Chapter 5

Traffic regulators:

Network interfaces, hubs, switches, bridges, routers, and firewalls

In the previous chapter, you learned how network hardware and software follow rules, called protocols, in order to convey information in an orderly fashion. In this chapter, we focus on one particular set of network hardware devices—network interfaces, hubs, switches, bridges, routers, and firewalls. These devices interconnect individual computers and ensure that they communicate efficiently.

In order to understand the role of these devices, it is helpful to consider once again our analogy between networks and information highways. In this analogy, the network manager is like a city planner. The manager defines the resources that will be available to the network town—for example, student laboratories, teachers' computers, accounting and grade systems, Web servers, electronic mail systems, and so forth. The manager also decides where to locate those resources and the capacity of the roads leading to them. For example, you network will contain small streets (less costly to implement and maintain) for areas with low traffic, large streets with traffic lights for areas of moderate traffic, and highways with limited access for areas of high traffic density. Finally, the manager considers how and when to restrict access to some resources that should remain private.

Network interfaces, hubs, bridges, switches, routers, and firewalls work together in a number of ways to create these different kinds of network roadways.

- First, these devices regulate the speed at which your network information travels. For example, you can interconnect individual computers through the use a device called hub (to which individual computers are connected like the spokes of a bicycle wheel). Depending on the speed of the hub that you purchase, the interconnected computers will operate at 10 Mbps (relatively low) or 100 Mbps (moderate) or 1000 Mbps (very high). Network interfaces, bridges, switches, routers, and firewalls, offer the same speed options.
- Second, these devices manage the flow of traffic, opening, closing, or directing it to specific streets as the need arises. For example, a device called a router suggests the most efficient route for network traffic to travel to its destination. As another example, a device called a switch opens a circuit, or connection, directly between two communicating computers and

keeps other computers from interfering with the connection. By rotating this connection among each pair of communicating computers, switches minimize contention and make very efficient use of the network roadways.

- Finally, some of these devices help you protect sensitive information. For example, bridges, routers, and firewalls can look at every bit of network information and decide, based on its destination or other internal information, whether to permit or deny the information to reach a specific resource.

5.1 Network interfaces

A network interface is a device that connects a client computer, server, printer or other component to your network. Most often, a network interface consists of a small electronic circuit board that is inserted into a *slot* inside a computer or printer. Alternatively, some computers, printers, or other devices include network interfaces as part of their main circuit boards (motherboards). In either case, the network interface provides two important services—it connects your computer physically to your network, and it converts information on your computer to and from electrical signals for your network.

The network interface connects to your network via a small receptacle called a *port*. For wired networks, you insert the network cable into this port. Alternatively, for wireless networks, the port includes a transmitter/receiver that sends/receives radio signals.

Besides providing physical connections, network interfaces convert information on your computer into electrical signals of appropriate shape and transmission speed for your network. All network interfaces on your network must conform to a common physical and data link level protocol in order for their electrical signals to be compatible (and therefore to exchange information successfully). For example, if you are running a 10BaseT Ethernet network, then all of your computers, printers, and servers must contain 10BaseT Ethernet network interfaces. Some network interfaces accommodate more than one physical or data link protocol—for example, combinations of 10BaseT and 100BaseTX protocols (most network interfaces automatically sense the appropriate speed). Combination 10BaseT/100BaseTX network interfaces provide growing room if you plan to upgrade your 10BaseT network to 100BaseTX sometime soon (by installing 100BaseTX components on the rest of your network).

On your computer, your network interface works closely with special software called drivers—software modules that control the network interface. When you install a network interface, you also install its drivers. Generally, drivers are available on diskette from the manufacturer of the network interface, on the master disks for your operating system (Macintosh OS or Windows), or from the Internet. The network interface usually includes instructions for installing both the interface and the drivers. From time to time, you should update the drivers on each computer or printer by installing new software from the network interface manufacturer. Updated drivers contain performance enhancements and bug fixes that may benefit your network.

Each network interface is associated with a unique address called its *media access control (MAC)* address. The MAC address helps route information within your local area network and is used by interconnecting devices such as switches and bridges. The exact role of network interfaces with regard to MAC addresses varies a bit among different networks. On Ethernet networks, each network interface receives a unique MAC address when it is manufactured. When the network interface is installed into a *slot* or onto the motherboard of a computer or printer, the interface MAC address becomes the address for the computer or printer. On LocalTalk networks, network interfaces work a bit differently. Instead of receiving permanent addresses, they assist in a more complex addressing scheme that assigns each Macintosh an address dynamically when it connects to a LocalTalk network. In any case, assigning the MAC address is an important task executed by the network interface.

The MAC address is just one of several network addresses assigned to each network computer, server, or peripheral. Another network address is the device's Internet (IP) address. This address helps route information between networks, and is used by interconnecting devices called routers.

Network interfaces are also called *NICs (Network Interface Cards)* or *adapters*, or sometimes just *cards*. You can see a picture of some network interfaces for Windows PCs at: http://www.3com.com/products/dsheets/400230.html .

Section 5.1 Technical Information Summary

Network interfaces

Network interfaces connect clients, servers, and peripherals to the network. Most network interfaces consist of a small circuit board that you insert into one of your computer's internal *slots*. Alternatively, modern computers sometimes include the network interface as part of their main circuit boards (motherboards).

