CHAPTER 10

Deployment Examples

In Chapter 9, "Deployment," a six-step deployment plan outlines some guidelines for migrating from 10Mbps shared Ethernet to Switched and Fast Ethernet. This chapter takes those basic principles one step further, giving practical examples. Each example described is a case study of a real network.

We will discuss the following specific examples:

- *Deploying Switched and Fast Ethernet in a new network*—This example describes how you can build a new network using state-of-the-art switching and Fast Ethernet products. Deployment examples cover client, server, workgroup, and backbone issues.
- *Adding Layer 3 switching capability to the backbone*—Using the previous example, this section shows how Layer 3 switching is practically configured, as well as how to mix in the use of virtual LANs.
- *Deploying Switched and Fast Ethernet in a LAN with an existing FDDI backbone*—This example describes how to convert or add on to an existing FDDI backbone. We discuss issues such as adding fast Ethernet workgroups to an FDDI backbone and expanding an FDDI backbone with 100BASE-FX.
- *Deploying a Gigabit Ethernet switched backbone*—This example gives a real-life deployment scenario for a Gigabit Ethernet backbone. We discuss deploying switched 100Mbps to the desktop along with aggregated Gigabit Ethernet backbone connections.

- *Deploying a 10Mbps switch as a switch of hubs*—In some smaller networks, 10Mbps switching may be more than enough bandwidth. This example shows one of these networks and how to leverage existing 10Mbps hubs.
- *Deploying Fast Ethernet in a branch office*—This example shows a branch office with a classic small LAN. The issues for this type of Ethernet installation include cost of equipment and routing to the WAN.

In general, this chapter applies what we've discussed in the previous chapters to real-world deployment examples. Although every network is different, you should find some similarities between your LAN and one of the instances discussed in this chapter. These examples should help crystallize the concepts and principles discussed up to this point.

Example 1: Deploying Switched and Fast Ethernet in a New Network

Rarely do you have the pleasure of building a new network from the ground up with few constraints on the type of network used. Even when this does happen, the budgetary considerations of a new network sometimes prove to be a bit stifling. 100BASE-T addresses this problem by combining good price performance with a scalable, high-speed networking technology. This example highlights the issues involved with building a brand new 100BASE-T network, including how to deal with the individual nodes, workgroup hubs, and backbone switches.

In order to work through this example, we must first make some assumptions. First, the new network will be installed in a building with four floors and a basement. Each floor is a square with 120-meter-long walls. Each floor will have roughly 200 offices per floor, and a CAD/CAM workstation cluster will reside on the top floor. Most of the servers and other critical network components will be placed in the basement of the building. High-performance Pentium II processor PCI-bus PCs will be purchased for the top three floors, but the bottom floor will be using existing Intel 80486 ISA-bus PCs. You only predict 150 users on the first floor, however. Each floor has two wiring closets, located directly opposite each other. Because plans call for future growth and the ability to upgrade, four pairs of Category 5 UTP cabling have been run to each office and 62.5/125μ two-strand fiber-optic cabling through the backbone. This building is shown in Figure 10.1.

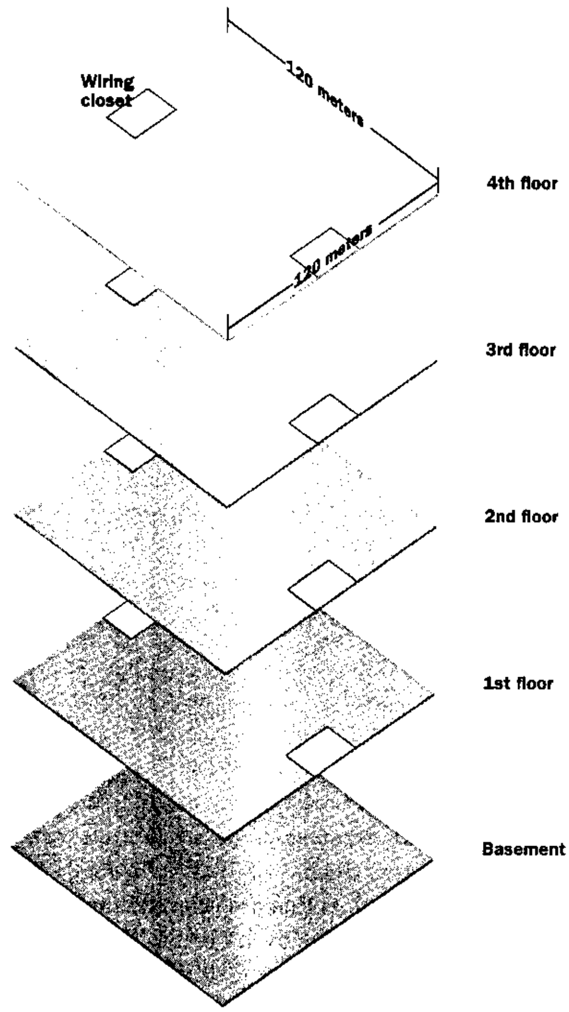

FIGURE 10.1 *A new building, ready for Switched and Fast Ethernet.*

The Backbone Solution: A 100BASE-FX Switched Backbone

After measuring the building for your formidable task, you discover that a 100BASE-T shared network won't be able to cover the entire area of the building. Not only is the building four stories tall, but each floor is 120 by 120 meters. The long distances involved also make a pure UTP installation messy because UTP must strictly follow the 100-meter rule. The only way to ensure less than 100-meter UTP drops to each office would be to situate the wiring closet for each floor in the center of the building. This is not practical because the wiring closets are already on the sides of the building.

With these facts in mind, a good solution for the backbone is switched 100BASE-FX. With Switched 100BASE-FX, you can connect to any other FX switch over fiber runs of up to 2km. Fiber is desired in the vertical runs because it offers more potential for future technologies (such as Gigabit Ethernet), it is less susceptible to noise, and it is recommended by the EIA/TIA standard for this type of application. Also, this network can be easily upgraded in the future because a 100BASE-FX backbone switch is highly scalable. Therefore, you decide to install a high-end 100BASE-FX backbone switch in the basement and develop a collapsed backbone. The switch should be modular in nature with several slots for 100BASE-FX modules. If your budget allows, you may decide to deploy a 100BASE-FX router (like the Cisco 7500) or a 100BASE-FX Layer 3 switch (like the Intel Express 550F) instead of a normal backbone switch. (We'll discuss this in more detail in the second example in this chapter.) For the purposes of this example, we'll choose the Cisco Catalyst 5000 with Switched 10/100 Base TX FE (Fast Ethernet) modules for the backbone switch. The 5000 provides 1.2Gbps of backplane capacity, so it should be able to support eight Fast Ethernet connections and is scalable to many more.

