

# Bridges, Repeaters, and Hubs

**SOME OF THE MAIN TOPICS IN THIS CHAPTER ARE**

---

**Simple Network Repeaters    2**

**Using Bridges to Connect Network Segments    4**

**Using Hubs to Centralize LAN Wiring and Minimize Cable  
Problems    10**

**Hub Ports    17**

**Troubleshooting Hub Problems    18**

Chapter 2, “Overview of Network Topologies,” covered the kinds of network topologies used to join the computers that make up a network. The simplest topology available is the *bus*, which consists of a single segment of cable to which all computers are attached. This topology was the first used for Ethernet networks and worked quite well when only a small number of computers were to participate in the network. When minicomputers came into the picture, and most especially when PCs made their appearance during the 1980s, it became obvious that a computer network would have to be able to handle a much larger traffic load and connect many more computers than the simple bus topology would allow. Taking Ethernet as an example, there comes a point at which the number of frame collisions becomes excessive and network throughput becomes unacceptable.

This chapter discusses some of the legacy devices that were used to create larger local area networks. Some of these devices, such as repeaters and bridges, simply allow you to extend the length of the network and add additional computers. Others, such as a bridge, can help to limit the broadcast domain, in addition to allowing for a larger LAN. Finally, hubs will be looked at as a solution that allows for all the preceding features with the added benefit of centralizing the wiring that connects the network. In the next chapter, we discuss switches, which take the concept of the hub and add to it by creating a broadcast domain consisting of only two nodes on the network—the switch and a single computer.

- ◀◀ Because Ethernet is based on multiple computers all using the same network media for transmissions, each computer must contend for access to the media, as explained in Chapter 14, “Ethernet: The Universal Standard.” Because it is possible for more than one computer to begin transmitting at the same time, the signals sometimes interfere with each other in what is termed a *collision*. The term *collision domain* is used to specify a portion of the network that is shared by several computers, all of which can possibly transmit packets that could cause a collision.

## Simple Network Repeaters

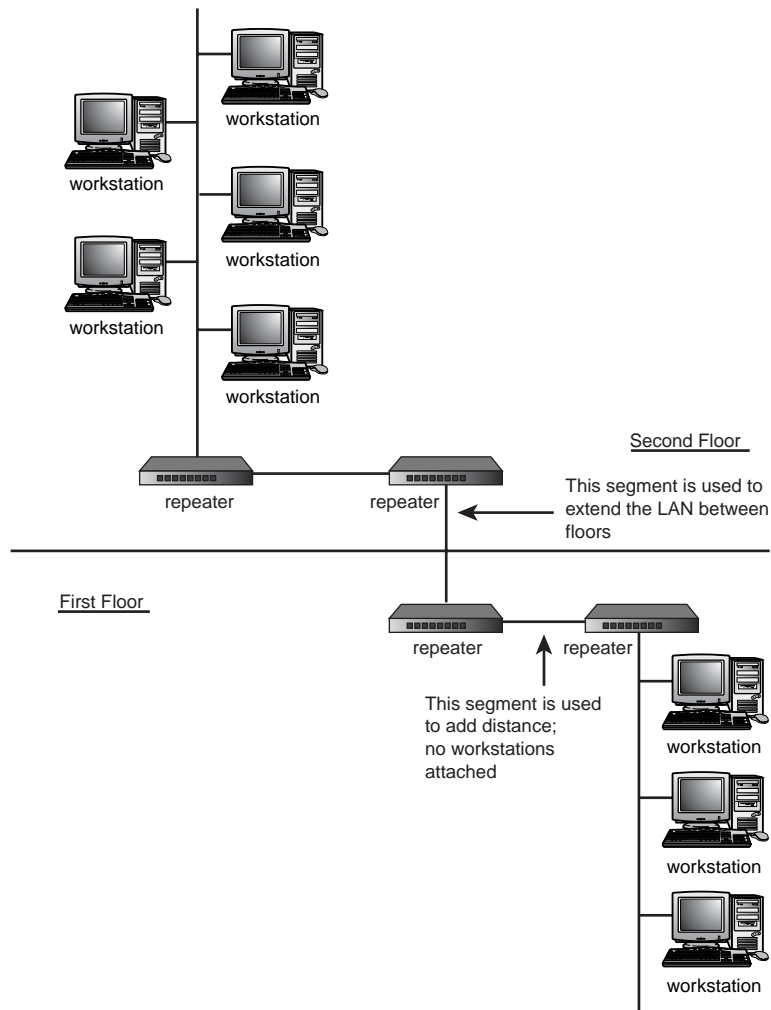
A *simple repeater* joins two physical segments, and it functions by amplifying the signal it receives on one port before transmitting the signal out the other port to the next network segment. Because a repeater only amplifies the electrical signal, and does not perform any actions on the information represented by the signal, it is totally ignorant of any addressing or routing information that might be encapsulated in the frame. This results in the repeater amplifying and sending out not only valid network frames that are transmitted on the network, but also fragmented frames resulting from collisions and background noise.

In a small network, this kind of device can be used advantageously to increase the size of a network based on a bus topology. For example, a thinnet network (10BASE-2) is limited to cable segment lengths of only 185 meters, or about 607 feet. The specification 10BASE-2 allows you to join as many as five cable segments, using four repeaters, which gives an overall length to the network of about 3,000 feet, as shown in Figure 1. Remember also that in this kind of network you can attach workstations to only three of the five segments. The other two segments can be used to extend the length of the LAN. In this drawing one segment is used to extend the LAN between floors in the building, and another is used on the first floor to add extra distance to get to the location of the users.