Each network interface is associated with a unique address called its *media access control (MAC)* address. The MAC address helps route information within your local area network and is used by switches and bridges.

The MAC address is just one of several network addresses assigned to each networked client, server, or peripheral. Another network address is the device’s Internet, or TCP/IP, address. This address helps route information between networks. Every networked device maintains multiple, simultaneous network addresses which are used for different purposes.

Practical advice

Purchasing Considerations

When you purchase a network interface, you should consider the following guidelines:

- Make sure that the network interfaces on all computers are compatible with the physical and data link protocol you have chosen. For example, if you are running a 10BaseT Ethernet network, then all network interfaces must also use this protocol.
- Make sure that the network interface is compatible with the *slot* into which it will be inserted. Slots provide places on your computer's main circuit board (motherboard) where you can insert daughter circuit boards that add functionality to your computer (for example, network interfaces, modems, and so forth). Common slot types include PCI (Peripheral Component Interconnect), ISA (Industry Standard Architecture), EISA (Extended Industry Standard Architecture), among others. Each slot type specifies the speed, number of data bits used in the signal, and the number and position of wires on the motherboard used for communication inside the computer. PCI is the newest and fastest of the slots, although EISA and ISA slots are sufficient for most common network interface cards such as those for 10BaseT Ethernet. Most computers include slots of several different types. Before you order a network interface, check your computer to determine which slots are available, and then check your motherboard manual to ascertain the slot type. Order a card appropriate for your slot.
- Purchase network interfaces from a known manufacturer whose support you trust. Make sure the manufacturer provides a competitive warranty.
- Macintosh computers usually come with network interfaces as part of their main circuit boards. Some Windows PCs, however, still require that you purchase a network interface (for new PCs, your vendor may install the interface for you).

5.2 Hubs

On 10BaseT and 100BaseTX Ethernet networks larger than two computers, each computer or printer (or other networked device) is connected to a *hub*. The hub is a small box that gathers the signals from each individual device, optionally amplifies each signal, and then sends the signal out to all other connected devices. Amplification helps to ensure that devices on the network receive reliable information. You can think of an Ethernet hub like the hub of a wheel, at the center of the spokes that connect each individual computer or printer. Hubs are also called *concentrators* or *repeaters*. Hubs come in various sizes, the most common being 12-port or 24-port (meaning they can connect to 12 or 24 computers/printers/hubs).

You can see a picture of a hub at:
http://www.farallon.com/products/fast/hubs/940dualspeed.html

A simple 10BaseT or 100BaseTX Ethernet network may consist of a few dozen individual computers, printers, or servers connected to a single hub. In a more complex network, many hubs can be interconnected. You can see a diagram of a network containing many hubs at http://www.networking.ibm.com/mse/mse0c01.html (see Figures 1 and 2). The number of connections is limited only by the capacity of your network and servers to accommodate many simultaneous users, and the distance and repeater limitations imposed by the 10BaseT or 100BaseTX protocols (see Chapter 4 for a list of limitations).

All of the clients, servers, and peripherals connected to a hub (or to a set of interconnected hubs) share the *bandwidth* (data delivery capacity) of your network. Technically, they form a single *collision domain*—an area of an Ethernet network in which data sent to or from a device may potentially collide with the data from other devices. (In Chapter 4, we discussed Ethernet collisions and suggested that a 10% collision rate as reported by your *network operating system* is normal.) As you add more clients, servers, and peripherals to an Ethernet network, the number of collisions increases and the performance of your network degrades. You can improve performance by isolating network traffic into many smaller collision domains. Unfortunately, hubs cannot divide a network in this fashion; they simply repeat every signal all to all connected devices. Instead, to divide networks into multiple collision domains you can deploy switches, bridges, or routers. Each switch port, bridge port, or router port forms a new collision domain. That is, devices connected to a single port share the network bandwidth, but they are protected from the interfering signals of devices on other ports. We talk about switches, bridges, and routers in the remainder of this chapter.

The optimal number of computers, printers, or servers in a collision domain varies greatly from one network to another. You should think about dividing your network into smaller collision domains when your network operating system reports that collisions routinely exceed 10%, utilization exceeds 80% (Cisco's article suggests that the network utilization should not exceed 50%—http://www.cisco.com/warp/public/779/smbiz/netguide/answers/), or when your clients, servers or peripherals have difficulty connecting to or retrieving information from the network.

Like network interfaces, hubs must be compatible with the physical and data link protocol that you use. For example, if you are running a 10BaseT Ethernet network, then all hub connections to this network must also use this protocol.

Hubs work directly with the network signal itself (and not the data within the signal). That is to say, they work at the physical protocol level in the OSI model of network levels (level 1 manages signal strength and propagation). We therefore refer to hubs as *level 1 devices*.

While hubs are required for 10BaseT and 100BaseTX Ethernet networks larger than two devices, they are not used on LocalTalk networks.

Section 5.2 Technical Information Summary

Hubs

A hub connects individual devices on an Ethernet network so that they can communicate with one another. The hub operates by gathering the signals from individual network devices, optionally amplifying the signals, and then sending them onto all other connected devices. You should use a hub or a switch on your Ethernet network if the network includes more than two clients, servers, or peripherals. (We discuss switches later in this chapter.)

You can see a diagram of a network containing many hubs at http://www.networking.ibm.com/mse/mse0c01.html (see Figures 1 and 2).

