CHAPTER 9

Deployment

In general, sources of information (servers), requesters of that information (clients), and communication paths for that information (the network) can represent an entire network. In network design, it is critical to understand the effects of server placement on various groups of clients. For instance, an Internet gateway shared by all clients on a network should be located at a central point. This is often the network backbone. A local mail, print, or file server that primarily services a set number of clients should be located in the workgroup. By locating the server close to the people who use it, you reduce network traffic on the other parts of the network. This concept is crucial to understand before embarking on a full-scale upgrade of your network.

Understanding the location of the sources and sinks of the network data will help you determine what kind of Switched, Fast, or Gigabit Ethernet network you need. This chapter focuses on seamlessly migrating your existing clients, servers, and 10Mbps shared Ethernet networks to incorporate Switched Ethernet, Fast Ethernet, Gigabit Ethernet, and even Layer 3 switching.

Switching and Ethernet go hand in hand because each needs the other for widespread deployment. Switched Ethernet is usually not of great benefit without a high-speed link, or fat pipe, to collect the various 10Mbps data streams. Fast and Gigabit Ethernet need switching to overcome basic diameter restrictions that would otherwise prevent them from being installed in all but the smallest of networks.

Switching and Fast Ethernet also combine well because switching is often thought of as a top-down technology, required in high-end solutions. Fast

Ethernet, on the other hand, is thought of as a bottom-up technology: a logical extension to the 10Mbps desktop of today. (A top-down technology is one that is implemented first in the backbone, whereas a bottom-up technology finds its way first into desktops and workgroups. Gigabit Ethernet is another top-down technology.) When used together properly, Switched, Fast, and Gigabit Ethernet products can increase performance in the desktop, the server, the backbone, and the network as a whole.

This chapter begins with the basic rules of deployment and installation for switched and shared 10BASE-T and 100BASE-T, as well as fiber rules for Gigabit Ethernet. We make many references to Part II, so we assume that you have a good understanding of the relevant standards and cable types. After discussing the general rules, weíll categorize the actual deployment of Switched, Fast, and Gigabit Ethernet into six steps. Each step has a specific role in the overall conversion from shared 10BASE-T to Switched 10BASE-T, 100BASE-T, and 1000BASE-SX/LX. Table 9.1 outlines these six steps.

Depending on how large or overloaded your network is, you may only want to implement some of the steps. For instance, if you implement 10Mbps switches, as discussed in step 1 and you find network performance acceptable, you may not need to proceed through the following steps. The steps are structured in such a way that you can implement them slowly, even over the course of a few years. In fact, implementing these steps slowly is often the best course of action when balancing network performance with risk mitigation. Note that as you implement the steps, network performance goes up, but the risk rises as well. Even though adding Gigabit Ethernet in the backbone is a high risk, it is still a lower risk than converting the backbone to a more complex technology, such as ATM.

TABLE 9.1 SIX STEPS TO DEPLOYING SWITCHED, FAST, AND GIGABIT ETHERNET

Step	Phase of Deployment	Performance Impact	Risk Relative to That of Step 1	Cost Relative to That of Step 1
1	Add 10Mbp sswitches to your current network	Minimal	None	None
2	Deploy Fast Ethernet for the first time (10/100Mbps+ NICs into new clients and servers)	Minimal	Minimal	Minimal
3	Convert workgroups to Fast Ethernet	High/ minimal[1]	Medium	Medium

Step	Phase of Deployment	Performance Impact	Risk Relative to That of Step 1	Cost Relative to That of Step 1
4	Deploy Switched and Fast Ethernet in the backbone	High	Medium	High
5	Implementing Layer 3 Switching, VLANs, and Fast Ethernet routing in the backbone	Medium	High	High
6	Adding Gigabit Ethernet switching to the backbone	Very high	High	High

[1] *High when accessing resources local to that workgroup. Minimal impact accessing centralized resources.*

The term relative risk is used. The risk outline in Table 9.1 is relative to that of implementing step 1. As you can see from the table, adding a Gigabit backbone switch is a high risk compared to adding a 10Mbps workgroup switch.

We'll break these steps down into deployment examples later in the chapter. (Many of the deployment options reference Chapter 8, "Network Components," where we more fully explain each network component and its associated features.) We then discuss specific deployment examples from actual networks in Chapter 10, "Deployment Examples."

This chapter begins by covering the general rules of 10BASE-T and 100BASE-T.

General Rules of 10BASE-T and 100BASE-T

The basic rules of deployment for Switched, Fast, and Gigabit Ethernet depend upon the choice of cabling. Category 5 unshielded twisted pair and multiple-mode fiber are common choices for today's networks. Although less common, single-mode fiber and inherited Category 3 UTP can enhance or reduce your ability to deploy Switched and Fast Ethernet network components. When it comes to Gigabit Ethernet, fiber is really the only reliable option today; although Gigabit Ethernet over Category 5 UTP may be just around the corner.

The 100m and 5m Rules

As we saw in Chapter 6, the EIA/TIA 568 cabling standard recommends 100m from hub to desktop in all UTP cabling infrastructures. This specification is mirrored in the international standard ISO 88023 (see www.eia.org for more details). The 100m is further broken down into the following distances:

- 5m from hub to patch panel
- 90m from patch panel to office punch-down block
- 5m from punch-down block to desktop connection

Most UTP installations conform to this 100m rule, as it is commonly called, which makes 100BASE-T very easy to install.