In addition, the Catalyst 5000 can be configured to support between 16,000 and 64,000 MAC addresses, which allows it to handle the large number of clients in this particular network. What does the Catalyst 5000 connect to in each floor's wiring closets? In the wiring closets, the most important features are the following: connectivity to multiple 100BASE-FX fiber ports and scalable 100BASE-TX switching capacity. These are important because they allow you to scale the performance of a floor as the bandwidth demands increase. In each wiring closet, place a 100BASE-TX switch such as the Bay 450T. The Bay 450T has 24 10/100BASE-TX switched ports, and two slots that are upgradable to 100BASE-FX. This will come in handy later when you deploy 100BASE-TX to the workgroups on each floor. When installing the fiber cabling, it is highly recommended to install multiple fiber runs to each floor because this will allow you to scale the bandwidth to any floor without reinstalling more fiber later. In this example, four fiber runs are made to each closet, for a total of eight to each floor. Figure 10.2 shows this.

Also, the connection to a WAN link is done through the basement chassis hub via a T1 (or another type of WAN connection) module. All access to the Internet and other regional offices will be made through this type of basement router.

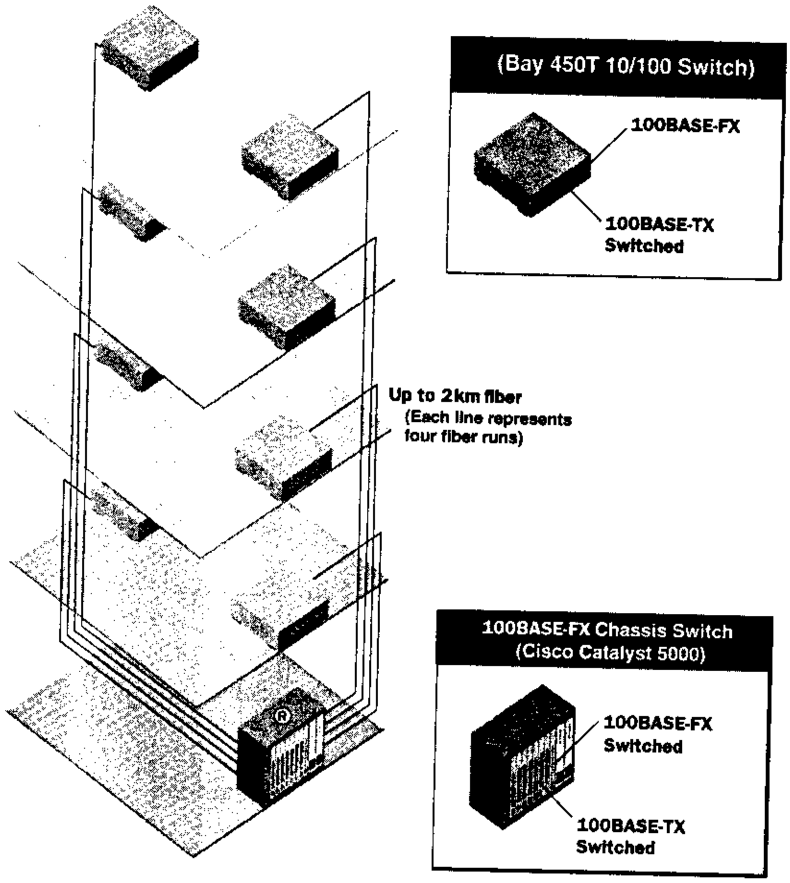

FIGURE 10.2 *Deployment of a new 100BASE-FX switched backbone. A chassis-based switch is installed in the basement and is connected to the wiring closets via 100BASE-FX fiber cabling.*

The Workgroup Solution: A Combination of Switched and Fast Ethernet

Now that your 100BASE-FX switched backbone is in place, you can address the deployment of 100BASE-T to the workgroups. The top three floors will all have new PCI-based PCs, so each of these can be outfitted with a 10/100 NIC. The first floor is using legacy ISA-based PCs, so they have to be installed with legacy 10Mbps ISA NICs. For this building, you have chosen 100BASE-TX for the 100Mbps option on your PCI NICs and 10BASE-T as the option for your 10Mbps-only NICs.

Top Floors: 10/100BASE-TX Stackable Hubs

Because many connections will be on each of the top three floors, you elect to deploy 10/100BASE-TX stackable hubs, such as the Intel 220T 10/100 Stackable Hub. Each of the stackable hubs is connected to one of the 100Mbps switched ports of the 100BASE-TX switch installed in the wiring closet. In turn, each one of the stackable hub ports is connected to a Category 5 UTP wiring segment that runs from the individual offices on the floor. By using stackable hubs that are capable of 10 or 100Mbps per port, you can automatically connect to both new fast Ethernet nodes and existing 10Mbps nodes.

It is also important to recognize a potential cabling problem here. If UTP cabling is only allowed to run parallel to one of the walls, some of the offices will be out of range of the 100m UTP limit. What is needed in this example is a little foresight in planning the UTP cable runs. In addition to the normal cable trays in Figure 10.3 (a), this building should be enhanced with diagonal cabling trays. With this design, you can reach all points on a floor with less than 100 meters of UTP cable. This is shown in Figure 10.3 (b).

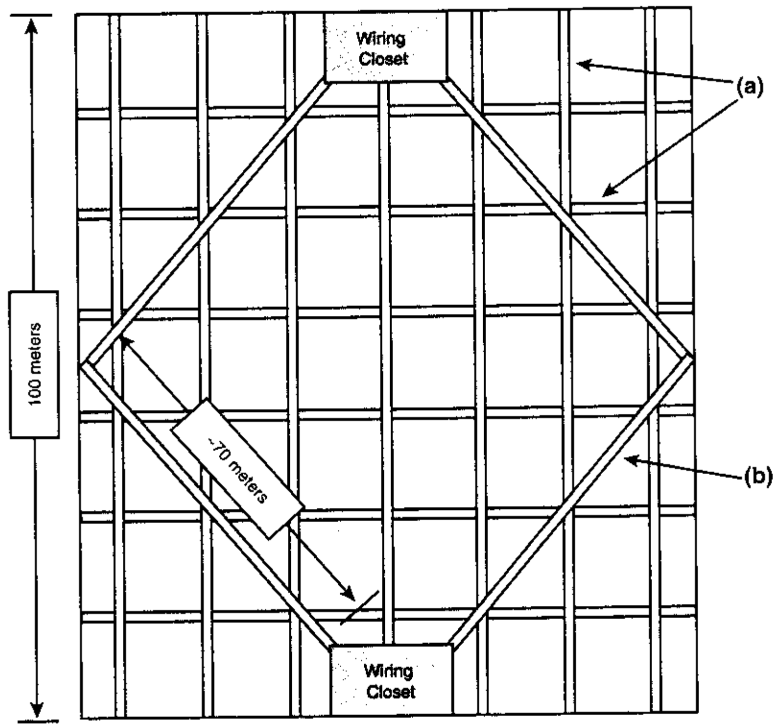

FIGURE 10.3 *(a) Normal cable runs in a building, which doesn't allow for 100 meters UTP, and (b) special diagonal cable runs that allow 100 meters UTP cabling.*

You might want to expand this network in the future (increasing the number of PCs per office, for example), so make sure you have room to grow. For this reason, you should buy 10/100BASE-TX stackable hubs that can be stacked at least six units high. Start by installing only two units per stack, which leaves room for the workgroups to expand beyond their current size. Each stackable unit has 24 10/100BASE-TX ports, so each stack of two units can service 48 users. Therefore, six stacks (three per wiring closet) of 48 ports each will service 200-plus users on each floor with ports to spare.