While you can connect dozens of clients, peripherals, and servers via hubs, your network performance may degrade if too many devices try to communicate within one area of the network. You can improve performance by adding switches, bridges, or routers to the network. Each switch port, bridge port, or router port regulates traffic so that devices on the port are protected from the interfering signals of devices on other ports.

Most hubs operate by examining incoming or outgoing signals for information at OSI level 1, the physical level.

Practical advice

Purchasing Considerations

When you purchase a hub, you may wish to keep the following information in mind:

- Like network interfaces, your hubs must be compatible with your physical and data link level protocols. If you are running a 10BaseT Ethernet network, then you must purchase 10BaseT hubs. Some hubs, called multiprotocol hubs, can accommodate more than one physical and data link level protocol. For example, modern hubs may accommodate both 10BaseT and 100BaseTX protocols.

- If you purchase a multiprotocol hub, then make sure that it automatically senses which protocol is being used on each port. Autosensing hubs ensure that you can connect any part of the network to any hub port. (Older hubs required that you attach each segment of the network to a port compatible with its physical and data link level protocol. Keeping the segments and ports straight presents a management headache.)

- Make sure that your hub includes an *AUI port* (connector). (AUI is an abbreviation for *attachment unit interface*.) AUI ports are intended to connect with a kind of cabling called thick *coaxial* cable (like that used for cable TV). While this cable is no longer used frequently for Ethernet networks, AUI ports are versatile in the sense that they can be fitted with adapters to connect to many different kinds of cable (for example, thin coaxial cable or fiber).

- Make sure that your hub includes a *crossover port*. Unlike regular hub ports, which connect hubs to clients, servers, or peripherals, a crossover port connects one hub to another. In order to understand this distinction, you must consider how network devices use the Ethernet cable to send and receive information.

  All devices on 10BaseT or 100BaseTX Ethernet networks send their information over one particular pair of wires within the cable. This particular pair of wires is called the transmit pair. Similarly, all devices receive information from a different pair of wires, called the receive pair. The location of each pair of wires within the cable is specified by the wiring standard—for example, T568B—that was selected when your network was installed. All devices on your network conform to the same standard.

  When regular ports on hubs receive incoming information, they transfer it from the transmit pair of the sending device to the receive pair of the destination device. Crossover ports work in a different manner than regular ports. When crossover ports on hubs receive information, they simply pass it on without transferring it between transmit and receive pairs. By refraining from any change of pairs, crossover ports ensure that the next hub on the connection receives the original information intact.

- Some hubs can be stacked. *Stackable* hubs look like one, giant hub to the network. That is to say, the Ethernet restriction on the number of hubs that can be traversed in a single network (discussed in Chapter 4) does not apply to stacked hubs.

- Purchase hubs from a known manufacturer whose support you trust. Make sure the manufacturer provides a competitive warranty.

- Install your hubs in a room that is cool and free of dust, if possible. Additionally, plug your hubs into an uninterruptible power supply (UPS) to ensure that they receive clean power. (We discuss uninterruptible power supplies in Chapter 2.)

5.3 Switches

Like a hub, an Ethernet *switch* is a device that gathers the signals from devices that are connected to it, and then regenerates a new copy of each signal.

You can see a picture of a switch at:
http://www.asante.com/products/p_sw6.html .

You can see a diagram of a switched network at
http://www.networking.ibm.com/mse/mse0c01.html (Figures 3 and 4)

Switches, however, are more powerful than hubs and can substantially increase your network performance. In order to understand how they perform this magic, it is necessary to understand first how they work.

Most common switches operate by learning the MAC addresses of all connected clients, servers, and peripherals, and associating each address with one of its ports. When a switch receives an incoming signal, it creates a temporary circuit between the sender and receiver. The temporary circuit provides two important benefits.

- First, the circuit allows the sender and receiver momentarily to exchange information without intrusion from other devices on the network. That is, each pair of communicating devices utilizes the full bandwidth (data carrying capacity) of the network instead of sharing that bandwidth, as they do in unswitched Ethernet networks. To say this another way, each switch port defines a collision domain containing only a small number of devices and thereby helps provide maximum performance for Ethernet networks.
- Second, the circuit ensures that information travels directly between the communicating computers. This behavior differs markedly from unswitched Ethernet networks. In unswitched networks, data from a transmitting computer is sent by the nearest hub to all connected devices (not just to the recipient) and therefore congests parts of the network needlessly.

Like all network equipment, switches benefit your network only if they are deployed in the proper manner. If your network is congested and if traffic pools in certain areas, then you can improve network performance by replacing hubs with switches, or by connecting hubs to switches in a hierarchical manner. (You can see a diagram of a school network that uses a hierarchy of switches and hubs at http://www.3com.com/nsc/500612c.html . The switches are gray boxes and the hubs are black boxes labeled with numbers to indicate how many ports they have.) For the pools of heavy traffic, switches increase bandwidth while segregating the traffic from the rest of the network. However, if your network is not congested or if your traffic patterns do not create pools of congestion, then switches may actually cause your network performance to deteriorate.

This performance degradation occurs because switches examine the information inside each signal on your network (to determine the addresses of the sender and receiver) and therefore process network information more slowly than hubs.

You can read an excellent description of switches at http://www.lantronix.com/technology/tutorials/switching.html .

Because switches depend upon MAC addresses, we say in the parlance of the OSI model that they are *level 2 devices* (level 2 manages the structure and MAC addresses within network signals). You must purchase a switch that is compatible with your physical network and your data link protocols.