Advantages to using simple repeaters include price and simplicity. Repeaters are relatively inexpensive devices when compared to bridges or switches, and they require no management console or intervention by the administrator for configuration. A repeater can be used as a quick remedy when your small LAN segment needs to be extended beyond its normal range and it's too soon to explore further upgrade options.

In Figure 1 you can see that each repeater joins only two cable segments. Passive hubs, which are discussed shortly, work in much the same manner but allow you to connect multiple cable segments to a single repeater. The simple repeater, while allowing you to extend the length of the bus topology,

repeats *all* network traffic; thus it does nothing to limit the broadcast domain. You can extend the length of the network and you can add more computers or other devices, such as networked printers, but the more devices you add to the network, the more traffic that is generated on all segments that make up the network. The more traffic on a single shared media Ethernet network, the greater the chance of collisions. Therefore, after a certain point, network throughput will suffer because of excessive collisions.



**Figure 1** Repeaters can extend the network beyond the length allowed for one cable segment.

### Note

What is a hub? In the examples we've just looked at, you probably can determine that it is a device that centralizes the cabling of your network. Don't worry; we'll get to hubs later in this chapter and explain the different types and how they work.

A repeater that contains more than two ports, called a *multiport repeater*, also can be used to connect many LAN cable segments. Most of these devices provide the capability to block a port that is causing problems so that it does not affect the other network segments. Using a multiport repeater, you can attach a single computer to each port or create LAN cable segments that contain multiple computers connected to the cable that connects to the repeater port. The problem with the latter method is easy to see. Each segment that contains more than one computer is essentially an independent bus. If the cable is broken or malfunctions anywhere along the line, more than one computer will be unable to communicate on the network.

Although multiport repeaters served their purpose well, for the most part they have been replaced by more modern devices, such as bridges, intelligent hubs, and switches.

## Using Bridges to Connect Network Segments

Simple repeaters operate at the *Physical layer* of the OSI layered networking reference model. They just connect the cables and amplify the electrical signal before retransmitting it. Another device, called a bridge, also can be useful when joining cable segments. A *bridge* operates at the Data Link layer and does a little more than amplify the incoming signal and send it out to the remaining ports.

A bridge is basically a repeater with a little intelligence. At the Data Link layer the *Media Access Control address* (often referred to as the MAC address or the hardware address) is used to sort out which computers are connected to each cable segment. When a bridge is used to connect cable segments, it actually examines header information contained in the frame and finds the source and destination MAC addresses of the frame. The bridge keeps a table in memory storing these hardware addresses and, over time, learns the location of each device on the network. Because a bridge becomes aware of which devices are attached to each of its ports, it can make decisions when it comes to transmitting an incoming frame on the other ports so that the frame is sent out only over the port that will get it to its destination.

When a bridge is first powered up, it has no knowledge of the layout of the network. There is no information in its internal tables telling it what devices are on the network, much less to which cable segment they are attached. This information is built up over time as the bridge learns the network topology. Consider the following scenario:

- Workstation A sends out a packet that it wants delivered to workstation C.
- The hub to which workstation A is attached repeats the frame on its other ports, one of which is connected to the cable that attaches the hub to the bridge. Because this is the first time the bridge has received a frame from workstation A, it stores this information as an entry in its internal table, recording the MAC address for workstation A and noting from which port it came.
- At this time, the bridge does not yet know to which segment workstation C is attached. So to deliver the frame, it is necessary for the bridge to repeat the frame on its other port. Workstation C recognizes that the frame is addressed to it, and communication between it and workstation A can then take place. However, when workstation C then sends a frame back to workstation A, the bridge does not repeat the frame on the other ports. The bridge knows that workstation A is on the same segment as workstation C. At this time the bridge creates an entry in its table for workstation C so that future communications on this segment from A to C are not repeated on the other port.
- Next, workstation B sends a packet destined for workstation F. Again, the bridge notes the port and MAC address for workstation B because it had not heard of it before, and makes an entry into its table. The frame is repeated onto the segment attached to the other port.

- When workstation F receives the frame, it responds and sends back a frame to workstation B. When the bridge sees the destination MAC address for workstation B in the frame, it recognizes the frame as one stored in its table and repeats the frame on the port that has workstation B attached.

As you can see from this example, *after a bridge learns the location of a computer*, it does not have to repeat the signal on every segment like a simple repeater does. After the network has been up and running for a short period, depending on which workstations are active, the bridge knows the layout of the network and transmits frames to another segment only when it receives a frame with an address that is not in its table.

### Note

Because this type of bridge forwards frames without having to interact with the source or destination nodes, the computers attached to the network are not aware that a bridge exists between them. For this reason, this type of bridge is usually referred to as a *transparent bridge*.

The benefits of a bridge over a simple repeater are obvious. The bridge cuts down on unnecessary broadcasts because it will not forward a frame after it learns the port to which a MAC address is attached, unless it is necessary to get the frame delivered to its destination. For this reason, a bridge limits the broadcast domain. By properly using bridges, you can create a network that distributes the network traffic so that no single segment becomes overwhelmed with traffic.

For example, you probably wouldn't want to put all file and print servers on one segment and all your user's computers on another segment. In a typical client/server network, the user's computers are interacting with the servers, and the network traffic must flow through the bridge most of the time.