Each floor has two wiring closets per floor, so you can connect half of each floor to the stackable hubs in one closet and the other half to the other closet. Each workgroup of 48 users shares 100Mbps of bandwidth, and each workgroup is allowed a 100Mbps dedicated port into the 100BASE-TX switch in the wiring closet. Therefore, you should place three sets of stackable hubs in each wiring closet. There are also some local servers for each workgroup. Connect these directly to a Fast Ethernet port on the Bay 450T workgroup switch. Figure 10.4 shows the 100BASE-TX stackable hub deployment.

Luckily, you have spent some time reviewing the traffic in similar networks throughout your company. You know from network management software traces that different users and applications create different kinds of traffic patterns. By studying the trend analyses and prospective user profiles of the new network, you discover that in any given workgroup, about 50% of workgroup traffic stays local: that is, it never gets forwarded to the backbone. Therefore, you deduce that the 100BASE-FX backbone link to the wiring closet needs to support a worst-case 50Mbps of data from each 48 user workgroup in that wiring closet (average workgroups generate much less, but you want to plan for a worst case scenario). Because there are three workgroups (stacks of hubs) per closet on floors 2, 3, and 4, the backbone connection needs to support a maximum of 150Mbps. Two aggregated, full-duplex, switched backbone connections that are provided by the 100BASE-FX switch will support this rate. If more workgroups are added or the existing workgroups are segmented so that there is more than 150Mbps of demand for the 100BASE-FX backbone connection, an additional Switched 100BASE-FX connection may be needed for that particular wiring closet. This is easily accommodated by connecting one of the extra fiber runs (remember we installed four) from the basement to an FX port in the wiring closet. The cost of running additional fiber at the time of initial installation is far less than doing it later.

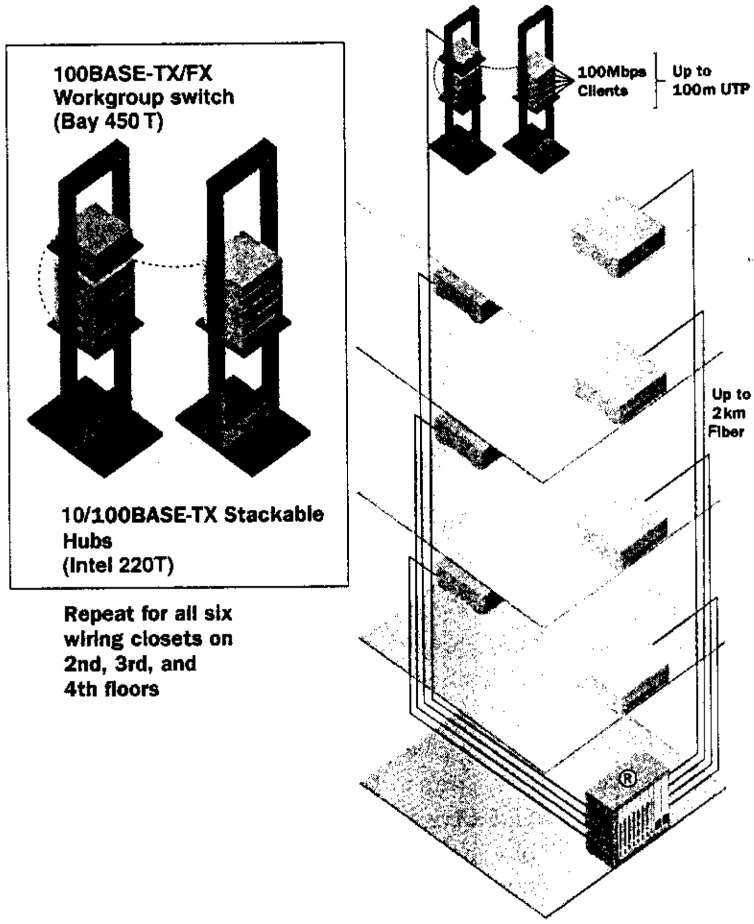

FIGURE 10.4 *Deployment of new 100BASE-TX workgroups. 100Mbps switches are connected to 10/100 stackable hubs in the wiring closet. This is ideal for clients with either 10 or 100Mbps NICs.*

First Floor: 10Mbps Switches

The workgroup situation on the first floor is very different. Because the bulk of the systems are ISA-based PCs, the majority of the connections will be 10BASE-T (ISA doesn't support 100-Mbps data rates very well). As we saw in Chapter 9, 10Mbps switches work well in workgroups with legacy 10BASE-T clients. The network installation, however, shouldn't preclude the possibility of adding newer PCI-based systems equipped with 10/100-Mbps NICs. The workgroup strategy for the first floor also has to account for local servers, which should already have 10/100-Mbps NICs installed.

10Mbps workgroup switches best handle this type of installation. Each workgroup switch has two or more 100Mbps switched ports and multiple 10Mbps ports. A good example of this type of switch is the Cisco Catalyst 2820. The Catalyst 2820 may be equipped with modules that have up to nine 100Mbps ports and 24 10Mbps ports. Every client is connected to a 10Mbps dedicated port on the 10/100 switch. Because each switch has 24 10Mbps ports, a total of eight switches are needed to connect the 150 clients on the first floor (four switches in each wiring closet). One of the 100Mbps uplink ports on the 2820 is connected to the previously deployed 100Mbps switch in the wiring closet. Figure 10.5 shows the first-floor installation.

There is one potential problem with this first-floor deployment scenario, however. Each workgroup switch collects data from 24 10Mbps clients. In this scenario, a workgroup switch could easily be forwarding 100Mbps of data at any given time. Because there are a total of eight workgroup switches, each providing 100Mbps of bandwidth, there is the potential for 800Mbps of traffic to be forwarded to the backbone. The 100BASE-FX backbone provides only one 100Mbps link, so it may become *oversubscribed*. Oversubscription occurs in switched environments when the sum of all data rates on the low-speed ports exceeds the data rate of the high speed port. Of course, not all traffic will be forwarded to the backbone, but if we assume half of the traffic is forwarded, as on the other floors, oversubscription could still occur. In order to prevent oversubscription, one of two things must be done. Either more switched fiber backbone connections must be run to the first floor, or the 100Mbps switch on the first floor must be equipped with a congestion control feature. As discussed in Chapter 4, "Layer 2 Ethernet Switching," a switch can use 802.3X full-duplex congestion control to tell its connections that it is in a state of oversubscription. The Bay 450T is a good example of a switch with this capability. In this example, congestion control will allow only 100Mbps of combined traffic to flow from the eight workgroups to the 100BASE-FX backbone connection. Because it is highly unlikely that all switches will be fully subscribed at 100Mbps at any given instant, congestion control provides a solid solution without the cost of additional switch backbone connections.