Recently, manufacturers have begun to offer switches that examine OSI level 3 (network routing) information such as that contained in the IP portions (rather than the data link portions) of a network signal. Later in this chapter, you will discover that routers also examine this information. Level 3 switches blur the distinction between switches and routers. Level 3 switches can replace routers within your network or between your network and the Internet (while level 2 switches can replace hubs, but not routers). You can read about switching and routing at http://www.networkmagazine.com/magazine/tutorial/internetworking/9705tut.htm .

Section 5.3 Technical Information Summary

Switches

Like a hub, a *switch* is a device that connects individual devices on an Ethernet network so that they can communicate with one another. But a switch also has an additional capability; it momentarily connects the sending and receiving devices so that they can use the entire bandwidth of the network without interference. If you use switches properly, they can improve the performance of your network by reducing network interference.

Switches have two benefits: (1) they provide each pair of communicating devices with a fast connection; and (2) they segregate the communication so that it does not enter other portions of the network. (Hubs, in contrast, broadcast all data on the network to every other device on the network.)

These benefits are particularly useful if your network is congested and traffic pools in particular areas. However, if your network is not congested or if your traffic patterns do not create pools of local traffic, then switches may cause your network performance to deteriorate. This performance degradation occurs because switches examine the information inside each signal on your network (to determine the addresses of the sender and receiver) and therefore process network information more slowly than hubs (which do not examine the signal contents).

Most switches operate by examining incoming or outgoing signals for information at OSI level 2, the data link level.

Practical advice

Purchase Considerations

When you purchase and install a switch, you should review and apply the following criteria:

- Your switches must be compatible with your physical and data link level protocols. If you are running a 10BaseT Ethernet network, then you must purchase a 10BaseT switch.
- Some switches can accommodate more than one physical or data link level protocol. For example, modern switches accommodate both 10BaseT and 100BaseTX protocols. It is wise to purchase a switch with at least one 100BaseTX port, since you can interconnect your switches via their high speed ports to improve network performance (even if the remainder of your network uses 10BaseT).
- If you purchase a switch that accommodates more than one protocol, then make sure that it automatically senses which protocol is being used on each port. Autosensing switches ensure that you can connect any part of the network to any switch port. (Older switches required that you attach each segment of the network to a port compatible with its physical and data link level protocol. Keeping the segments and ports straight presents a management headache.)
- Purchase switches from a known manufacturer whose support you trust. Make sure the manufacturer provides a competitive warranty.
- Install your switches in a room that is cool and free of dust, if possible. Additionally, plug your switches into an uninterruptible power supply (UPS) to ensure that they receive clean power. (We discuss uninterruptible power supplies in Chapter 2.)

5.4 Bridges

A *bridge* is a device that connects two or more local area networks, or two or more segments of the same network. For example, suppose that your network includes both 10BaseT Ethernet and LocalTalk connections. You can use a bridge to connect these two networks so that they can share information with each other.

In addition to connecting networks, bridges perform an additional, important function. They filter information so that network traffic intended for one portion of the network does not congest the rest of the network. (You may remember from the previous section that switches also perform these functions. Only to add to the confusion, you will find in the next section that routers perform similar filtering functions. Later in this section, we compare and contrast these three network traffic regulators in more detail to clarify their specific strengths and weaknesses at various network tasks.) Bridges may consist either of standalone hardware devices or of software running on a client or server.

When bridges were introduced in the 1980's, they typically joined two homogeneous networks (for example, two kinds of Ethernet networks). More recently it has become possible for bridges to connect networks with different physical and data link level protocols. For example, you can use a bridge to connect a LocalTalk network to an Ethernet network, or an Ethernet network to a TokenRing network.

Like switches, bridges learn the MAC addresses of all connected clients, servers, and peripherals, and associate each address with a bridge port (network connection). When a bridge (or switch) receives an incoming frame, it opens and reads its destination MAC address. If the port that will receive the frame is different from the port connected to the sender, then the bridge forwards the frame to the destination port. If the port that will receive the frame is the same as the port connected to the sender, the bridge drops the frame. (Since the bridge is by definition at the end of the network segment, the receiving computer presumably intercepted a copy of the frame on its way to the bridge.) If the bridge cannot determine which port is associated with a destination address, it passes the frame along to all ports.

Traditional bridges connect a single workgroup to another workgroup. More recently, however, manufacturers have produced multiport bridges. Multiport bridges allow network managers to connect more than two network segments to each other. Additionally, you can reconfigure or expand networks because simply by replacing one network interface card inside the multiport bridge with another (for example, adding a LocalTalk interface to a multiport Ethernet bridge).

Bridges generally inspect data link level information within a network signal—information like the Ethernet or LocalTalk (MAC) destination address. They do not attend to network routing or transport protocol information such as that carried within the TCP/IP, IPX/SPX, or AppleTalk portions of the signal. However, bridges can be fitted with custom filters that enable them to read this information—including network routing or transport source address, *packet* size, or type of protocol—and reject or forward information based on it. Custom filters enable network managers to isolate particular areas of the network and control which protocols enter or leave each area. For example, custom filters might allow requests from the Internet (outside the school district) not to enter certain areas of the network.