Cable installers also recommend that hub-to-hub UTP connections be made with 5m or less of cable (the 5m rule). Short cables in a noisy wiring closet translate to less induced noise on the wire and less crosstalk in large multiple-cable bundles. Short cables however, may restrict hub location in large wiring closets, so this guideline is often overlooked. In 10BASE-T networks, this rarely causes a problem, but when installing Fast Ethernet shared workgroups, the 5m rule should be strictly followed. We'll explain the reasons for this later in the chapter.

10BASE-T Shared Network Rules

10BASE-T requires that all collisions be resolved within 512 bit times (also called one slot time). In a 10BASE-T shared network, each component, including the cabling, adds transmission delays, which account for a shrinkage in the total network diameter. Today's technology allows for a worst-case 10BASE-T UTP network with roughly four repeater hops and three populated segments. Why only four hops? In order to lock on to an incoming signal, each repeater eats up bits of the signal. This can be accounted for as a loss of bit budget or network diameter. Each cabling segment and repeater represents a certain transmission delay, and the total round-trip delay cannot exceed one slot time, or 512 microseconds.

With current repeater technology, this results in 10BASE-T networks of no more than four repeater hops. So even though Ethernet's collision domain is specified at 2500m, in 10BASE-T form it rarely exceeds 500m. The 10BASE-F (Ethernet over fiber) allows for much larger shared 10Mbps networks due to the extended transmission length of fiber, but 10BASE-F is typically implemented as a switched connection.

10BASE-T Switched Network Rules

The 10BASE-T switched networks are no different from 10BASE-T shared networks, except that a new network diameter calculation begins at each switch port. Because switches provide dedicated connections, fewer collisions occur, and no real collision domain exists to restrict diameter. The 10BASE-T switches are only limited by the same EIA/TIA 568 rules that govern current installations: the 100m hub-to-node and 5m hub-to-hub rules. Therefore, a 10Mbps switched network will work in any existing EIA/TIA 568 compatible UTP cabling infrastructure with no network diameter constraints. Refer to Chapter 6 for more information on cabling standards.

100BASE-T Shared Network Rules

As discussed in Chapter 3, "Ethernet, Fast Ethernet, and Gigabit Ethernet Standards," shared 100BASE-TX and 100BASE-T4 networks require a much smaller collision domain: only 205m. This allows for a 100Mbps shared network of two repeaters with 100m cabling to each node and 5m cabling between repeaters. As you can see, purely shared Fast Ethernet networks require exact compliance to EIA/TIA UTP cabling specifications. If a network diameter greater than 205m is required, you must use switches, bridges, or routers somewhere.

Class I 100BASE-T repeaters and stackable hubs are further limited to only one repeater hop because they incur the additional delay of converting incoming analog data to the digital MII interface. The 100BASE-TX-to-100BASE-T4 translational repeaters are Class I by definition. Most Fast Ethernet stackable hubs are Class I.

Class II 100BASE-T repeater ports are usually all of one type (either TX or T4) and therefore allow two repeater hops. Figure 9.1 shows maximum shared network diameters for Class I and Class II stackable hubs.

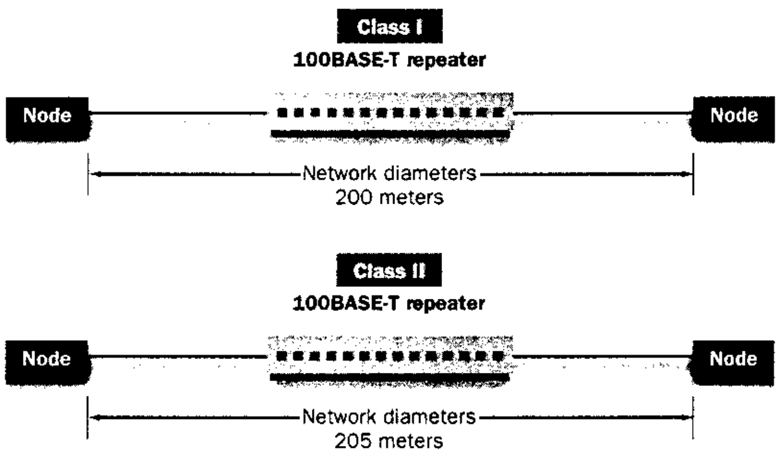

FIGURE 9.1 *Class I and Class II shared media network diameters are limited to 200m and 205m, respectively.*

In general, 100BASE-TX is less fussy about shared network restrictions than 100BASE-T4 because its bit budget has more flexibility. Some 100BASE-T4 components, however, allow for extended cable lengths over 100m or more repeater hops due to innovative designs. You should do a bit-budget analysis on any 100BASE-T installation that exceeds the IEEE 802.3u specification: more than two repeater hops or greater than 205m in network diameter (refer to Chapter 3 for details). 100BASE-T4 is also lower frequency than TX and therefore less susceptible to noise and crosstalk. 100BASE-T4 never really caught on in mainstream networks and therefore is only marginally supported by major networking vendors today. In general, you should avoid installing 100BASE-T4 in a network unless you absolutely cannot upgrade the Category 3 UTP cabling infrastructure.