Local servers and clients with new systems can still connect at 100Mbps through one of the 100Mbps uplink ports on the Catalyst 2820 workgroup switch. Each new system deployed on the first floor should include a 10/100 NIC, so it can eventually be connected to a 100Mbps port. Another option for the new systems on the first floor is to connect them via 10/100 stackable hubs similar to the other floors.

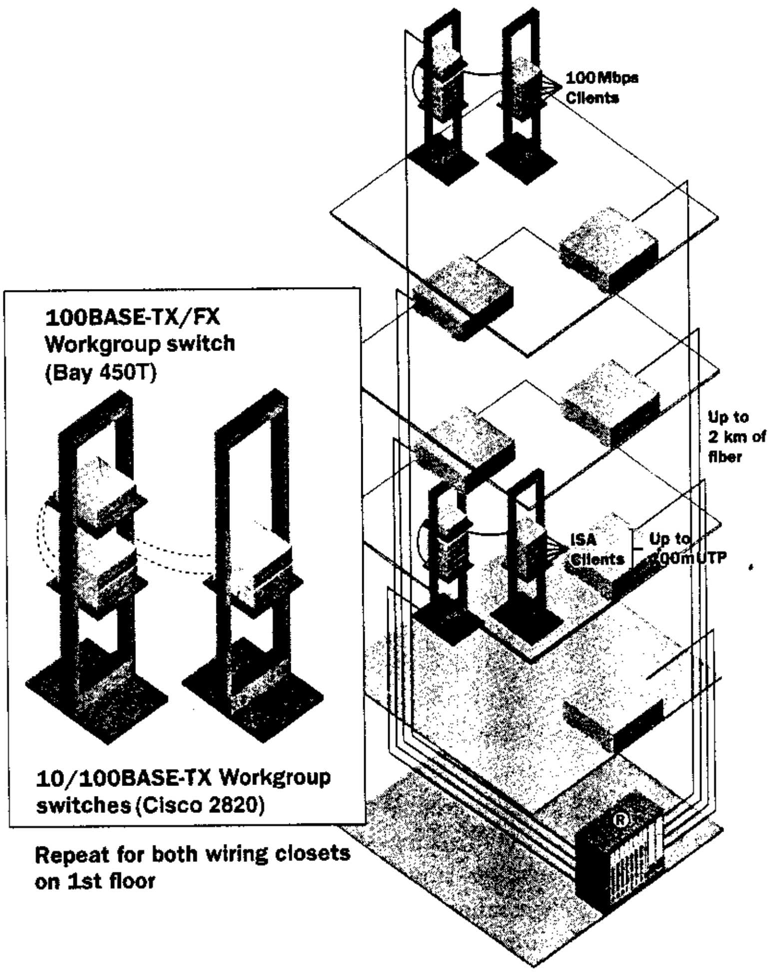

FIGURE 10.5 *Deployment of new 100BASE-TX workgroups with 10Mbps ISA clients. 10Mbps switches with 100Mbps uplinks connect these clients to the 100Mbps network.*

Cost Analysis

The cost of deploying Fast Ethernet in this new network example will be higher than the cost of deploying a similar 10BASE-T solution. The performance gained, however, is roughly ten times that of 10BASE-T. The cost of deployment for 100BASE-T has dropped and will continue to drop, quickly. As volumes and silicon integration go up, network vendors will be able to sell 100BASE-T at costs close to 10BASE-T. In fact, in the time since we published the first edition of this book, the price for 100BASE-T network equipment has literally fallen to within 10% of similar 10BASE-T equipment. For any new network, cost is easily determined by analyzing the cost of the network equipment in each case combined with the cost for installing and maintaining that equipment. Because installation and maintenance of a 100BASE-T network is basically identical to that of a 10BASE-T network, one should look closely at the initial purchase price. Below is a cost worksheet that compares the cost of deploying 10BASE-T and 100BASE-T. This worksheet can be used as a starting point to determine how much Switched and Fast Ethernet will cost. In Table 10.1, the equipment from this example is used. To analyze the cost of your new network, create a similar table and plug in the associated equipment costs.

TABLE 10.1 A SAMPLE COST WORKSHEET FOR JUSTIFYING 100BASE-T DEPLOYMENT

Network Equipment	100BASE-T Equipment	Equipment Example	100BASE-T Cost	10BASE-T Equipment	10BASE-T Cost
Backbone switch	100BASE-FX chassis-based backbone switch or router	Cisco 5000 24+ 100FX ports		10BASE-T backbone switch	
	8 100BASE-TX backbone switches with 100BASE-FX port	Bay Networks 450T 24 ports		8 10BASE-T backbone switch (nonmodular)	
Workgroup switch (2nd–4th floors)	6 100BASE-TX stackable hubs (6 units each)	Intel 220T 10/100 Stackable Hub 24 ports		6 10BASE-T stackable hubs (6 units each)	
Workgroup switch (1st floor)	8 10/100 switches	Cisco Catalyst 2820 24 ports		8 10BASE-T repeaters	
Workgroup NICs (2nd–4th floors)	600 10/100BASE-TX PCI NICs	Intel Corp. EtherExpress PRO/100+		600 10BASE-T PCI NICs	
Workgroup NICs (1st floor)	150 10BASE-T ISA NICs	Intel Corp. EtherExpress PRO/10+		150 10BASE-T ISA NICs	
Total					

Example 2: Adding Layer 3 Switching Capability to the Backbone

In this example, we'll look at adding Layer 3 switching to the network from example 1. Two factors are driving the need for Layer 3 in this case. First, the Catalyst 5000 routing module in the basement is becoming more and more saturated with routing tasks. This is because each floor has been segmented into one or more IP subnets. Second, a few existing NetWare servers are using IPX, and we want to isolate the IPX traffic only to those workgroups that require access. This will be accomplished by using protocol-based VLANs.

Adding Layer 3 switching capability is relatively straightforward in this case. First, the basement switch is upgraded with 100BASE-FX modules that support IP Layer 3 routing capability. Second, the wiring closet switches are upgraded to support Layer 3. These switches can be upgraded one by one, as the need to filter Layer 3 traffic becomes greater for that particular workgroup. In this example, shown in Figure 10.6, the Intel Express 550T has been used in the wiring closet. The Intel 550-T has been used in place of the Bay Networks 450T because of its support for IP- and IPX-based Layer 3 switching.