Bridges are relatively simple and efficient traffic regulators. However, in some networks they have been replaced by their more powerful cousins—hubs, switches, and routers. Each of these traffic regulators brings a unique set of strengths and weaknesses to its work:

- Hubs, switches, bridges, and routers can interconnect two different kinds of networks such as 10BaseT Ethernet and 100BaseTX.
- Hubs (unlike switches, bridges, and routers) do not filter traffic between the two networks.
- Switches have the unique capability to enable communicating devices momentarily to utilize the full bandwidth (data carrying capacity) of the network.

- However, switches (and hubs) cannot accommodate the variety of protocols and cabling types that bridges can.
- Routers are much more expensive and much more difficult to install and manage than hubs, switches, or bridges, but they can filter and route information much more precisely. (We discuss routers in more detail later in this chapter.)

When you purchase equipment, make sure you understand how each of these details affects your network. Then work with your technical staff or network integrator to choose the best equipment for each situation.

Because bridges (like switches) generally depend upon MAC addresses, we say in the parlance of the OSI model that bridges are level 2 devices. You must purchase a bridge that is compatible with your physical network and your data link protocols.

Section 5.4 Technical Information Summary

Bridges

A bridge connects two or more networks, or segments of the same network. These networks may use different physical and data link protocols. For example, you can install a bridge to connect a small lab of Macintosh computers using LocalTalk to the school's main Ethernet network.

Bridges filter network traffic. They examine each set of data, transmitting only appropriate data to each connected segment. (Hubs, by contrast, broadcast all information to each connected computer, whether or not that computer is the intended recipient.) In this manner, bridges help reduce overall network traffic.

Bridges are relatively simple and efficient traffic regulators. However, in most networks they have been replaced by their less expensive or more powerful cousins—hubs, switches, and routers.

Most bridges operate by examining incoming or outgoing signals for information at OSI level 2, the data link level.

Practical advice

Purchase Considerations

When you consider purchase of a bridge, you should follow these guidelines:

- Before you decide on your purchase, take a moment to clarify what you wish to achieve (connecting a Macintosh LocalTalk lab to Ethernet? connecting two Ethernet segments?). Then work with your technical staff, or with manufacturers and consultants, to determine your options. You can often use a hub, switch, or router in the same places that you can use a bridge. Each device brings its unique set of strengths and weaknesses to the job.
- Make sure that the bridge is compatible with your physical and data link protocols.
- Purchase bridges from a known manufacturer whose support you trust. Make sure the manufacturer provides a competitive warranty.
- Install your bridges in a room that is cool and free of dust, if possible. Additionally, plug your bridges into an uninterruptible power supply (UPS) to ensure that they receive clean power. (We discuss uninterruptible power supplies in Chapter 2.)

5.5 Routers, firewalls, and proxy servers

5.5.1 Routers

Like bridges, *routers* are devices whose primary purpose is to connect two or more networks and to filter network signals so that only desired information travels between them. For example, routers are often used to regulate the flow of information between school networks and the Internet. However, routers can inspect a good deal more information than bridges, and they therefore can regulate network traffic more precisely. They also have another important capability: they are aware of many possible paths across the network and can choose the best one for each data packet to travel.

Routers operate primarily by examining incoming data for its network routing and transport information—for example, information carried within the TCP/IP, IPX/SPX, or AppleTalk portions of the network signal. This information includes the source and destination network routing addresses. (Remember that every client, server, and peripheral on the network maintains multiple addresses, including both a data link and network routing addresses. The two addresses are used for different purposes. Among other things, the network routing address provides information on which routers base traffic management decisions.) However, most routers also include the same functionality as bridges. That is, they can inspect the data link level portions of the network signals for such information as the Ethernet or LocalTalk destination address.

Based on complex, internal tables of network information that it compiles, a router then determines whether or not it knows how to forward the data packet towards its destination. If the router has been configured with sufficient information to know which of its ports is en route to the destination, it transmits the packet. If the router has not been so configured, it typically drops the packet. Dropping unknown packets provides an important service to your network by eliminating restricted, wayward, or damaged information from your network. Bridges lack this capability (they forward unknown packets to all ports), and the misinformation they forward often creates extra network traffic.

Routers can be programmed to prevent information from being sent to or received from certain networks or computers based on all or part of their network routing addresses. If you have sensitive student records on a server, for example, you can use a router to filter packets headed for the server so that only authorized personnel—for example, personnel whose network addresses match a specified list—can connect to it.

Routers also determine some possible routes to the destination network and then choose the one that promises to be the fastest. As network traffic patterns change during a day, routers can adjust their route recommendations. (Very large routers route your information across the Internet in this manner.) There is a good deal of jargon associated with the art and science by which routers select paths, and a whole host of protocols that define their methods (for example, OSPF [Open Shortest Path First], RIP [Routing Information Protocol], or RTMP [Routing Table Maintenance Protocol]). Defining these schemes—and the costs and benefits of each—could occupy another volume the same size as this one; we suggest that you review other network references for routing protocol information.

Routers must learn formidable amounts of information about your network in order to inspect network routing and data link level portions of network packets, and to route information along the best path to its destination. Unfortunately, routers do not learn this information without human intervention. Installing routers is a complex task that involves configuring each network interface that connects the router to your network. First, you must enable support for the desired protocols on each network interface. Then, for each interface-protocol combination, you must either define routing tables or configure support for an automatic routing table update protocol. If you have defined enterprise-wide policies for security (rules that define the kinds of information that must be restricted), you must also define filters that implement these policies for each interface-protocol combination. Needless to say, you should make sure that qualified personnel (either your network integrator or your staff) install and manage them.

