking after the fact. Token Ring avoids errors, requiring devices to wait until the safe delivery of the packet can be guaranteed. The token passing protocol, essentially, assigns each device on the network a numbered position on the logical ring. A token packet circulates around the ring. Each device can transmit or receive data only when the token passes its position. Although this approach avoids errors, a portion of the potential bandwidth is consumed by devices waiting their turn.

In the end, Token Ring networks faded away, and these days Ethernet dominates.

In its early prototype, Ethernet operated at 2.94 Mbps. The first commercial version operated at 10 Mbps. Since then, Ethernet's bandwidth capabilities have climbed from 10 Mbps, to 100 Mbps, to 1 Gbps, with the fastest version in use today at 10 Gbps. Speed increases will continue. Ethernet's approach–of filling the media as fast as possible and fixing errors after the fact–has proven to be a winning strategy.

It's interesting, however, to observe that WiMAX (IEEE 802.16), a wireless technology emerging for broadband Internet access, follows a set of media access rules similar to the deterministic Token Ring protocols.

- **Layer 1**–Physical (electrical characteristics of cabling)
- **Layer 2**–Data Link (Ethernet) Ethernet cards, hubs, switches; (802.11)
- **Layer 3**–Network (IP) Routers
- **Layer 4**–Transport (TCP/UDP) error recovery, transfer of data
- **Layer 5**–Session
- **Layer 6**–Presentation
- **Layer 7**–Application

Figure 1
OSI Reference Model

Network Segmentation

One of the fundamental concepts in the Ethernet world involves segments. The early Ethernet networks used a thick, rigid coaxial cable (ThickWire). About the same thickness as a garden hose, this type of cable was infamously difficult to manage. This type of Ethernet–called "10Base5"–operated at 10 Mbps supporting a maximum length of 500 meters. These networks followed a bus, or linear, topology, meaning that each device connected directly to the cable as it snaked through the room or building. To function properly, a 50-Ohm terminator was placed at each end of the cable. Each device connected to the network through a physical tap drilled into the cable, onto which a transceiver could be connected. Each length of cable could support up to 100 devices and was called a "network segment." Multiple Ethernet segments could be tied together through repeaters, bridges, or routers.

Ethernet is a shared-media network. All the devices within a segment have access to all the data as it travels through the segment. In order to pick out the packets it's supposed to receive, each device must scan the destination address field in the headers of all packets. Ethernet rules forbid a device from opening packets if the destination address does not match its own MAC address, unless it is a broadcast packet intended for all the devices on the segment.

Eavesdropping: The Bane of Security

Theoretically, a device on a network should only receive and open packets that have its own MAC address in the header–just like it was polite to pick up the party-line telephone *only* when you heard the distinct ring for your own house. In the real world, however, eavesdropping happens. This is the characteristic of Ethernet networks that must be taken into consideration in regard to privacy and security.

In order to guard against eavesdropping, early Ethernet transceivers were built to open packets that only matched their own MAC addresses. These devices prevented you from viewing packets intended for others on your segment. Early networks were fragile, and network engineers needed the ability to look at all the packets to troubleshoot problems and perform diagnostics. To support this need, a special type of Ethernet device was created that could operate in "promiscuous" mode, allowing it to open and view the contents of all the packets.

IEEE 802 LAN/MAN Standards Committee
www.ieee802.org/

Part of a network engineer's standard equipment was a network "sniffer," a device that could capture all the packets on the network, calculate performance and error statistics, and display the contents of each packet in human-readable form. At first, only a small number of cards could be set to operate in promiscuous mode, but eventually all cards were built with that capability.

Today, the software needed to view and analyze Ethernet packets is commonplace. Given the widespread capability to intercept and view Ethernet packets, one cannot assume that data transmitted on an Ethernet is private or secure unless additional protective measures are taken.

Privacy and security will be covered in detail later in this report, but keep in mind that the eavesdropping "factor" influences many security issues.