Routing capability in the workgroup will prevent broadcast traffic reaching the backbone from that workgroup. In this example, the Catalyst 5000 in the basement is freed up to handle the switching of data packets instead of determining where a certain packet should go.

This design also requires that the IPX traffic from the NetWare server be isolated from the clients that don't need to access the NetWare servers. We do this by configuring the switches to support protocol-based VLANs. As discussed in Chapter 4, this will be a manual process because most switches don't support a standards-based way of automatically supporting protocol-based VLANs. In the future, this may be accomplished with switches and NICs that support the emerging 802.1Q VLAN tagging standard.

The Catalyst 5000 and the Express 550F each have to be set up to separate IPX traffic into a different VLAN. Broadcast traffic generated by the NetWare servers will only be forwarded to ports on this VLAN. This means that each switch must be set up in advance with the ports that are connected to clients and servers that might send IPX traffic. In this example, no one on the first, second, or third floors needs access to the NetWare servers in the basement, so IPX broadcast traffic will never reach those floors. The top floor, will, however, forwrd IPX broadcast traffic to the predesignated switch ports. This is not the only way to use VLANs in this scenario, but it achieves the design goal of separating the IPX traffic from the bulk of the network. If more flexibility is needed, the IPX traffic could be limited to three or four select VLANs instead of just one. Figure 10.7 shows this.

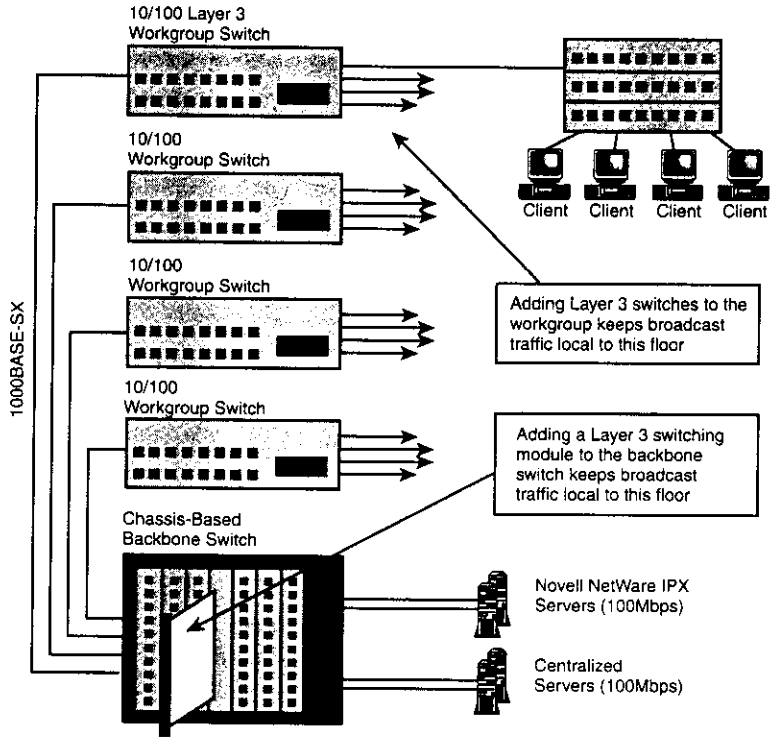

FIGURE 10.6 *Adding Layer 3 switching support to both the backbone switch and the workgroup switches. Broadcast traffic is contained in the workgroup and prevented from entering the backbone.*

Overall, VLAN segmentation can be beneficial, if not a little cumbersome. In this example, it is worth the effort because the number of clients that need access to the NetWare server is small, and they are isolated to one physical floor.

In this example, we have continued to use aggregated 100BASE-FX as the backbone technology of choice, however, the same model applies when Gigabit Ethernet is introduced. Instead of a 100BASE-FX Layer 3 switch module for the

Catalyst 5000, we could have selected a Gigabit Ethernet Layer 3 switch module. Gigabit Ethernet modules would also be needed in the workgroup switches (Intel 550Ts in this example). We'll discuss this more in example 4 of this chapter.

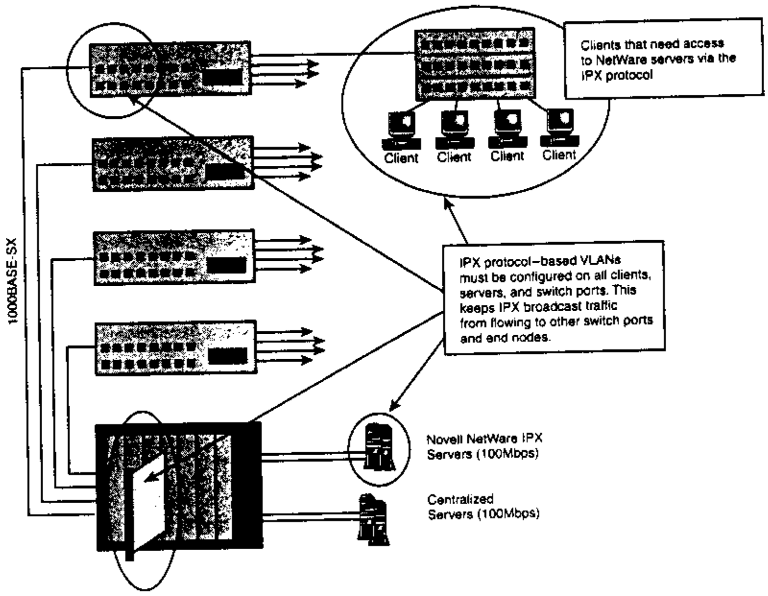

FIGURE 10.7 *IPX traffic generated by NetWare servers and clients is limited, by design, to a specific, predefined VLAN. Although this improves network performance, it requires some manual configuration of each switch involved.*

Example 3: Deploying Switched and Fast Ethernet in a LAN with an Existing FDDI Backbone

Many networks today feature an FDDI backbone ring. FDDI provides a high-speed, robust, semi–fault-tolerant technology that is ideal for backbones. FDDI, however, is considered expensive and difficult to maintain and manage. 100BASE-FX, on the other hand, is not designed with built-in fault-tolerant hardware; however, many 100BASE-FX switch and router products provide this function with port aggregation schemes, load balancing, and specialized software. 100BASE-FX and 100BASE-TX are less expensive than FDDI and provide the same effective data rate: 100Mbps. In addition, 100BASE-FX allows for full-duplex connections of up to 200Mbps data rates. In most cases, it is unwise to replace existing FDDI backbones with 100BASE-T. 100BASE-T, however, can be used as an inexpensive extension to an existing FDDI ring.

This example examines the issues involved in adding 100BASE-T to a network that already has an FDDI backbone. 100BASE-FX can extend the fiber backbone, and 100BASE-TX can distribute data from the FDDI ring down to the workgroups. We will first look at using 100BASE-T to improve the throughput of an FDDI ring to the workgroups.