Instead, it makes better sense to place clients on the same segment with the server or servers that they most often use. Thus, the network traffic generated will be localized on that segment. In Figure 2, you can see a simple LAN that uses a bridge to join network segments that span two floors in a building.


In this example, you can see that the hub on the second floor has both user workstations and file and print servers attached to it. As long as users need to communicate only among themselves or these servers, the traffic is localized to the hub and does not pass through the bridge. A similar situation exists on the first floor, but two hubs are used instead. In this scenario, the two workstations on the second hub are used for training purposes and are located in a different room. By using a separate hub for these training workstations, the administrator can easily add or remove workstations in the training room without having to run new cables through the ceiling every time an additional computer is needed. Instead, all that is required is that the cable be plugged into the existing hub.

## Segmenting a LAN

As Figure 2 depicts, bridges can be used to take a simple network and divide it into segments to help isolate broadcast traffic. A general rule of thumb referred to as the 80/20 rule states that you should segment a LAN with bridges so that 80% of the traffic will be destined for devices attached to the local segment, and 20% might need to pass through a bridge to reach its destination.

### Note

The 80/20 rule was designed before switches and other modern network hardware were created. The rule is described here because you might find a question about this in a test, or if you are unlucky, you may find yourself in a very old LAN environment! Today, switches pretty much make this rule unnecessary. It is not uncommon today for a client to access servers in the next room, or a thousand miles away at another location.



**Figure 2** A bridge can be used to connect departmental LANs.

By taking a single LAN and placing a bridge between nodes that do very little communication, you effectively increase the available bandwidth because the entire LAN does not have to hear broadcast traffic for every other node. If you have one set of users who make heavy use of a particular server, and another set of users who rarely use the server but send a lot of print jobs to a print server, then separating these two groups with a bridge cuts down on the overall network traffic. Yet, because the LAN is connected by a bridge, any workstation can communicate with any other computer or device on the network if it needs to.

- ◀ There are many ways you can use connectivity devices—such as hubs, switches, routers, and bridges—to segment a network and reduce network congestion. Chapter 2 is recommended reading if you're new to the physical layout of a network.

The 80/20 rule assumes that in a typical business environment most communications take place between users who work in a particular department of the business. For example, most of the network traffic generated in the accounting department is most likely targeted at servers or other workstations in the accounting department. Occasionally it will be necessary to communicate with computers in other departments—email is a good example of this—but for the most part workers in the accounting department are busy interacting with a file server that stores the data (and in many cases the application software) that is used for the accounts receivable/accounts payable and general ledger. Putting


these users on the same LAN segment as users in the engineering department would make little sense. The intense amount of traffic generated by CAD software and other applications used by the engineers could result in complaints from users about the “slow network” they are on.

### Note

The 80/20 rule applies only to legacy devices and topologies associated with early Ethernet implementations. Today it is more likely that high-performance servers are stored in a computer room and the users who access the data the server stores do so from a different subnet or even across a wide area network. Switches and routers now operate at speeds that make this possible. To understand how switches work, and the problems they solve, see Chapter 8, “Network Switches.” Chapter 10, “Routers,” discusses when routers need to be used and how they operate.

So far, only the connection of two LAN segments has been discussed. Although you can see three hubs in Figure 3, keep in mind that two of these hubs are on the same LAN segment (on the first floor), so all traffic is repeated to all other computers attached to both of these hubs. Hubs, unlike bridges, do not limit the broadcast domain.

The network is more complicated in most situations. You can use multiple bridges to break up a LAN into more than two broadcast domains. In Figure 3, users in segment A make heavy use of the file and print servers located on their segment. Occasionally, however, they need to access the file server that resides on the other side of a bridge in segment B. Users in segment C make the most demand on the print server on their segment, but sometimes need to access the file server located on segment B also.



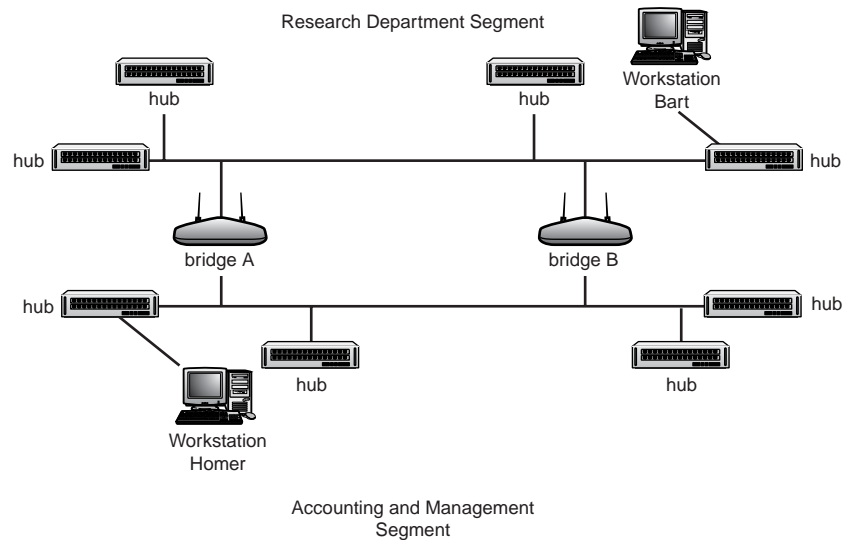
**Figure 3** You can use multiple bridges to break up a LAN into smaller broadcast domains.