Since routers play a key role in connecting networks, they can cause significant problems if they malfunction. As part of your network plan, you should consider how you might deal with the failure of key routers on your network. Many sites include redundant connections—additional routers and network cable connections—configured to take over if one router or connection fails.

You can see a picture of a router at:
http://www.cisco.com/warp/public/cc/cisco/mkt/core/7100/index.shtml

Because routers depend upon network routing addresses, we say in the parlance of the OSI model that they are *level 3 devices* (level 3 manages routing information between different networks).

Section 5.5.1 Technical Information Summary

Routers

Like bridges, routers connect two or more networks. However, routers are much more powerful than bridges. Routers can filter traffic so that only authorized personnel can enter restricted areas. They can permit or deny network communications with a particular Web site. They can recommend the best route for information to travel. As network traffic changes during the day, routers can redirect information to take less congested routes.

If your school is connected to the Internet, then you will most likely use a router to make that connection. Routers ensure that your local area network traffic remains local, while passing onto the Internet all your electronic mail, Web surfing connections, and other requests for Internet resources.

Routers are generally expensive to purchase and difficult to configure and maintain. Be sure that your staff have the resources necessary to manage them well.

Routers quickly become critical components of your network. If they fail, your network services will be significantly impaired. As part of your network plan, you should consider how you might deal with the failure of key routers on your network. Many sites include redundant connections—additional routers and network cable connections—configured to take over if one router or connection fails.

Most routers operate by examining incoming or outgoing signals for information at OSI level 3, the network addressing level.

Practical advice

Purchase Considerations

When you purchase and install a router, you may wish to keep the following points in mind:

- It is best to purchase all routers from a single manufacturer. Purchasing routers from a single manufacturer ensures that the software you use to configure and manage the routers via the network will be compatible across devices (it is very important to be able to monitor and manage routers across the network if you want to keep things running smoothly). Additionally, your staff will find it easier to learn about and operate devices that are relatively uniform because they come from a single source. Make sure that your router manufacturer offers a wide variety of routers, including models for local area networks, dial-up connections, and wide area networks so that you can continue to purchase from the same manufacturer as your network grows. Consult with other educators to see which router manufacturers they have used and liked.

- Before you purchase a router, you should draw a picture of your network, including the place where intend to put your router. Then label the segments on either side of the router with the kind of cable used as well as with the protocols that will travel across the router to/from each segment. Your router must accommodate the cable types on all adjacent segments. In addition, the router must be compatible with protocols that appear on both sides of the router.
- Work with your router manufacturer or network integrator to choose the router model(s) that you need. Be sure that you can describe to your manufacturer/integrator not only the protocols in use, but also the kind of information that will be exchanged on the attached network, the kinds of information that may be restricted, the number of users, and their patterns of usage. You must match your router's capabilities to your particular network needs.
- Routers are often expensive. Your router should be easily upgraded so that you need not replace the entire device as your network incorporates additional kinds of cable or protocols. Ask your manufacturer about the particular expansion modules they offer, and what is involved in purchasing, installing, and maintaining them.
- Some managers plan to deliver multimedia applications over the Internet. These applications require a fast, steady stream of data to function properly. To deliver this increased performance, Internet standards organizations have defined options that allow routers and other network devices to reserve the bandwidth they need on the Internet. Such equipment is assures *quality of service,* or *QoS,* for specified purposes. Not all routers are capable of providing QoS services. If you are planning for multimedia delivery over the Internet, you may wish to make sure that your router does so.
- Price should not be the determining factor in purchasing a router. Routers, like servers, are key components of your network. It is far better to purchase durable equipment from premium manufacturers than to suffer equipment breakdowns or malfunctions.

5.5.2 Firewalls and proxy servers

A firewall is a device that prevents unauthorized electronic access to your network. The term *firewall* is generic, and includes many different kinds of protective hardware and software devices. Routers, discussed in the previous section, comprise one kind of firewall. A different kind of firewall might be created by installing software on a network server that is dedicated to the task of monitoring network activity. Yet another firewall consists of a standalone box (that is, a computer with no keyboard or monitor) which watches all the traffic on your network. All firewalls have one thing in common: they guard your network, examining information inside every network packet. Based on a list of restrictions which you provide, the firewall allows or disables each packet from traveling any further.

All the different kinds of firewalls mentioned in the previous paragraph can provide good protection. In order to select the best firewall solution for your school or district, you should match its capabilities against the particular security restrictions you require. To define these restrictions, start by asking yourself the following questions. What are the different types of data on the network (student grades, attendance, payroll, student projects, Web pages, and so forth)? Who are the different groups of people on your network (teachers, administrators, students, community members)? For each set of data, which groups should have access? What kind of access should they have (ability to read but not write, ability to delete, weekend or 24-hour access, and so forth)? Do teachers and students need video? audio? chat? Web pages? How often, for how long, and at what times of the day?

Once you’ve defined these policies, you are ready to consider the technical differences among firewalls and select the one appropriate for you. Keep in mind that the tighter your security policy, the slower your firewall will operate (because of the quantity of information it must examine to enforce complex rules), and (most likely) the more expensive it will be.