100BASE-T Switched Network Rules

The restrictions of shared 100BASE-T networks further reinforce the importance of switching in Fast Ethernet deployment. Rarely do 100BASE-T workgroups have more than one repeater hop between the end-station and a Fast Ethernet switch. In other words, shared 100BASE-T networks are most likely deployed in workgroups via stackable hubs connected to switches. This means that in a properly constructed shared/switched 100BASE-T network, the 100BASE-T network diameter of 205m will never come into play. Switching is imperative to the successful deployment of Fast Ethernet. In fact, the bulk of this chapter is dedicated to explaining how switched and shared 10Mbps, 100Mbps, and 1000Mbps media can be blended together to form small and large LANs.

100BASE-FX Fiber Network Rules

Another necessity of today's LANs is the fiber optic connection. Fiber optic cabling rarely connects desktops but is commonly found in switch-to-server, switch-to-switch, and backbone applications (see Table 9.2). Due to its capability to carry signals for distances of up to 2km, fiber allows multiple campuses to connect to the same backbone. Currently, you'll commonly find FDDI running on multimode fiber optic cabling, which is why 100BASE-FX is designed for the same type of multimode fiber. The 100BASE-FX specification, also categorized under IEEE 802.3u, allows for many levels of extended distances, depending on the type of connection.

TABLE 9.2 TYPES OF 100BASE-FX FIBER CONNECTION

Type of 100BASE-FX Connection	Description of Rules
Shared-to-shared	In a true 100BASE-FX fiber repeater setting, the maximum distance from hub to node is 160m in accordance with the 100BASE-T bit budget. For this reason, 100BASE-FX repeaters will be in limited use.
Shared-to-switched	If one side of the connection is to a 100BASE-FX fiber switch, the same 100BASE-FX repeater mentioned above will be able to send signals over 210m of fiber. Again, this will limit the practical use of the properties of 100BASE-FX.

Type of 100BASE-FX Connection	Description of Rules
Switch-to-switch	A 100BASE-FX switched port connected to another 100BASE-FX switched port can transmit over 412m of fiber optic cable. This will be the entry point for most 100BASE-FX backbone products.
Switch-to-switch full-duplex 100BASE-	The 100BASE-FX specification calls for a special full-duplex switch-to-switch connection that allows for 2km of fiber cabling between switches. This allows FX to be run anywhere FDDI is used today. This represents the bulk of the 100BASE-FX switch market.

100BASE-FX is mainly used to extend networks to multiple floors or buildings. 100BASE-FX deployment beyond this function is a rarity.

Gigabit Ethernet Network Rules

Gigabit Ethernet also relies on fiber optic cabling. You can connect Gigabit Ethernet over single-mode fiber or one of two types of multiple-mode fiber. The main requirement for Gigabit links is a switched, full-duplex connection on each end of the fiber. Although Gigabit Ethernet can be deployed in a buffered-repeater design, most network managers opt to skip the shared approach and go directly for Switched Gigabit. We'll discuss the reasons for this later. For now, consult Table 9.3 for network diameters associated with Gigabit Ethernet.

TABLE 9.3 NETWORK RESTRICTIONS FOR 10BASE-T, 100BASE-T, AND 1000BASE-SX/LX

Technology	Network Diameter Shared Medium	Network Diameter Switched Medium	Switch-to-Switch Cable-Length Switched Medium
10BASE-T	About 400m	Unlimited	100m
100BASE-TX/ 100BASE-T4	205m	Unlimited	100m
100BASE-FX	320m	Unlimited	400m
100BASE-FX Full-Duplex	N/A	Unlimited	2000m
1000BASE-SX	Only possible with a buffered repeater	Unlimited	260m multimode fiber
1000BASE-LX	Only possible with a buffered repeater	Unlimited	550m multimode fiber, 3km single-mode fiber

The IEEE in the IEEE 802.3ab subcommittee is also looking at Gigabit Ethernet over UTP. This will likely result in a Gigabit Ethernet–over copper IEEE specification sometime in mid-1999. Several vendors have made a solid commitment to bring out Gigabit Ethernet copper products, but those products are still forthcoming. The risk-conscious network manager should avoid Gigabit Ethernet over UTP until the standard has been fully ratified. Don't take this to mean that you should avoid Gigabit Ethernet over UTP cabling; just make sure your vendor adheres to the eventual IEEE standard.

The Golden Rules of Switched, Fast, and Gigabit Ethernet

The same specifications, the EIA/TIA 568 and the ISO 88023 standards, govern the general rules of both 10BASE-T and 100BASE-T deployment. Each specification outlines basic guidelines on how UTP and fiber-based network technologies should be deployed. In general, we can summarize the guidelines for Switched and Fast Ethernet deployment in five golden rules:

- Use 100m maximum UTP connection from desktop to hub.
- Stack Fast Ethernet hubs; don't cascade them.
- Use 100m maximum UTP connection from Fast Ethernet hub to Fast Ethernet switch.
- Use 160m maximum fiber connection from Fast Ethernet hub to Fast Ethernet switch.
- Use 2km maximum full-duplex fiber connection from Fast Ethernet switch to Fast Ethernet switch.