Adding 100BASE-T Workgroups to an FDDI Ring

In this example, assume that an FDDI backbone exists with 10BASE-T-to-FDDI routers (or advanced bridges) providing the connection from the FDDI ring to the workgroups. Figure 10.8 shows that each workgroup shares 10Mbps of bandwidth provided by an FDDI-to-10-Mbps Ethernet switch, such as the Cisco Catalyst 5000.

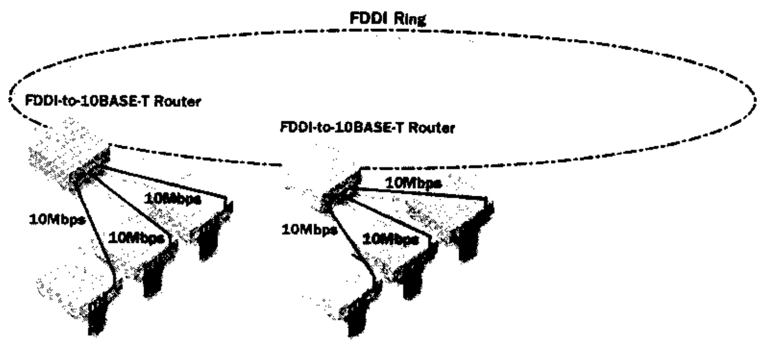

FIGURE 10.8 *A 10BASE-T network with an FDDI backbone. 10BASE-T workgroups access the FDDI ring via routers.*

When adding new workgroups to this network, 100BASE-T should be considered. Assume this particular network uses Category 5 UTP cabling for each workgroup. In this case, 100BASE-TX is the logical selection for 100Mbps connectivity to the desktop. In Figure 10.9, some new desktops have been connected to a 100BASE-TX stackable hub. This workgroup now needs to be connected into the FDDI backbone. Several products that perform this type of connection are available.

One type of product is a standalone FDDI-to-100BASE-TX switch or router, available from Bay Networks, Cisco, Cabletron, 3Com, and many other vendors. Because this is a multiport device, its addition allows multiple 100BASE-TX workgroups to be connected to the FDDI backbone. As in example 1, intelligent logic within the switch or router may use 802.3X congestion control mechanisms to prevent the 100BASE-TX workgroups from oversubscribing the FDDI ring with too much traffic.

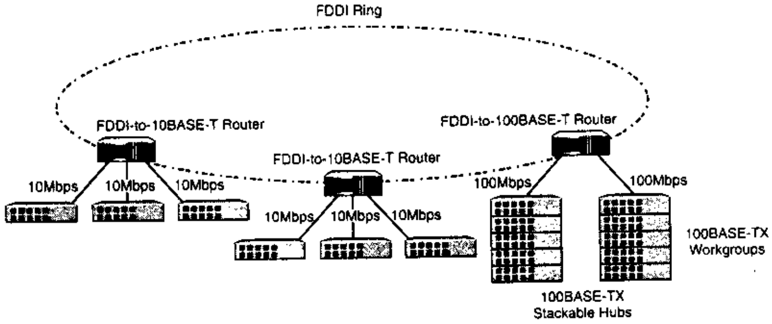

FIGURE 10.9 *Adding 100BASE-TX workgroups to an FDDI backbone. 100BASE-TX is routed to FDDI much like 10BASE-T.*

Extending an FDDI Ring with 100BASE-FX

Another way to add 100BASE-T to an FDDI ring is to consider extending the fiber backbone. In the previous example, the network utilized a distributed FDDI ring, which, in turn, had many FDDI-to-10BASE-T connections. A simple way to extend the backbone to 100BASE-FX or 100BASE-TX is to add an FDDI-to-100BASE-T router in a strategic location. By doing this, the backbone can be extended in a star configuration from the location of the router. If this router is modular in nature, the FDDI ring can be routed to 100BASE-FX, 100BASE-TX, or 10BASE-T via the addition of new router modules. In the example shown in Figure 10.10, the FDDI backbone has been extended to 100BASE-FX via a standalone router.

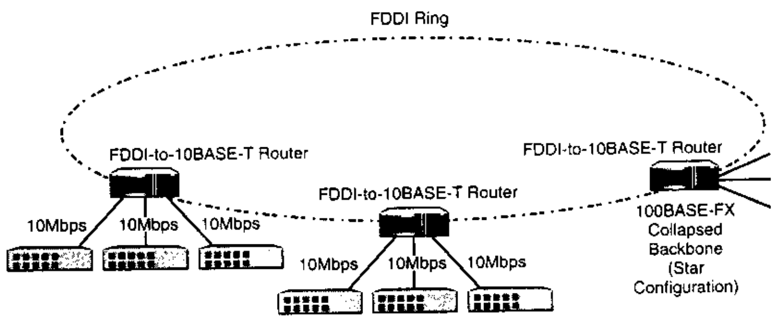

FIGURE 10.10 *Adding a 100BASE-FX to an FDDI backbone allows for the continued use of fiber in the backbone. 100BASE-FX expands on a star topology while FDDI is a ring topology.*

Proponents of FDDI often point to its fault-tolerant, dual-ring design, which helps prevent fiber faults from bringing down the backbone. 100BASE-FX can be set up similarly by using aggregated links. In this way, two switches can be connected with redundant 100BASE-FX links, thus providing similar robustness to the dual-ring architecture of FDDI. Also, there could be a tendency to oversubscribe the FDDI ring with several full-duplex 100BASE-FX ports. It is wise to consider the relocation of certain servers that may have resided on the FDDI ring. It might be more prudent to connect them directly to the 100BASE-FX switch.

Example 4: Deploying a Gigabit Ethernet Switched Backbone

In Chapter 9, we discussed how a Gigabit Ethernet backbone is built, so in this example, we'll focus on specific Gigabit Ethernet products available. This example assumes the same basic structure as the first two examples, but with a few notable differences. First, 100Mbps switching has been carried forward to the desktop. That is, many of the desktops are connected directly to a switched Fast Ethernet port. Second, a large amount of traffic is consistently present on the network. This comes from massive file transfers, high-access database servers, and multimedia applications. Last, network performance is paramount, but budgetary considerations still prevail. These three factors point squarely at Gigabit Ethernet.

Consider Figure 10.11. The Fast Ethernet desktop switch is a 3Com 3300, which provides 24 ports of dedicated Fast Ethernet connectivity and room for a Gigabit Ethernet uplink. Each floor in this example has between 30 and 40 clients, so two 3300s are required per floor. These switches are connected together via a special switching interface that allows two 24-port switches to act like one 48-port switch. This provides 100Mbps connectivity to the high-performance desktops.