Network Connectivity Infrastructure

Linear Topology—Previously, it was noted the earliest Ethernet networks used ThickWire cabling. This cabling was organized in a linear topology; all devices tapped into a length of cable that spanned the network. The immediate successor to ThickWire 10Base5 was ThinWire Ethernet, or 10Base2. This flavor of Ethernet used flexible RG-58 cabling, using simple twist-on BNC connectors. These networks still followed a linear topology, with each device connected to the next in a daisy-chain configuration.

Although this version of Ethernet was much easier to work with, it was still fragile. One faulty connection would disrupt the entire segment. Like its predecessor, each end of the segment had to be terminated, and multiple segments could be connected through repeaters, bridges, or routers.

Fortunately, networks based on runs of coaxial cables have long since been retired. This bit of history is worth mentioning, however, because many of the networking principles of those early networks carry forward to the networks in place today.

Star Topology—Wired networks today follow a "star" topology, meaning that each device on the network has its own dedicated cable. Typically, each floor in a building will have a wiring closet, and cables will be installed between a patch panel in the wiring closet to a wall plate near each computer location. Short patch cables connect each device to the wall plate.

Ethernet operates on unshielded twisted-pair (UTP) cabling. These cables consist of four pairs of twisted copper wire, with each of the four pairs also twisted around each other. The tightness of the twists, as well as the gauge and quality of the wires, affect the electrical characteristics of the cable and its ability to effectively transmit network signals. Category 5 UTP describes cable that conforms to a set of specifications designed to carry Ethernet traffic up to 100 Mbps. Category 5e or Category 6 is recommended for 1Gbps and faster Ethernet.

While this star topology involves considerably more cable than the linear bus topology, it is much more fault tolerant. A cable fault usually affects only a single device.

Digital Stoplights

Several different types of devices manage the flow of traffic.

An **Ethernet hub** is a multi-port unit that connects computers to a network. Hubs physically connect computers together in a star topology, each attached through a dedicated cable. Logically, an Ethernet hub functions as an Ethernet segment. Each of the devices connected to the hub share the media, compete for collisions, and potentially eavesdrop. Older Ethernet hubs operated at 10Mbps. Today 100 Mbps and 1Gbps hubs are available.

Ethernet switches are higher-performance devices. From the outside, an Ethernet switch looks much like a hub. Internally, though, the Ethernet switch works quite differently.

I previously noted that all the devices attached to a hub form a logical Ethernet segment. Each port in a switch, however, effectively functions as an independent segment. Using high-speed switching technologies, packets are delivered only to the specific port, unlike hubs that broadcast all packets to all ports. Ethernet switches come in performance levels varying from 100 Mbps to 10 Gbps and are priced accordingly.

A **router** is a device that directs the flow of packets from one network to another. Technically, a router is a decision-making device that operates at OSI Level 3 (see figure 1). It's the router's job to determine the most efficient path for data packets to travel to reach their destinations. Routers can be used to connect multiple LANs in an organization or to connect a LAN to the Internet.

```
C:\>ipconfig /all
Windows IP Configuration

  Host Name . . . . . . . . . . . . : Breeding4
  Primary Dns Suffix  . . . . . . . :
  Node Type . . . . . . . . . . . . : Hybrid
  IP Routing Enabled. . . . . . . . : No
  WINS Proxy Enabled. . . . . . . . : No
  DNS Suffix Search List. . . . . . : BellSouth

Ethernet adapter Local Area Connection:

  Connection-specific DNS Suffix  . : BellSouth
  Description . . . . . . . . . . . : Realtek RTL8139/810x Family Fast Ethernet NIC
  Physical Address. . . . . . . . . : 00-11-2F-92-2B-5E
  Dhcp Enabled. . . . . . . . . . . : Yes
  Autoconfiguration Enabled . . . . : Yes
  IP Address. . . . . . . . . . . . : 192.168.0.109
  Subnet Mask . . . . . . . . . . . : 255.255.255.0
  Default Gateway . . . . . . . . . : 192.168.0.1
  DHCP Server . . . . . . . . . . . : 192.168.0.1
```

Figure 2
Sample: Author's Windows IP Configuration

Large networks require many different pieces of communication equipment. Today, most networks consist of hubs and switches. A low-performance option would consist of all end-user devices connecting to hubs.