In addition, three new rules generally govern Gigabit network infrastructures. They may not give you all the flexibility that Gigabit Ethernet allows, but by following them, you'll be sure to stay within specification. These rules are as follows:

- *Multiple-mode SX*—Use 260m maximum multiple-mode, full-duplex fiber connection from 1000BASE-SX Gigabit Ethernet switch to GE switch or server NIC.
- *Multiple-mode LX*—Use 550m maximum multiple-mode, full-duplex fiber connection from 1000BASE-LX Gigabit Ethernet switch to GE switch or server NIC.
- *Single-mode LX*—Use 3km maximum single-mode, full-duplex fiber connection from 1000BASE-LX Gigabit Ethernet switch to GE switch or server NIC.

Note

By using a 50μm core multiple-mode fiber optic cable, you can extend Gigabit Ethernet connections over longer distances. See Chapter 6 for details.

An understanding of the general rules for switched and shared 10BASE-T, 100BASE-T, and 1000BASE-LX/SX deployment is essential before discussing how to implement them in a production network. The rest of the chapter is dedicated to the six steps of deploying Switched, Fast, and Gigabit Ethernet.

Deployment Step 1: Adding 10Mbps Switches

The first step toward upgrading your network involves determining where to implement 10Mbps switches. This section focuses on implementing 10Mbps switches in two applications: workgroups and switches of hubs. (Note that step 1 deals with 10Mbps-only switches. We discuss 10Mbps switches with 100Mbps uplinks in step 2.)

In the workgroup, you can effectively implement a 10Mbps switch in two ways:

- *Standard workgroup*—A workgroup switch is used as a replacement for 10Mbps repeaters in workgroups with clients and local servers.
- *Switch of hubs*—A workgroup switch is configured to provide 10Mbps of dedicated bandwidth to individual stackable hub units. In this configuration, a 10Mbps switch may look like a small, switched backbone.

We explore these two deployment options, plus some of their derivatives, more thoroughly in the pages to follow. Because this is the first step in an overall deployment of Switched and Fast Ethernet, the performance gained will be moderate, but the timing of the implementation can be immediate.

When buying 10Mbps switches for any kind of deployment, always look for products that can be upgraded later with 100BASE-TX or 100BASE-FX uplinks. You can usually accomplish this with upgrade slots. An upgrade slot, as described in Chapter 8, allows for the addition of newer technology into the switch, long after the original purchase. Not only does this allow a 10Mbps switch to connect to a 100Mbps backbone, but it is typically flexible enough to allow other types of modules (ATM, Layer 3 routing, and so on). A 10Mbps switch without 100BASE-T uplink capability is pointless and not a good investment.

The obvious benefit of deploying a 10Mbps switch is that you don't have to replace the 10BASE-T NICs already installed in desktops and servers. For this reason, we can consider 10Mbps switching a drop-in replacement. 10BASE-T, in both switched and shared form, runs equally well over Category 3 and Category 5 UTP cabling. For this reason, many companies with suspect cabling, or no re-cabling budget are opting for 10Mbps switching as their core network architecture.

We have already heard that ISA-based desktops and PCMCIA-based laptops will not gain much from the addition of a 10/100Mbps NIC. Therefore, any new systems based on these buses may be configured with a classic 10Mbps NIC. These types of systems will get more than enough network bandwidth from a 10Mbps switched connection. Also, if the system has particularly slow subsystems, like a slow server hard disk, you may determine that a 100Mbps connection is not needed. For instance, a server with a higher speed bus, like an EISA bus, but with a slower processor (Intel 80486 or lower) may not be a good candidate for a 100Mbps connection. The flexibility gained from a 10/100 solution is sometimes worth the cost, though. For instance, if you use 10/100 PCMCIA cards, you may never reach 100Mbps data rates, but you won't worry about finding the right speed hub port to connect to.

Standard Workgroup Deployment

Many clients, a few local servers, and a 10Mbps repeater compose the typical workgroup of today. By adding a 10Mbps switch, you can convert a shared workgroup to a higher speed switched workgroup. Consider the diagrams in Figure 9.2, where a workgroup has four clients and two servers connected through a 10Mbps repeater. All nodes share the repeater's 10Mbps of bandwidth. By replacing the repeater with a 10Mbps switch, you can increase the available bandwidth of the workgroup by several fold. In the switched example, each client has a dedicated 10Mbps connection, and each server has two dedicated 10Mbps connections. Each server can provide up to 20Mbps of data so that the new workgroup bandwidth is 40Mbps, or four times the original. In addition, the concept of simultaneous server access now becomes a reality with switches. One client can be accessing one local server while another is accessing a different local server. This type of switch deployment works well in small networks or isolated workgroups where most of the traffic stays local.