It is also possible (and even likely) that your school or district will include more than one firewall. For example, you might use a router and a separate computer fitted with firewall software at the entrance to your network, configuring the second firewall specifically to guard student records and accounting information inside your network. Your technical staff and firewall vendor should be able to review your security policy and propose the best combination of firewall devices for your purposes. Finally, keep in mind that firewalls form only one portion of your network security. Physical security, passwords, filtering software, and intrusion detection (the latter two topics are discussed below) complete the picture.

Types of firewalls

Firewalls can be divided into three major categories: packet-screening firewalls, proxy servers, and stateful inspection proxies.

Packet-screening firewalls. Packet-screening firewalls operate by examining incoming or outgoing signals for information at OSI level 3, the network addressing level. For example, you can configure your firewall to examine incoming packets for their Internet (IP) source address (the place where the information originated); you can deny access to your network if the packet comes from a network(s) that you have identified as unauthorized. Alternatively, your firewall can examine information leaving your network for its Internet (IP) destination address (where the information is being sent); you can deny access if users on your network attempt to connect to unauthorized sites. Besides source and destination addresses, firewalls can filter information based on the type of protocol used, the port number to which it is addressed (each Internet service such as electronic mail, TELNET, or FTP has a unique number, called its port number, which it uses to identify itself on the network), or content type (for example, JavaScript, Java, and so forth). You can use firewalls to control the information that enters your local area network from the Internet, leaves your site for the Internet, or travels from one part of your site to another.

Packet-screening firewalls have traditionally been implemented as add-on services within routers. On the positive side, they are among the fastest firewalls (because they examine a more restricted set of information than other firewalls). On the negative side, they are limited to examining network address and related information; they cannot implement complex security rules (for example, allowing the use of Web browsers but restricting the use of video). Finally, packet-screening firewalls leave you vulnerable to malicious information in portions of the packet that they don't examine—the data area beyond the packet's network address information.

Proxy servers. Proxy servers (also called application-level gateways) operate by examining incoming or outgoing packets not only for their source or destination addresses but also for information carried within the data area (as opposed to the address area) of each network packet. The data area contains information written by the application program that created the packet—for example, your Web browser, FTP, or TELNET program. Because the proxy server knows how to examine this application-specific portion of the packet, you can permit or restrict the behavior of individual programs. For example, you can configure your proxy server to allow Web browsing but to deny requests from FTP programs such as Fetch or WS_FTP. Alternatively, you can configure your proxy to permit FTP requests, but only if they read (not write) information. As a third example, proxies can deny Web browsers access to unauthorized Web sites.

You must configure a separate (software) proxy servers for each application you wish to screen. For example, your proxy server (hardware) will include multiple proxy servers (multiple software programs) if you want to screen information based on the common Internet applications—Web browsers, FTP, TELNET, and electronic mail. You should note that not all application programs can be proxied. For example, your accounting system may have no available proxy. For application programs without proxies, you must protect the program through packet-screening firewalls or other network services (passwords, for example).

Besides filtering information, proxy servers perform several other useful tasks. First, proxies hide the Internet (IP) addresses used by your organization so that intruders with malicious intent cannot easily determine the addresses to attack. Second, proxies cache information. That is, if the proxy grants permission to retrieve an Internet resource such as a Web page, it keeps a local copy. The next time that someone on your network wants to browse the same page, the proxy server checks its local cache. If the page is there, the proxy server checks with the originating Web server to see if the page has been updated. If not, the proxy delivers its local copy of the Web page to the user. This sequence of events is much faster than retrieving the entire Web page from the original server. In classroom situations where several computers simultaneously browse the same Web pages, this performance improvement can be substantial.

The specialized server that runs the proxy software is made as secure as possible by stripping it of all but essential services. For example, regular network servers may offer login, file- and printer-sharing capabilities, but secure proxy servers (and, for that matter, secure firewall servers) allow none of these services; all unnecessary or risk-prone services are turned off. Additionally, *operating system* updates, which often contain security fixes, are applied religiously. Stripping the proxy server (or firewall server) of extraneous services and keeping its operating system updated makes it more difficult for unauthorized users to gain access.

Stateful inspection proxies. *Stateful inspection proxies* examine the data within network packets to ensure that they are a legitimate part of a sensible, ongoing conversation between computers rather than a random insertion of (possibly malicious) material. Stateful inspection proxies fall midway between packet-filters and proxy servers in terms of security, but they offer relative ease of use and high performance. Like proxy servers, stateful inspection servers hide your internal Internet (IP) addresses from would-be intruders.

Security beyond firewalls

Besides firewalls, two additional types of security are increasingly common: software filtering and intrusion detection.

Software filters. In addition to packet-screening firewalls and proxies, some school districts implement software filters that screen information based on lists of restricted sites, types of application (for example, newsgroups), or types of offensive content (such as drugs). Generally, these software applications are designed specifically for use in schools and libraries, and they include lists of sites that are recommended for restriction.

There is a good deal of controversy surrounding the use of software filters, including debate over whether such software violates civil liberties and community standards, whether they restrict more useful than inappropriate information, and whether they work at all. Like many other worthy topics, a reasonable discussion of filtering software is beyond the scope of this Primer. You can read one educator's opinion of filtering software at http://fromnowon.org/fnomar96.html . In addition, CyberLibNet includes a list of useful resources at http://www.ala.org/editions/cyberlib.net/2penso02.html .