Still, both groups of users need to exchange data with a remote office periodically. In Figure 3, the wide area network is accessible through a router connected to segment B. By placing the router on this segment, you keep down the level of traffic the router must process because it will not be seeing

every frame that contends for network bandwidth on segments A and C. Because the file server on segment B is used by both groups of users, but not on a frequent basis, the traffic to and from it is not likely to place a strain on the router's capabilities to examine and discard frames destined to the server.

## The Spanning Tree Algorithm

It also is possible to connect two network segments with more than a single bridge. This can be done to provide for fault tolerance. If one bridge ceases to function properly, another path exists between the segments so that users are unaffected, as you can see in Figure 4.



**Figure 4** You can use more than one bridge to connect the same segments.

However, using multiple bridges could result in a lot of confusion if bridges were not configured to cooperate. For example, consider the following situation:

- If the workstation named Homer sends out a frame that has the workstation Bart as its destination, both bridges A and B see the frame on the segment used by the accounting and management divisions. They both retransmit the frame onto the segment used by the research department.
- Both bridges then add an entry into their internal tables to remember that Homer is connected to the accounting and management segment.
- When bridge B sees the frame that bridge A transmitted onto the research segment, it does not understand that this frame is a duplicate of the one it also transmitted. Instead, because the bridge looks only at the MAC address, it sees it as if the frame were just broadcast by Homer on this segment.
- Bridge B then decides that Homer must have been moved and changes its internal table so that Homer now is attached to the research segment. Bridge A suffers the same fate, based on the frame that was originally retransmitted by bridge B.



Depending on timing issues, it also is possible that Bart will try to send a frame to Homer. If it is detected by either bridge while they still have Homer listed incorrectly in their internal tables, then Homer will never get the frame because the bridges will assume that workstation Homer is on the same cable segment as workstation Bart. If this is the case, the bridges assume that there is no need to send a copy of the frame over to the segment on which Homer is actually located.

But it could get worse. After both bridges have retransmitted the other's retransmitted frame back to the first segment, they will again see the other bridge's newly retransmitted packet and once again make changes to their internal tables and once again forward the duplicate frames back to the research department's LAN segment!

This infinite looping would make it impractical to use more than one bridge to connect two segments. In a large environment, chaos could result if an administrator accidentally connected two segments with a second router without realizing the problem.

---

**Note**

A bridge does not maintain an entry for a node in its internal table forever. In actuality, each bridge has a Time to Live (TTL) value that it uses. When an entry has been in the table longer than this value without being refreshed by the bridge again seeing it in a source address, the entry is removed.

---

To prevent this specific type of situation, the IEEE 802.1D specification defines the *spanning tree algorithm*. This specification allows bridges to interact with each other to initially create and then maintain a loop-free network. Although there can exist many routes that a frame can take through a multiple bridged network, the spanning tree algorithm establishes only one route a frame can take between any two points. When a new bridge is added or a bridge is removed from the network, the bridges recalculate the paths, and nodes are still able to communicate, provided that more than one bridge can be used to forward the frame.

Several values are used to provide the information needed by the algorithm to calculate the paths that will make up the tree. The administrator of the network can assign these values:

- **Bridge ID**—A unique identifier for each bridge.
- **Port ID**—A unique identifier for each port on each bridge.
- **Port priority**—A relative value designating a port's priority.
- **Port cost**—A value that designates a "cost" for a port. The higher the bandwidth of the link, the lower the cost should be for the port.

When a bridge is added to the network, it multicasts a message called a *Bridge Protocol Data Unit (BPDU)*, which contains information about it, including its ID and cost information. The bridges in the network evaluate these messages to calculate the correct paths that make up the tree of bridges.

The *root bridge* is selected based on the bridge ID value. The bridge that has the lowest bridge ID value becomes the root bridge. If two bridges have been assigned the same ID, the one that has the lowest hardware (MAC) address becomes the root bridge.

Every other bridge then must calculate the lowest cost path that connects it to the root bridge. The port that provides the lowest cost path to the root bridge is designated the root port for that bridge. If two ports have an equal cost path to the root bridge, the one with the lowest priority to the root bridge is designated the root port.

A designated bridge must be specified for each LAN. If only one bridge is connected to the LAN, it obviously becomes the designated bridge. If more than one bridge is connected, the bridge that has the lowest cost path to the root bridge becomes the designated bridge for that LAN.

For bridges that are not a designated bridge for a LAN, each port that is not a root port on the bridge is set to a blocked state so that it does not transmit any data. These blocked ports still listen for BPDU messages, however; so when the network topology changes, they can be used to begin the tree calculating process again. Bridges continue to exchange BPDU messages periodically. When bridges detect that a designated bridge has failed because they do not receive a BPDU message from it within the specified time limit, they begin to recalculate the tree topology.

## When to Bridge


As this chapter has discussed, bridges do more than just join segments. They are an inexpensive device that can be used to help grow a network until you reach the point where a switch or router would make more sense. By limiting the area in which a network frame can travel, you can use bridges for the following functions:

- **Reducing network congestion**—Grouping workstations, servers, and other devices that most frequently interact with each other reduces the total number of frames broadcast throughout the entire network.
- **Extending the length of the network**—Bridges allow you to connect multiple segments until you reach the limit for the network topology in use.
- **Minimal security purposes**—Place all workstations that exchange confidential data on a single network segment and use a bridge to connect them to the rest of the LAN. This way, only those workstations on that local segment are able to intercept the packets exchanged. A network sniffer on another segment, for example, would be unable to intercept these packets because they never get to that segment. Although this approach does provide some security, keep in mind that until a bridge learns the hardware addresses of the computers connected to the secure segment, it still passes them on to all other segments! A better solution is to use a firewall, which is discussed in Chapter 49, “Firewalls.”
- **Fault containment**—One malfunctioning device on the LAN is less likely to cause a problem beyond the segment to which it is connected.