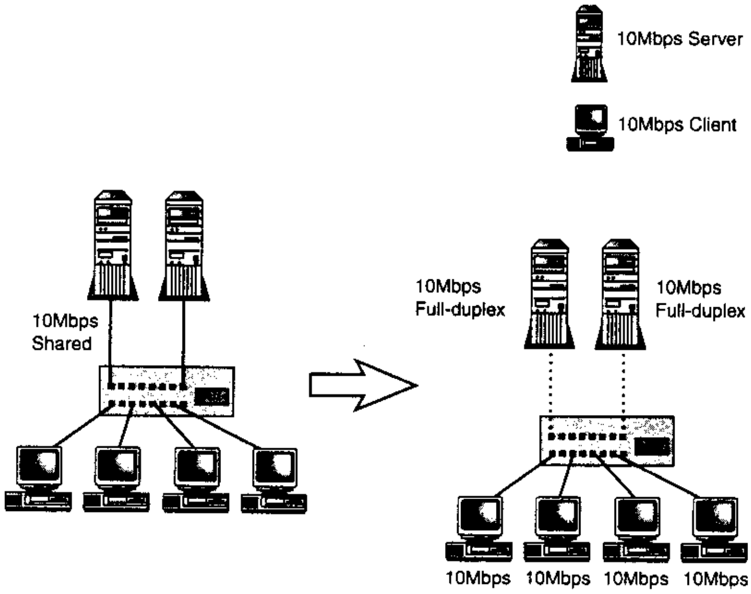

FIGURE 9.2 *10Mbps switches in a standard workgroup.*

When does it make sense to install a 10Mbps switch in this configuration? A few key indicators will give you a good idea of whether 10Mbps switching will help. These include counting the number of local servers in the workgroup and examining the amount of traffic that stays locally within the workgroup. Table 9.4 explains how these indicators determine whether 10Mbps switching is needed. If a workgroup exhibits one or more of these characteristics, treat it as a candidate for a 10Mbps switch upgrade.

TABLE 9.4 DETERMINING THE NEED FOR 10MBPS SWITCHING IN A STANDARD WORKGROUP

Workgroup Characteristic	Description
Multiple local servers in the workgroup	If the workgroup has only one server, the cost of a 10Mbps switch may not be worth the small performance gain. In cases where only one server is present in the workgroup, other features such as multiple server NICs, port aggregation, and full-duplex are needed to increase performance.
Primarily local traffic	At least 80% of the activity is between the clients and local servers.
Large amount of traffic	The 10Mbps shared workgroup is at least 30% utilized (that is, more than 3Mbps is utilized).

Once the need for a 10Mbps switch has been determined, the next question is what type of switch to deploy. Some key switch features to look for in a standard workgroup deployment are full-duplex support, port aggregation, cost per port, and slots for add-in uplink modules. This type of workgroup switching solution allows network administrators to make the most out of what equipment they already have. No new NICs are required for the clients and a 10Mbps switch is a quick replacement for a 10Mbps repeater.

Using a Switch of Hubs

A 10Mbps switch also improves performance in larger networks based on a number of 10Mbps stackable hubs.

Many large networks are configured with stackable hubs, which are typically used when connecting a large number of users in a workgroup. In this case, you can use a 10Mbps switch as a switch of hubs (as shown in Figure 9.3) to improve the performance of a 10Mbps stack. Instead of sharing a stacking bus, each unit of the stack is provided with 10Mbps of dedicated bandwidth from the switch. Once again, you can accomplish this with little impact on the network infrastructure. In fact, you can often physically place a 10Mbps switch right on top of the existing stack.

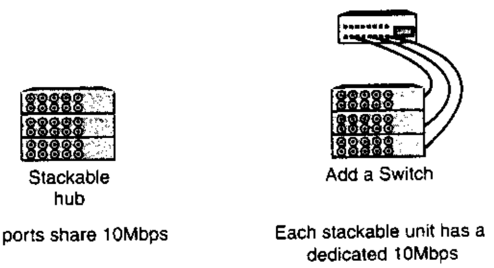

FIGURE 9.3 *Upgrading a 10Mbps stackable hub with a 10Mbps switch. This is often referred to as a "switch of hubs" configuration.*

The performance gained is typically proportional to the number of switching ports used. In Figure 9.3, there is a five-fold performance gain (50Mbps versus 10Mbps). The main things to look for when using a switch in a switch of hubs configuration are port densities (are there enough switched ports for each stackable unit?) and the implications for network management. Because the management module in the main stackable unit can no longer "see" all network traffic (the switch blocks it), an alternative type of management must be used. For more on the management of switches in this configuration, see Chapter 11, "Managing Switched, Fast, and Gigabit Ethernet Networks."

Aggregating Ports

Many networks are constructed entirely of 10Mbps repeaters. In this case, all end nodes share the bandwidth available on the medium. When advanced network components such as bridges, routers, or switches are deployed, traffic can begin to be filtered and divided, thus increasing network performance. Networks with bridging, routing, or switching isolate sub-LANs because they can filter unwanted traffic. In this respect, a 10Mbps switched backbone can be extremely useful in improving the overall performance of a 10Mbps shared network. The deployment of a switch in such a network is analogous to the switch-of-hubs concept used for stackable hubs. Each independent 10Mbps repeater is connected to a dedicated 10Mbps pipe, as shown in Figure 9.4.