Intrusion detection. Intrusion detection software detects suspicious network behavior based on rules that you define, and then takes automatic action to terminate the behavior and trace its source. Suspicious network behavior can arise in a number of ways—via the exploitation of manufacturer-reported security holes, especially in operating system software; Java applets that are downloaded and commandeer a user's system; or password cracking (electronically guessing passwords). You can configure the intrusion detection software to recognize these difficulties and correct them automatically.

Section 5.5.2 Technical Information Summary

Firewalls and proxy servers

A firewall is a device that prevents unauthorized electronic access to your entire network. The term *firewall* is generic, and includes many different kinds of protective hardware and software devices. Routers, discussed in the previous section, comprise one kind of firewall. Most firewalls operate by examining incoming or outgoing packets for information at OSI level 3, the network addressing level.

Firewalls can be divided into 3 general categories: packet-screening firewalls, proxy servers (or application-level gateways), and stateful inspection proxies.

Packet-screening firewalls examine incoming and outgoing packets for their network address information. You can use packet-screening firewalls to restrict access to specific Web sites, or to permit access to your network only from specific Internet sites.

Proxy servers (also called application-level gateways) operate by examining incoming or outgoing packets not only for their source or destination addresses but also for information carried within the data area (as opposed to the address area) of each network packet. The data area contains information written by the application program that created the packet—for example, your Web browser, FTP, or TELNET program. Because the proxy server knows how to examine this application-specific portion of the packet, you can permit or restrict the behavior of individual programs.

Stateful inspection proxies monitor network signals to ensure that they are part of a legitimate ongoing conversation (rather than malicious insertions).

Besides firewalls, other types of security software may also be useful. For example, intrusion detection software monitors your network for particular kinds of malicious activity (attempts to steal passwords, for example). Filtering software maintains lists of Web sites that are permitted or restricted for students, and enforces those restrictions.

Many schools combine one or more of these solutions to create their network security system. Each solution has strengths and weaknesses. In order to choose a solution, you should begin by defining your security policy (the resources you wish to share or restrict, and the personnel who will have access to each resource). Then work with your manufacturer to ensure that your security solution meets your needs.

Practical advice

Firewalls (packet-screening, proxy, and stateful inspection) provide logs of traffic which you should monitor frequently. Logs indicate the people and resources that are active on your network. Also, the firewall should contain two network interfaces—one connected to the outside world and the other connected to your private network; the firewall operates by controlling the flow of information between the two.

If you choose to implement a firewall, most managers recommend that you implement a combination of packet-screening and proxy or stateful inspection services. Packet-screening firewalls are the fastest performers but the least powerful in terms of the kinds of information they can filter; proxy and stateful inspection servers are more powerful but slower.

When you add a firewall to your network, you must situate the firewall equipment so that it is the single point of access to all those resources on your network that you consider private. To visualize such a configuration, you can examine a picture of typical firewall and network at http://www.as400.ibm.com/tstudio/firewall/overview.htm . When the firewall forms a single point of access, it can review all inbound network traffic to determine whether it should reach private data, and it can review all outbound traffic to determine whether it is bound for an acceptable destination.

Computers that contain public information, such as Web servers, are not usually protected by firewalls. While it is relatively straightforward to define which outsiders should access your private information and to configure your firewall appropriately, it is very hard to define which outsiders should be excluded from your public information. In order to protect Web servers, you generally approach the problem from an entirely different point of view. You try to ensure that no one can add information to your server unless they have specific privilege to do so, and that malicious users cannot disrupt its activities. To enforce these restrictions, you configure the Web server operating system, Web server software, and associated software. These topics are beyond the scope of this Primer, but you can find an excellent discussion at the World Wide Web Consortium, http://www.w3.org/Security/Faq/www-security-faq.html .

You should avoid mixing security products from different manufacturers. Incompatibilities among equipment can cause unnecessary work and security risks.

Security is a very complex topic, and you must understand the possible solutions in order to make a good selection of hardware and software. There are many possible types of equipment, network configurations, and manufacturers. Take your time and research the area thoroughly before you purchase. For additional information about firewalls, see *Kicking Firewall Tires,* http://www.networkmagazine.com/magazine/archive/1998/03/9803mkt1.htm .

5.6 Putting it all together

You might wish to look at these articles:

Local area network information for schools:

- *A Guide to Networking a K-12 School District*
  http://www.ncsa.uiuc.edu/edu/nie/overview/handbook/handbook.html
- *Smart Valley Technical Guidebook for Schools*
  http://www.svi.org/netday/info/guidebook.html
- *The Virginia Department of Education's K-12 Technology Handbook (especially the link to The Technical Side of Networking):*
  http://www.pen.k12.va.us/go/techbook/toc.shtml

Network primers from manufacturers:

- Tips for education networks from Asante:
  http://www.asante.com/solutions/s_edu_c1.html
- The *Educator's Guide* from Asante: http://www.asante.com/solutions/pdf/edu_guide.pdf
- *Networking: A Primer* from Bay Networks:
  http://www.baynetworks.com/products/Papers/wp-primer.html
- *Network Essentials for Small Businesses* from Cisco:
  http://www.cisco.com/warp/public/779/smbiz/netguide/index.html

Advanced technical information:

- *Network Computing* magazine (series on networking technologies):
  http://www.networkcomputing.com/netdesign/series.htm
- *Network Magazine* (white papers on network technologies)
  http://www.networkmagazine.com/
- *techguide.com* (white papers on communications and internet technologies)
  http://www.techguide.com/