However, in many cases a bridge makes sense only when used in a moderately sized network unless it is used in conjunction with other devices, such as switches or routers, and is used to limit traffic on a few segments. When the network outgrows the distances imposed by the topology limitations, you must use other means as an upgrade path.

## Using Hubs to Centralize LAN Wiring and Minimize Cable Problems

So far we have discussed using repeaters and bridges to connect LAN segments to create a larger network and to help relieve network traffic congestion problems. However, another problem with LANs that use a bus topology is that a break or another problem in the cable can cause all computers on a particular segment to lose their connection to the network. As discussed in Chapter 2, in a simple Ethernet bus topology a T-connector attaches to a BNC connector at each workstation. The coaxial cable that makes up the network connects to both sides of the T-connector. At each end of the bus (that is, the computers at the far ends of the network), a terminator is used to prevent signals from reflecting back and causing the signal to become garbled. Removing this terminator—as any network administrator knows, users will try anything—can cause the entire segment to fail or experience a dramatic slowdown in communications. In Figure 5 you can see the T-connector and a terminator. This figure also shows you the coaxial thinnet cable, with two twisted-pair cables for comparison.



**Figure 5** Twisted-pair cables simply plug into the card or hub port. Thinnet coaxial cables require a BNC T-connector and a terminator at each end of the network segment.


Using a multiport repeater to join segments that each have only one computer eliminates this major limitation of the bus topology. For example, if any cable becomes faulty or a terminator is removed from the end of the segment, only the computers on that segment are disconnected from the network. Other segments connected to a multiport repeater will continue to operate as usual.

Another nice feature of a multiport repeater is the fact that it allows you to concentrate LAN connections using a single device. The hubs that have been shown in diagrams in this chapter are essentially nothing more than multiport repeaters, with one major exception. Just like a multiport repeater, a hub concentrates LAN connections at a single point. Unlike a multiport repeater, however, hubs were developed to use twisted-pair wiring instead of the original 10BASE-2 and 10BASE-T coaxial cable. When a hub is used, each computer is attached to a separate port. This means that if one cable is defective or broken, the remaining computers on the LAN can still communicate. Because only one computer is attached to each hub port, there is no need for a terminator. Instead, a plug similar to a phone jack plugs into the hub port at one end and to the computer's network adapter card at the other end.

Hubs can be connected, however, so there still does remain a single point of failure if a cable that connects a hub to the network backbone becomes a problem. However, because a hub is usually placed in a secure location, such as the wiring closet, it is unlikely that a user will accidentally disconnect a cable or do anything else that will cause problems for other users on the hub.

A simple passive hub operates at the Physical layer of the OSI reference model. It does no addressing or framing; instead, it just receives data signals from one port and sends them out to all other ports. Because incoming signals are retransmitted on all the other ports, every computer in the broadcast domain can receive the data. This kind of hub is ideal for a very small LAN. However, today's more advanced hubs provide functions above and beyond the simple repeater.

In Figure 6, you can see that the hub creates a network based on a star topology from the physical point of view. That is, all the computers are connected by cables that are routed back to the hub, which acts as a wiring concentrator. However, from a logical view, the network still resembles the bus topology because each workstation in the LAN can intercept every packet that is broadcast on the network.



**Figure 6** The simple hub creates a LAN with a physical star topology.

Although the network is still a logical bus (every node can hear every other node), the hub has several features that make it superior to the single-cable bus topology:

- Only a single workstation suffers if a cable problem develops.
- A hub makes it easy to move computers because the entire network does not have to be disrupted when a computer is moved. When using a 10BASE-2 cable type, you must “break” the cable and install a T-connector to add a computer at a new point in the network. Using a hub, you just plug the new computer into any available port on the hub.
- The hub centralizes wiring. Instead of having a single coaxial cable snaking throughout the office or building, each workstation is connected by an easy-to-install twisted-pair cable. Just plug the cables in at each end and you’re ready to go.

## What Kind of Hub Do You Need?

In general, hubs can be classified by their function or by their construction. By physical construction, you can classify hubs as one of the following:

- Standalone
- Stackable
- Modular

Standalone hubs are small units that can be used to build a very small workgroup LAN. These are the most inexpensive hubs and can be purchased as a commodity item at most computer stores. These devices usually have a small number of ports (2–10). You can pick up a small hub of this sort for less than \$50—if you can still find one, because switches have now become the standard for small LANs. Figure 7 shows an example of a small eight-port workgroup hub. Notice also that you can see a BNC connector sticking out of the back of this hub. Although this hub uses twisted-pair cables to connect user workstations, it allows for a thinnet 10BASE-2 connection to other hubs or to the network backbone.


For a large LAN, stackable hubs provide for easy expansion as the network grows. This is because stackable hubs, as the name implies, can be linked so that they operate much like a single, larger hub. In addition, stackable hubs can be managed as a single unit when using SNMP- or RMON-based

management console applications. As you can see in Figure 8, creating additional ports is a simple matter of connecting another stackable hub unit to the existing units.




**Figure 7** A small workgroup hub can connect users to a larger network.

- ◀◀ The Simple Network Management Protocol (SNMP) and RMON are discussed in detail in Chapter 53, "Network Testing and Analysis Tools."