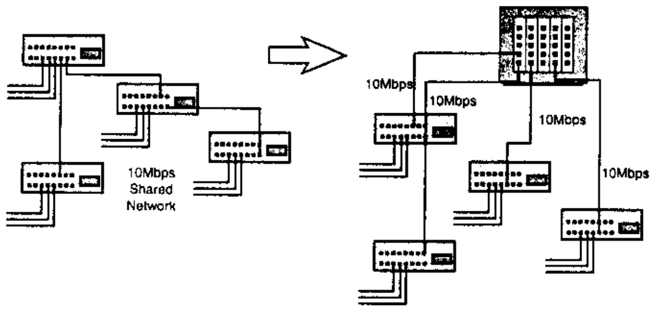

FIGURE 9.4 *Upgrading a 10Mbps shared network with a 10Mbps switched backbone. Total network bandwidth has increased by a factor of four.*

The theoretical network performance of the example in Figure 9.4 has increased from 10Mbps to 40Mbps by adding a single switch. This is very similar to adding four bridges to the network to isolate traffic between segments. In addition, you can cascade 10Mbps switches together with aggregated 10Mbps switched, full-duplex connections to improve throughput between switches. Port aggregation techniques such as EtherChannel or the upcoming IEEE 802.3ad standard are useful for this purpose. This type of deployment is really the beginning of a distributed backbone built from 10Mbps switches. This architecture is scalable as long as additional switch-to-switch connections are available and the port aggregation technique is similar throughout the network. An example of this is shown in Figure 9.5, where a new switch is connected via two 10Mbps, full-duplex switched connections. The interswitch bandwidth is therefore 40Mbps: 20Mbps in each direction.

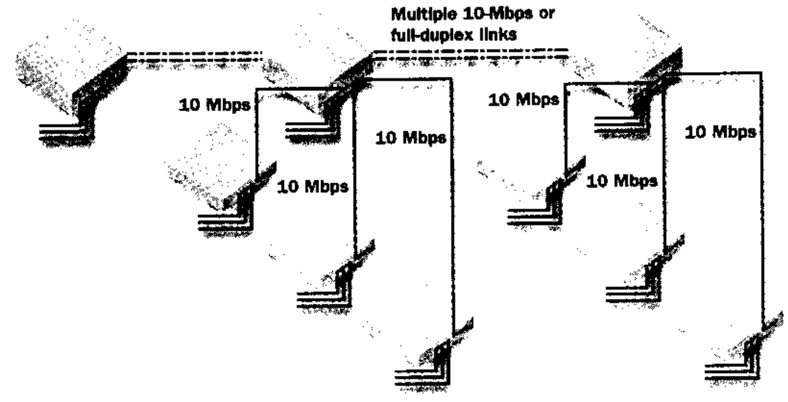

FIGURE 9.5 *10Mbps switches can be configured for multiple backbone connections. These connections may also be full-duplex.*

Server Farms: A Practical Appli-cation of 10Mbps Full-Duplex Ethernet Switching and Port Aggregation

A *server farm* is a collection of servers that reside on the backbone of a network. Servers in server farms are usually high-end systems that service a great number of users. Most servers in server farms are accessed indiscriminately by a large number of users, so the traffic patterns are fairly constant, as opposed to those of servers in the local workgroup. Data tends to come in regularly from all sorts of sources. Installing a 10Mbps switch in a server farm is yet another way to improve network performance with 10Mbps switching. Adding a 10Mbps switch in a server farm provides each server with dedicated bandwidth. The switched architecture also allows for aggregated connections to the backbone in order to handle the increased traffic flow from the server farm.

A typical 10Mbps shared server farm is upgraded by adding a 10Mbps switch. The connection to each server is scalable by connecting additional 10Mbps dedicated lines to the server farm switch. The trick of scaling the connection is to weigh the amount of traffic typically generated by the switched server farm with the number of backbone connections. For instance, if the server farm typically generates an average load of 20Mbps with peaks of

continues

continued

30Mbps, you should use three aggregated 10Mbps backbone connections. In the case of multiple backbone connections, the switch must incorporate some sort of traffic balancing, similar to the load balancing of the workgroup scenario, to allow for the most efficient use of the switched ports. The IEEE 802.3ad port aggregation standard is taking these features into consideration.

Because the data stream contains traffic from many different nodes, the possibility of corrupted packets is higher than normal. Therefore, store-and-forward switches with some measure of error-checking and packet filtering is preferred in server farm applications. Error-checking reduces the amount of work the server NIC and network operating system have to do to filter unwanted packets. A 10Mbps switch might be a good short-term solution for a server farm, but as you'll see later in step 4, a 100Mbps switch is the solution of choice. Whenever possible, try to deploy 100Mbps switched connections to individual servers in server farms.

Managing Switches

Managing 10Mbps switches, or any switch for that matter, is a subject that has already spawned several books. Chapter 11 deals with this issue in more detail, but in this section, we'll focus on what to look for in step 1 of your deployment process. First, only select switches that support SNMP and RMON management. Because this is typically built into the switching silicon, it should be included in the base price of the switch. You'll learn more about this later, but suffice it to say that RMON Groups 1, 2, 3, and 9 are important for 10Mbps workgroup switches. Other important features to look for are support for your software management application (including Hewlett-Packard's OpenView and Bay Network's Optivity) and support for WWW-based configuration and management. Many switches will include an embedded WWW-server that will allow you to access the switch directly from your favorite browser.