Also in Figure 10.11, each stack of two switches is then connected to a Gigabit Ethernet switch. This is done via a Gigabit Ethernet uplink available from 3Com for the 3300. The Gigabit backbone switch used in this example is the Black Diamond 6800, from Extreme Networks. The 6800 offers up to 48 ports of Gigabit Ethernet with 64Gbps of backplane capacity. All four workgroups can be supported with a full gigabit of bandwidth if needed. A lower cost version of this switch could be Extreme's Summit1 switch, which comes configured with eight Gigabit Ethernet ports.

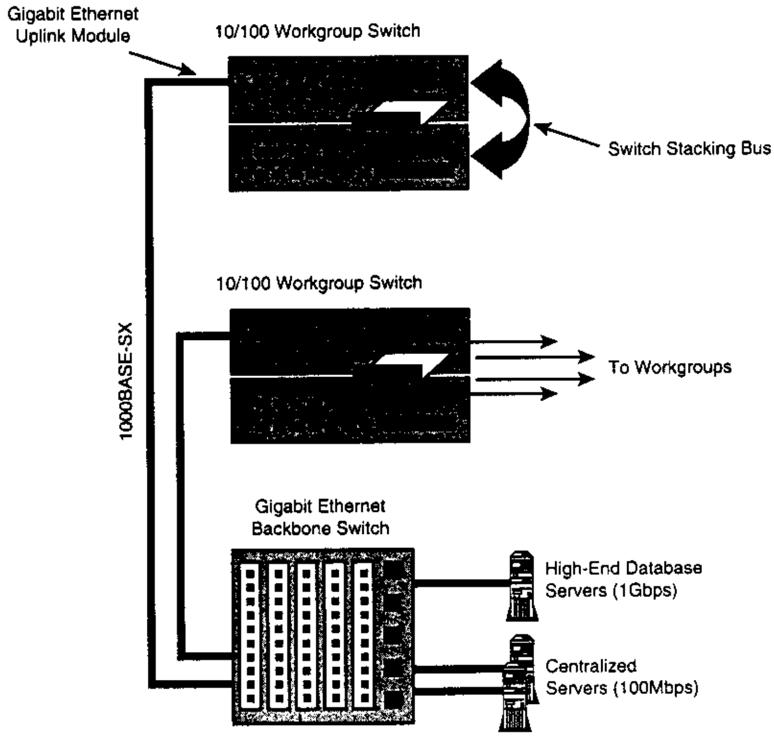

FIGURE 10.11 *A Gigabit Ethernet switch provides the backbone power for Fast Ethernet desktop switches. Two 24-port Fast Ethernet switches stacked together provide 100Mbps dedicated bandwidth to the desktop from a 48-port virtual switch.*

In addition, two very high-end database servers are connected directly to the Black Diamond 6800 Gigabit Ethernet switch. Both of these servers are equipped with an Intel PRO/1000 Gigabit Server adapter. These servers have large amounts of memory, have very fast multiprocessors, and support 64-bit PCI slots.

One could expect this network to outperform the networks deployed in examples 1 and 2 by a factor of 10. Not suprisingly, the actual performance increase will be on the order of 4 to 6 times. The operating system software on both the server and clients will likely limit the performance. The network in this example, however, is ready for faster, future generations of PC and server technology.

Example 5: Deploying a 10-Mbps Switch as a Switch of Hubs

In many cases, neither Gigabit nor fast Ethernet is required. A simple upgrade from 10Mbps hubs to 10Mbps switches might do the trick. In this example, we will examine a case where many clients and servers make up a 10Mbps shared network. The majority of the systems—roughly 50 users—are clients. The rest—about six systems—are servers. This particular network, shown in Figure 10.12, is experiencing average wire utilization above the 40% mark, but is nowhere near requiring 100Mbps of bandwidth. This has caused some concern and a small amount of funds have been budgeted to do something to reduce wire utilization and add bandwidth. An analysis of the network showed that most of the traffic was associated with two of the servers (A and B).

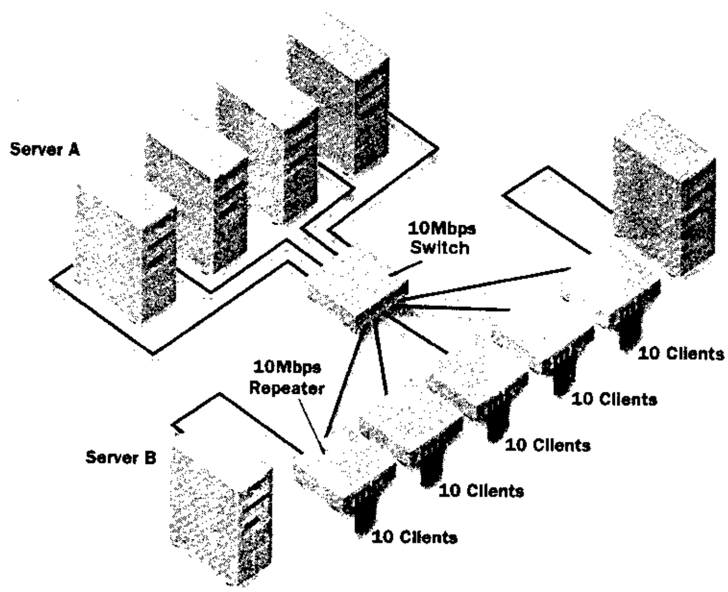

FIGURE 10.12 *A typical shared 10Mbps network.*

By applying the rules learned in step 1 of Chapter 9, it was determined that most of the traffic problems could be alleviated by adding a 10Mbps switch of hubs to this network. Each hub will now have its own dedicated segment with 10Mbps of bandwidth to share among all clients. The servers are also logically relocated on the network. The two busy servers are each given their own dedicated 10Mbps connection, and the other four servers share a 10Mbps connection. The newly configured network is shown in Figure 10.13.

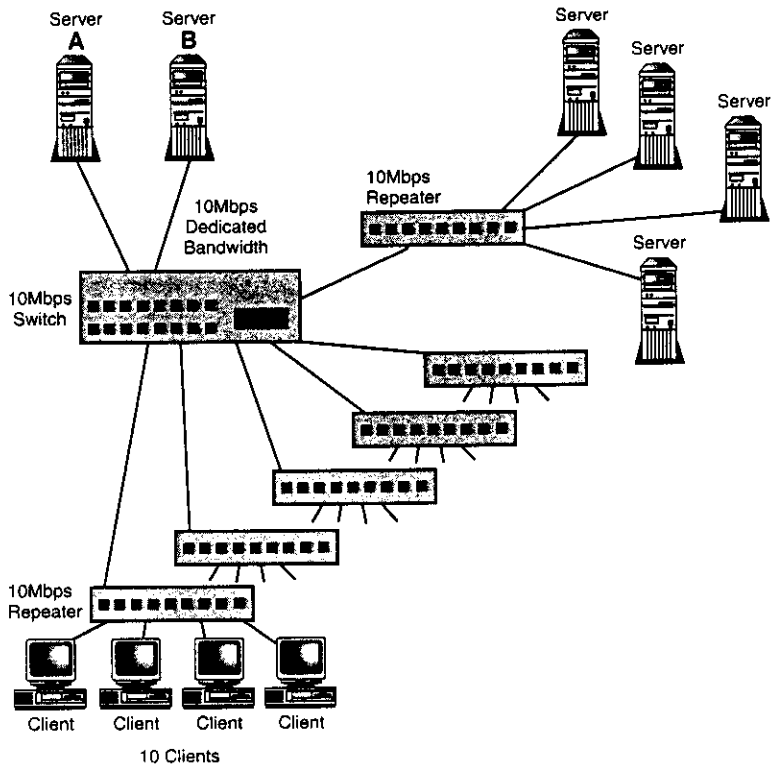

FIGURE 10.13 *The same network, after a 10Mbps switch of hubs has been added. The performance of the network has increased by a factor of 8.*