**Figure 8** Stackable hubs make it easy to expand the network one hub at a time.

In Figure 9 you can see an example of a rack-mounted stackable hub. Although it's a simple matter to locate a small workgroup hub on the desktop or on a shelf, using rack-mounted stackable hubs allows you to conserve space by placing a large number of hubs in a single cabinet. Notice that in addition to the 24 user ports available on this hub, there is a single port off to the right side used to link the hub to the next hub in the stack.



**Figure 9** Stackable hubs are usually available as rack-mounted equipment that can be stored in a cabinet in the wiring closet.

If you have a network that uses different technologies at the Physical and Data Link layers, such as a mixed Ethernet/Token-Ring network, then a modular hub can be used to connect computers to the network. In a sense, a modular hub is one step above the stackable hub.


The advantage is that instead of separate units that need to be linked together by short lengths of cable, the modular hub consists of a *cage* or *chassis* with a *backplane* or *motherboard*. Hubs are implemented as cards, with ports that can be inserted into the backplane. This type of hub can be quite sophisticated, incorporating Ethernet ports, Token-Ring ports (MAUs), management modules, and other devices (see Figure 10). Whereas stackable hubs each have a separate power supply and therefore require a separate electrical plug in the wiring closet, a modular chassis unit provides a common power supply for all the cards it holds.

Some modular hubs use a standard *backplane*, which consists of a single bus shared by all the cards inserted into the chassis. Some modular hubs house multiple buses, with each bus being specific to a particular network segment or LAN. A *segment bus* is similar to the multiple bus design, but the administrator can configure ports to be part of different LANs (virtual LANs, which are discussed in Chapter 9, "Virtual LANs"). Finally, a *multiplexed bus* allows for a multiple bus design using a single bus, with each separate LAN multiplexed on a separate channel. These last two types make it easier to move a workstation from one LAN to another. The administrator needs only to use the hub's management software to reconfigure the port to be on a different virtual LAN. This is much simpler than having to physically move a cable from one port to another or from one hub to another.

In spite of the advanced features of a modular hub, you might find that stackable hubs are better suited to your situation. Considering the fact that all cards in the modular hub share a common power supply, you can potentially lose more clients when a power supply malfunctions. Although an easy solution to this problem is to purchase a modular hub that has dual-redundant power supplies, you cannot plan for all possible disasters when using one piece of equipment to house so many things.

Another thing you should think about when considering modular hubs is whether LED indicators are present for each port. Like network adapter cards, most hubs have activity LEDs that can be used to diagnose problems with the port. Because the cards that make up a modular hub can be quite compact, it is often the case that there is no room to put in an LED for each port. Of course, the management application that is used to monitor and configure this kind of hub can make up for this, but

when the monitoring console is in one room and the actual physical hardware is a room away (or perhaps miles away), the actual LEDs can be a big help.



**Figure 10** A modular hub can house cards for different network types.

With stackable hubs, you can patch around a malfunctioning hub and keep your most important clients connected. Also, modular hubs are meant to be space-savers. Some of the features you expect on a hub, such as the LEDs that indicate the status or error condition, might not exist on individual cards that make up a modular hub.

By function, hubs can be classified into three major types:

- Passive
- Active
- Intelligent

### Passive Hubs

A passive hub is the simplest kind of hub. It acts like a multiport repeater and simply repeats incoming frames on all other ports. No signal processing is done when using this type of hub, so you need to be sure to stick with the cable lengths specified by the standard (such as 10BASE-T/100BASE-T) used for your network.

In a small business office that has no need to connect to a larger network, the standalone passive hub is an excellent choice. If you need to connect only 2–10 workstations, you can use this type of hub. For home networking, most computer stores sell these hubs in a kit format (“network in a box”) that includes the hub, two network cards, and the cables needed to connect the computers to the hub. For \$30–\$100, you can set up a network at home or in a small office in less than an hour.

### Active Hubs

Attenuation and other factors cause the signal to degrade as it passes through the wire. An active hub has electronic components that can amplify or clean up the signal it receives from a port before sending the data out on the remaining ports; this process is called *signal regeneration*. Some active hubs also have a capability called *store and forward*, in which individual packets can be examined and some simple corrections can be made to packets that have become corrupted during their journey.

Active hubs also can help compensate for ports that are connected to cables or workstations that are more likely to produce timing errors. Packet loss can be compensated for on these slower links because the active hub can retiming their delivery. This might at first seem like a bad idea because it effectively slows down communications when faster nodes communicate with the slower one. However, an advantage can be gained because the error-prone computer does not have to make repeated attempts to rebroadcast packets, which itself can cause network slowdowns in an already heavily loaded LAN.

---

#### Note

One thing to check for when looking at the specs for an active hub is the type of signal regeneration that is performed. For example, simply amplifying the signal can be a good thing for LANs that cover a long distance. Signal regeneration can be a bad thing when the active hub also amplifies any noise that occurs on the line. Simple things such as fluorescent lights can cause problems on copper wires, and this noise can be passed on by an active hub.

---

### Intelligent Hubs

The signal regeneration function of an active hub is included in an intelligent hub. Network management functions, which enable the administrator to gather information about the hub or each port on the hub, are also included. Statistical data about network traffic and error detection can make a big difference when you are troubleshooting. Some vendors provide their own proprietary management software. If you work in an environment in which the network is expected to keep growing, look to a vendor that supports industry standards such as the Simple Network Management Protocol (SNMP) or RMON, an extension of SNMP. As you add other devices later, they can be incorporated easily into any management console application that supports industry standard protocols.