Knowing When to Upgrade

Although you have many ways to improve network performance with 10Mbps switches, definite trade-offs are associated with deploying this technology. For example, many network management issues associated with switching must be addressed. Also, depending on how your networking traffic grows, 10Mbps switches may only support your network for a short time before high bandwidth devices are required. New, high-performance servers may quickly swamp a 10Mbps switch with data. Even with these potential drawbacks, however, the advantages of 10Mbps switches are numerous. You can attain increases in performance with little or no impact on the current network architecture (for instance, NICs rarely need to be replaced when 10Mbps switches are deployed). Also, 10Mbps switching is relatively inexpensive compared with other switching alternatives, such as switched Fast Ethernet, Gigabit Ethernet, and ATM.

In general, 10Mbps switches enhance network performance in two types of applications, which are summarized in Table 9.5.

TABLE 9.5 A SUMMARY OF 10MBPS WORKGROUP SWITCH DEPLOYMENT EXAMPLES

10Mbps Workgroup Switch Installation Type	Description	Deployment Issue
Standard workgroup	Provides dedicated pipe(s) to local clients and servers in any given workgroup.	Helps when most traffic is local. Workgroups with multiple local servers benefit the most. Full-duplex NICs and switch ports give added benefit. Load balancing software/hardware is sometimes needed.
Switch of hubs	Provides dedicated pipe to individual repeaters or stackable hub units.	Improves performance in almost all cases. Switch can be physically stacked with rest of stackable units. May affect stackable hub management strategy. Allows server farms to be deployed. Aggregating switch-to-switch connections starts a distributed backbone.

Deployment Step 2: Adding Fast Ethernet NICs

The second step in deploying Switched and Fast Ethernet is putting the first Fast Ethernet network components in place. The natural entry point for Fast

Ethernet in most networks is at the desktop and server. This is primarily because almost all NICs are 10/100Mbps, which means that they can operate at either 10Mbps or 100Mbps. A 10/100Mbps NIC installed today will typically run at 10Mbps for some time before a Fast Ethernet hub or switch is purchased. This section describes how to prepare new desktops and servers with 10/100Mbps NICs and how to plan your network architecture to accommodate a mixture of 10Mbps and 100Mbps workgroups.

Figure 9.6 is International Data Corporation's view of the 10/100Mbps NIC, hub, and switch markets. It is easy to see the causal relationship the 10/100Mbps NIC market has on the 10/100 Switch and Hub market. 10/100Mbps NICs were deployed into existing 10Mbps networks in order to future-proof new PCs. This, in turn, drove demand for 10/100Mbps hubs to connect those desktops and for 10/100Mbps switches to connect those hubs. Now, the prevalence of 10/100Mbps switching is causing demand for Layer 3 switching and Gigabit Ethernet in the backbone.

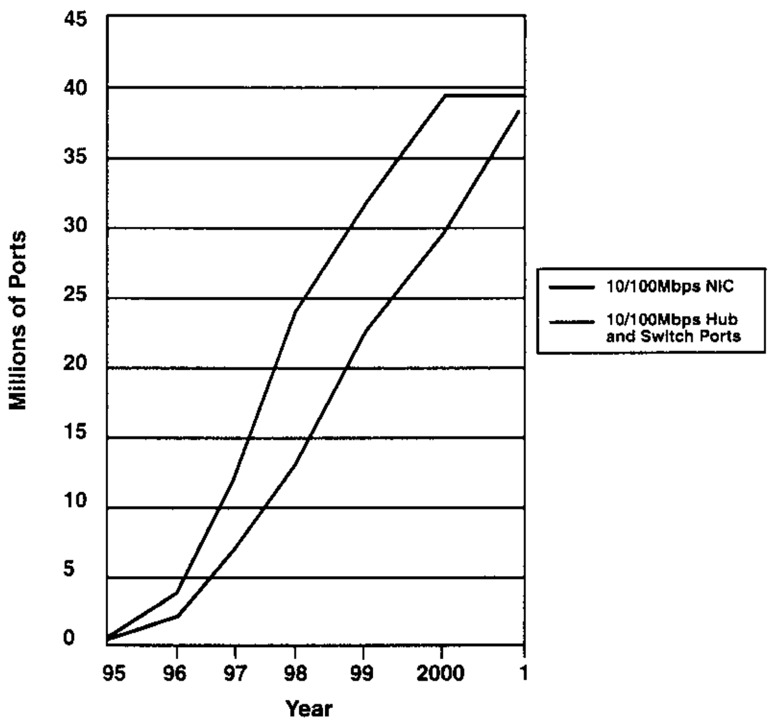

FIGURE 9.6 *IDC's predictions of the 10/100Mbps NIC, Hub, and Switch markets. Note how the 10/100Mbps NIC market leads the hub and switch market by approximately two years.*

10/100Mbps Future Proofing

A wide assortment of desktops and servers are connected to 10Mbps Ethernet ports today. With such a variety, it is practically impossible for new 10/100Mbps cards to fit into every existing system. That is why a practical plan for deployment of Fast Ethernet begins with installing desktop and server 10/100Mbps NICs in newly purchased or newly installed systems. As discussed in Chapter 6, older systems based on the ISA or PCMCIA bus are not ideal for Fast Ethernet upgrades; thus, older desktops and laptops should retain their 10Mbps-only NICs until they are retired from service. The next section (step 3) will discuss how to connect these systems to your new Fast Ethernet network.