With larger networks, better performance can be accomplished by using one or more switches to divide the network. Switches form the network backbone, with hubs connected to each port of the switch.

The highest level of performance and security can be gained by using switches throughout the network. Each end-user device connects to its own port on a switch, virtually eliminating all possibilities for collisions on the network.

TCP/IP

TCP/IP stands as the dominant suite of network protocols. Made popular on the early Internet and on Unix systems, today practically all local and wide-area networks rely on TCP/IP. Options such as Novell's IPX, AppleTalk, and other proprietary protocols once provided considerable competition for TCP/IP, but more and more organizations are migrating to an IP-only environment.

TCP/IP Configuration Details

In a TCP/IP network, each device needs to be configured with a few essential settings. Each device on the network must be assigned its own IP address. In an independent self-contained network, each device will have a unique device within that network. If the network is connected to the Internet, it must have a globally unique IP address.

Other configuration details include specifying the proper subnet mask, which tells the device what portion of its IP address describes the sub network on which it resides (versus its own unique address). Each station must also be configured with a default gateway, an IP address necessary if the computer's packets are to reach external networks. The default gateway usually refers to the IP address of the nearest available router.

Dynamic Host Configuration Protocol

Having to manually configure each computer on a large network with IP settings can be an unwieldy task for a network administrator. In the early days of networking, that's exactly what had to be done. I can remember having to *physically* visit every computer on the network in order to use a text editor to adjust its net.cfg file. I had to be very careful to avoid assigning the same IP address to two computers, an error that would keep one or the other from connecting to the network.

Eventually, network administrators were able to use a process called "bootp" to automatically assign IP addresses. Using the bootp process, an administrator simply maintained a table (hosted on a server) that configured each computer's MAC address to the administrator-assigned IP address.

Today, a much more powerful process called "Dynamic Host Configuration Protocol," or simply DHCP, performs the heavy lifting of network configuration. DHCP makes the administration of an IP network amazingly easy for the network administrator and for each computer user.

For the network administrator, a DHCP server can be set up to make IP assignments out of a pool of eligible addresses. For networks that need more control, it's also possible to lock each MAC address to a specific IP address. DHCP does a lot more than dole out IP addresses; it also sets each computer's subnet mask, default gateway, default DNS server, and even performs some tuning parameters to ensure efficient communications.

For the end user, DHCP makes connecting to a network a snap. If your computer is equipped with a DHCP client—and all recently manufactured computers are—just about all you have to do is plug in. Behind the scenes, the DHCP client will send out a packet requesting information, and if all goes well, the DHCP server will respond with its configuration assignments. In some cases, you might need to restart the computer. In some more tightly controlled networks, the DHCP server might only respond to pre-registered computers, so you may have to consult the network administrator.

In the DHCP world, you don't buy your place on the network, you lease it. When a DHCP server makes an assignment, it stipulates the lease terms. Its duration may be a few hours or a few days. At the expiration of that lease, the client computer must renew the lease. In most cases the DHCP server renews the existing parameters without the computer user having to be aware that anything happened behind the scenes.

For the rare occasions in which something goes wrong, most computers have commands you can use to view your IP configuration and manually request a new lease. As an example, let's take a look at what's available for Microsoft Windows (XP, W2K, NT, Win95).

These are console commands, so begin with Start → Run, and type "cmd" into the field. The "ipconfig" command is available to view your current settings and make requests to a DHCP server.

- **ipconfig /all:** This shows the details of each network interface on your computer (see Figure 2).
- **ipconfig /release:** This tells your computer to relinquish its current settings.
- **ipconfig /renew:** This sends out a request for a new lease.

In practical terms, when you can't seem to connect to a network, an "ipconfig /release" followed by a "ipconfig /renew" usually works wonders.