---

#### Note

The Simple Network Management Protocol is a standard protocol that enables management software to manage and monitor devices throughout the network. SNMP provides a communication path between management stations and management agents that reside on the network device. RMON was created to enhance the capabilities of SNMP and allow for additional capabilities on the remote system. You can find more information about both of these standard monitoring protocols in Chapter 53.

---



Intelligent hubs often include not only monitoring features, but also management features that enable you to do things such as shut down an individual port connected to a computer that is causing problems on the network. For example, if a device on the network is malfunctioning and sending out a storm of broadcast packets, you can detect which port is causing the problem, and then selectively disconnect it from the network until the problem (cable, network card, or perhaps the hub port itself) is fixed.

## Hub Ports

The most typical hub has a row of sockets for RJ-45 connectors, which are currently the most widely used connectors for cables connecting workstations and network devices to a hub. When you're linking hubs to create a larger network, twisted-pair wiring might not be the cabling of choice if you need to span a distance longer than UTP can support. For this reason, you will find additional sockets on most hubs that can be used to connect the hub to a backbone. The RJ-45 connector ports are normally used to connect the cables that connect to a termination point at the user's location. This is typically a faceplate with another socket that can be used to connect the network card to the faceplate with a short patch cable. For more information about the EIA/TIA 568 structured wiring standard for using twisted-pair wiring and how the cables make their way from a hub to the user's environment, see Chapter 6, "Wiring the Network—Cables, Connectors, Concentrators, and Other Network Components."

For connections to other hubs, switches, or network devices that span a greater distance, you'll typically find that hubs intended for larger networks (that is, not a local home network) use ports that allow for connection to coaxial cables or fiber-optic cables.

### Note

Smaller hubs, such as the kind that were mentioned for a home or small office network, usually also use an RJ-45 port as an "uplink" connector to another hub.

## UTP, AUI, and BNC Ports

UTP ports are used with RJ-45 jacks to connect twisted-pair wiring to the hub. For connecting a hub to a network backbone, some hubs come with an AUI port, which can connect an Ethernet transceiver to 10BASE-5, thicknet cabling. More common is the BNC port, which can be used to connect the hub to a 10BASE-2 thinnet cable. For a high-speed connection you will find various connector ports that can be used to connect the hub to a fiber-optic cable. Indeed, there are hubs used for fiber-optic networks that use mirrors to split signals (instead of electronics, as is done with a standard hub).

For management purposes, you might find a port for a DB9 connector that is used to attach a management terminal to the hub. This type of asynchronous serial port is used either for the exchange of simple ASCII characters between the hub and a terminal when performing monitoring or management functions, or to connect the hub to a central unit that manages multiple hubs.

### Note

A hub doesn't have to have a serial port that can be used to attach a management console. In fact, most of today's advanced hubs can be configured with an IP address like any other network device, so management applications can use Telnet and other protocols and utilities to perform management and monitoring functions.

## Cross-Over Ports

Cross-over ports can be used to connect one hub to another to expand the LAN broadcast area. Sometimes referred to as an Uplink port, this port is wired differently than the others so that the send and receive lines match up when communicating between hubs. Some hubs have a switch next to the cross-over port so that you can toggle it to work as a regular port. This feature can be used to connect an additional workstation to the hub if you don't need to attach it to another hub. When it comes time to expand the network, toggle the switch, link the hubs using a twisted-pair cable with RJ-45 connectors, and use the next hub to connect additional clients.

## Autosensing Ports

If you are in the process of expanding an existing network and want to take advantage of faster network speeds, then be sure to check the specifications for any hub you purchase to be sure it supports both 10Mbps and 100Mbps. You also can simply buy separate hubs to accommodate these two speeds, but it's much simpler to have a single hub that can be used for both. As you begin to replace the network adapters in user workstations with faster cards, you won't have to transfer the cable at the hub end to another hub if you have purchased a hub that supports both speeds. Another feature that is pretty much standard for these kinds of hubs is that the ports (and the network cards) either can be manually set to a particular speed (using application software that comes with the hub) or can sense the network speed. An autosensing port just makes your job easier. Plug in the new network card, reboot, and users are back to work and can work faster.

## Troubleshooting Hub Problems

Sometimes, it can be difficult to troubleshoot problems with simple hubs that lack management software capabilities. You can be reduced to having to check the LEDs to ensure that the port is actively working or switching the particular computer to another port to determine what is causing the problem. For example, suppose that a workstation suddenly stops working on the network. What has happened recently that might have caused the problem? Has new software been configured on the workstation or on the network? Has someone damaged the cable (it's easy to damage twisted-pair cabling if it's lying around the user's area and his chair rolls over it a few times). It's probably best to start your troubleshooting efforts with the network configuration on the workstation, and then check the physical components, from the network adapter card to the cabling.

The following are some of the things you can easily check if you suspect problems with a hub:

- LEDs on the hub
- Whether a new connection or change has been made recently
- Whether a configuration change using management software has been done recently
- Whether the hub itself or just one port is totally dead

## Check Those LEDs

After ensuring that no changes have been made to the workstation's network configuration (which you can easily discern if you have a spreadsheet that is used to track this information), check the LEDs on the network adapter. If the link LED is lit, that simply means that you are connected to the hub and data can be exchanged. Next check the LEDs on the hub itself, if the hub provides this function. If the hub LED is not lit, try switching the user's cable to another port and trying again. Of course, you should do this using a port you know is in good working order, which usually means disconnecting another user for a short period.