The new switch should have extra features, such as error-checking and advanced packet filtering, so collided frames and runt frames are not propagated throughout the network. This is a great example of using a well-placed 10Mbps switch to boost network performance without adding Fast Ethernet. Overall, each of the five workgroup segments shares its own 10Mbps of bandwidth, four servers share 10Mbps, and two other servers have a dedicated 10Mbps each. Thus, the total network bandwidth available increases by a factor of 8.

Of course, with the addition of switches come new challenges in network management and troubleshooting. These issues are addressed later in Chapter 11, "Managing Switched, Fast, and Gigabit Ethernet Networks."

Example 6: Deploying Fast Ethernet in a Branch Office

The last example addresses the concept of deploying Fast Ethernet in a branch office. The term *branch office* is often used interchangeably with *small network* or *isolated network* to refer to a network where the number of users is typically small and the amount of traffic is very localized. In addition, a branch office usually has a connection to a WAN link for sending and receiving information from a remote place, such as corporate headquarters. The WAN links are mostly used at certain points in the day, such as the close of business for a bank branch office. The following list points out some of the common traits of a branch office:

- There is a small number of users, typically fewer than 50.
- The physical network diameter is small.
- Replacement of network components requires a special visit to the site.
- The majority of daily traffic is localized.
- There are one or more WAN connections.
- WAN connection use is typically high during specific times and low at other times.

The branch office is a situation in which a shared Fast Ethernet solution excels. Because a branch office can be thought of as one large workgroup (although it is separated from the corporate backbone by the WAN link), the rules of workgroup deployment apply. Think back to example 1 where new workgroups were deployed using 10/100BASE-T stackable hubs. Stackable hub deployment applies in this example just as well. Consider the branch office shown in Figure 10.14. There are eight users, two local servers, and one router that provides a connection to the WAN.

In this particular example, assume that the branch office is an older building with a mixture of Category 3 and Category 5 UTP wiring. We don't know how much of each is installed, though. In this case, a 10/100 hub should be used throughout the network because of the uncertainty in cable quality. There are eight clients and two servers that have been equipped with either 10-only or 10/100-Mbps NICs. It's impossible to tell in advance without making a special trip to the office and examining each PC.

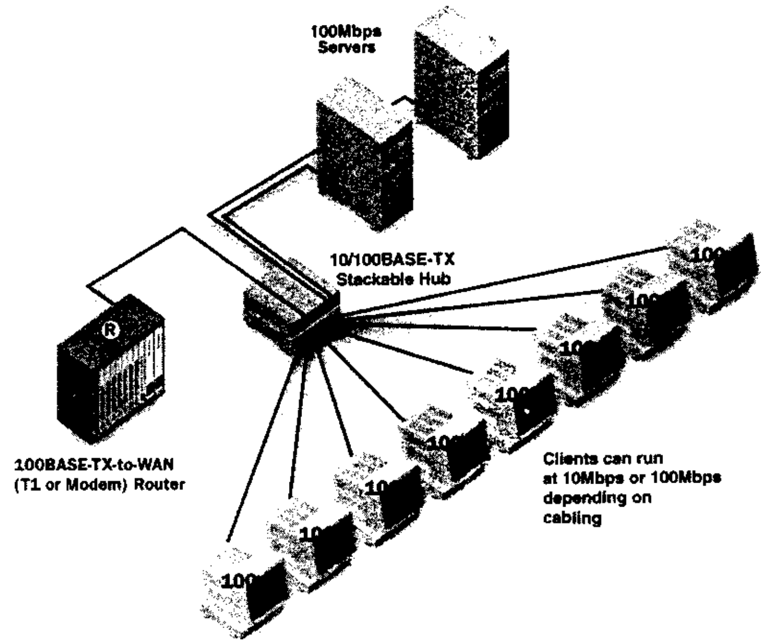

FIGURE 10.14 *A 100BASE-TX branch office. Branch offices usually have a small number of users and a need to access the WAN at regular times. 10/100 capability on the hub is especially important in a branch office because the client connectivity is much less regulated.*

Because it's impossible to make assumptions about how many nodes are 100Mbps-capable, a 10/100 repeater or stackable hub is recommended. In this example a 12-port Compaq Netelligent 2824 will do nicely. Nodes with 10Mbps NICs or connected via Category 3 cabling can operate at 10Mbps, whereas others can operate at 100Mbps.

We can use a simple standalone router to connect to the WAN. Typical WAN connections are T1- or modem-based and therefore much slower than even 10Mbps Ethernet. The routing capabilities of this device will prevent local traffic from being propagated on the WAN and wasting valuable WAN bandwidth. This is especially critical in this example, where the branch office is sharing 100Mbps of bandwidth and the WAN link is a 1.5Mbps T1 line (or worse yet, ISDN or simple dial-up modem access). Performance of the LAN/WAN router is not critical in this case, because the WAN links are slow compared to the LAN links.

In situations like this, we can also use PC-based routing. PC-based routers, which use software such as Novell's NetWare MPR (multiprotocol router), may provide the right level of performance at a low cost. In this example, however, we'll use a standalone WAN access router that supports 10/100 Ethernet connectivity: the Intel Express 8100 router.

In the Real World, Every Network Is Different

Examples allow us to gain insight into deployment issues, but in reality, every network is different. Although your entire network will not be exactly like one of the six cases discussed in this chapter, it is likely that parts of your network will be similar. These examples are designed to illustrate basic principles and highlight the deployment issues associated with Switched, Fast, and Gigabit Ethernet. The examples discussed in this chapter should not be used as strict rules, but rather as guidelines for your own Switched, Fast, and Gigabit Ethernet deployment strategies.

So far, we have looked at what defines Switched, Fast, and Gigabit Ethernet and how to deploy it in your network. In Chapters 11 and 12, we will discuss what to do with your Switched and Fast Ethernet network after it is installed: specifically, how to manage and troubleshoot a network that is both similar to and different from your 10Mbps Ethernet network of today.