If you find that the problem is with the hub port, you have several choices. If it's a cheap hub, connect the user to a free port, if one is available, or toss the hub and replace it. If it's an expensive hub,

be glad you signed that maintenance contract (if you don't have one and you're operating a large network, you'll probably be looking for another job sometime soon), and have it fixed or the hub replaced. In a large LAN environment, where downtime is intolerable and you can't wait for a repair person to arrive, having a spare hub stored away for just this kind of problem can be a lifesaver (or, better put, a job saver).

## Check for New Connections

If you are having problems connecting a new workstation to the hub, perform the LED check at both ends as just described and try another hub port. One of the most common problems encountered with new connections results from excessive cable length or sources of electrical interference somewhere along the path that the cable takes on its way to the faceplate at the user environment. Remember that the maximum distance a cable can be extended from a hub to the faceplate is about 90 meters, with another 10-meter cable running from the faceplate to the workstation. Of course, this depends on the technology you are using and will differ from 10Mbps Ethernet to Gigabit Ethernet. Check the specs! Also, don't forget that Category 5 twisted-pair cables are the preferred wiring for 10BASE-T and faster technologies. If you are still using Category 3 cables and are having problems, use a shorter cable length or, better yet, upgrade to Category 5.

- ◀◀ You can find out more about the restrictions imposed by Ethernet topologies in Chapter 14, "Ethernet: The Universal Standard," and those imposed on Token-Ring networks in the chapter "Token-Ring Networks," located on the [upgradingandrepairingpcs.com](http://upgradingandrepairingpcs.com) Web site.

Remember also that many environmental conditions can affect twisted-pair wiring. When first used in a Token-Ring environment, shielded twisted-pair cables were used because it was thought that the shielding would help prevent sources outside the cable from interfering with the signal that the cables carry. Ordinary Category 5 twisted-pair cables aren't shielded and depend on the number of twists-per-inch to help cancel out crosstalk and other factors that can cause problems. Trace the path that the cable takes and move it away from sources that generate electrical fields, such as fluorescent lights, electrical distribution points, and other similar things. You might solve a problem by simply moving the cable a few feet away from such a source of interference.

## Check the Hub or Port Configuration

Sometimes the problems you find in a hub are not due to failure of the hardware components. Instead, they can result from an improper configuration. Check the speed and the full- or half-duplex settings at both ends of the cable. In a large organization it might be that one group is responsible for installing and maintaining users' computers, and another group is responsible for hub management. In this case, make sure that the communication lines are open (between the people involved) and that the proper settings are established for both sides of the connection.

Finally, although most modern network adapter cards and hubs support "autosensing" the speed and other settings involved in the connection, they might not work well together. That is, although the hub and the network card might both say they support autosensing, they might use different methods. In such a situation, manually fix the method you want to use at each end.

## Check the Hub's Management Software

If you are using (expensive) hub application software to manage your hubs, don't overlook the information you can obtain from this application. Use the features of this application to assist you in troubleshooting efforts. If you don't use a central management console application, then connect a laptop or character-cell terminal to the serial port on the hub and use the commands available to try to determine what the cause of failure is. If the hub has a good software layer you can use, there will be tests that can be performed for each port and the hub in general. Make use of these tools.

Some vendors adhere to standards, such as SNMP, and others use proprietary software solutions. Whichever your vendor uses, be sure to check the documentation and try to use the tests or probes that are available to track down problem ports or hubs. It goes almost without saying that it is best to use hubs that conform to industry standards. This way you can use a central management console application to manage all your hubs, switches, routers, and other network devices from a single place, making troubleshooting efforts much easier in a network that is geographically dispersed. If you do run into a problem, you're more likely to find information that will be helpful for troubleshooting purposes if you are using equipment that adheres to the standards.

## General Hub Failure

Many problems you encounter with hubs involve only the failure of a single port. However, occasionally the entire hub could cease to function, disconnecting all its users from the LAN. All electronic equipment has a maximum temperature (and a minimum temperature) beyond which it will not operate properly. If you suspect that heat is the problem, try leaving the unit powered off for a half-hour or so to give it time to cool off, and then try powering it up again. If this solves the problem, try reorganizing the equipment installed in and around the device to make sure that temperature is not the problem.

There are all sorts of factors you can check to ensure that ventilation isn't being obstructed. Have you recently added new rack-mounted devices in the same cabinet? Have you checked to see that the ventilation fans inside a particular cabinet are functioning? Has someone inadvertently placed documentation on top of a ventilation area? Is there enough space between each device in a rack-mounted configuration for air to circulate freely? Finally, has the air-conditioning equipment in the computer room been checked recently? As a network administrator, you must consider these kinds of possibilities in troubleshooting seemingly obscure problems on your network.

Finally, if a hub outwardly appears to be functioning—LEDs are lit, for example, but in reality nothing is working—power-cycle the hub. When powering up, most modern hubs run through a self-check power-on sequence and report problems, either through management software or by a certain sequence of how the LEDs light up or flash. Examine your documentation. If the hub passes its power-on self tests (POSTs) and still does not respond, check the power supply. Although you can't do this easily for a small hub that has an internal power supply, you can check external sources such as an external UPS. Try troubleshooting or replacing any external power sources. If this does not work, call your vendor and get the darn thing replaced.