Consider the simple network shown on the left side of Figure 9.7. In its current state, each existing desktop and server is connected by a 10Mbps NIC. The network hubs are 10Mbps repeaters. Now, two new servers and four new clients are to be added to this network to connect new employees to the LAN. These systems are enabled with 10/100Mbps-Mbps NICs and connected to the existing 10Mbps network, as shown on the right side of Figure 9.7. Note that these systems have been future-proofed in that they are 100Mbps-ready yet are running at 10Mbps today.

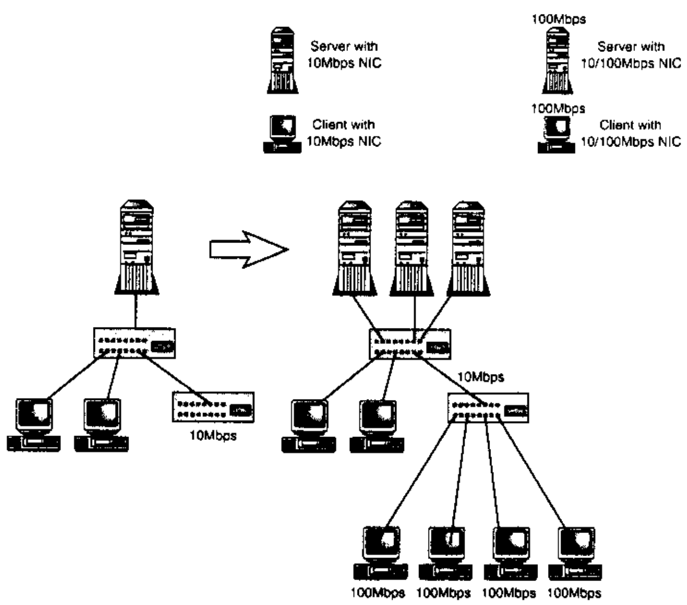

FIGURE 9.7 *A 10Mbps shared network before (left) and after (right) adding new systems with 10/100Mbps NICs. All clients and servers in this example are running at 10Mbps even though some have 10/10 NICs installed. This provides a network that is primed for a fast Ethernet upgrade.*

By adding 10/100Mbps NICs now, you are preventing an expensive and time-consuming future NIC upgrade. (Most network servers are so critical that they can't afford to be powered down. Desktops are usually too numerous to allow for a mass replacement of NICs once they are installed.) Lack of proper cabling to support 100BASE-T should not be a deterrent to installing 10/100Mbps NICs in new systems. Who knows where they may be moved to in the future. You can always upgrade cabling when the budget allows, and you don't want to have to upgrade all your desktops as well. Now that 10/100Mbps NICs are virtually the same price as 10-only NICs, 10/100Mbps should be the rule in all new systems.

Using 100BASE-TX Versus 100BASE-T4

The type of 10/100Mbps NIC to install depends primarily on the type of cabling used for the particular LAN. With Category 3 UTP, 100BASE-T4 is the only Fast Ethernet option. With Category 5 UTP, either 100BASE-TX or 100BASE-T4 deployment is possible. The best guideline to follow when choosing between TX and T4 is to determine which supports the majority of your cabling. Most network managers, however, are upgrading their cabling to Category 5 when they are upgrading their infrastructure. Combine this with the lack of 100BASE-T4 product availability, and it's no surprise that 100BASE-TX is the solution of choice in today's Fast Ethernet networks. In our opinion, you should avoid 100BASE-T4 whenever possible.

If you don't know what category your cabling is in or how many pairs are available, consult your cable contractor or try one of the many portable cable testers discussed in Chapter 6. Table 9.6 outlines the various TX and T4 options dictated by your cabling infrastructure. Pay special attention to Category 3 cabling with only two pairs available. Neither TX or T4 can run on this cabling so your best upgrade option is Switched 10Mbps Ethernet.

TABLE 9.6 CABLING OPTIONS

Cable Type	Connector Type	Number of Pairs Available for LAN	100BASE-T Signaling Scheme to Deploy
Category 3 or 4 UTP (voice grade)	RJ45	2	100BASE-T not deployable. Use Switched 10.
Category 3 or 4 UTP (voice grade)	RJ45	4	100BASE-T4
Category 5 UTP (data grade)	RJ45	2	100BASE-TX

continues

TABLE 9.6 CONTINUED

Cable Type	Connector Type	Number of Pairs Available for LAN	100BASE-T Signaling Scheme to Deploy
Category 5 UTP (data grade)	RJ45	4	100BASE-T4 or 100BASE-TX
Type 1 STP	DB9	2	100BASE-TX
Coaxial Cable	BNC	N/A	100BASE-T not deployable

Upgrading Servers to Fast Ethernet

Servers have a special place in most networks as they hold the data and information to which most people want access. This often makes them vulnerable to overloading, especially in a 10Mbps shared environment. Therefore, the first place to look when adding 10/100Mbps NICs is in your servers. Once you have installed the 10/100Mbps NICs, upgrading the line speed to 100Mbps is basically automatic. In addition to adding the basic 10/100Mbps capability to servers, step 2 includes some other considerations about which type of NIC to use.

Specialty NICs include intelligent NICs and multiport NICs. Intelligent NICs have an intelligent processor or subsystem that relieves the host CPU of many of its tasks. Intelligent NICs should be considered mainly in four situations:

- Multiple-segment Fast Ethernet servers
- Application servers where CPU power is at a premium
- PC-based routers such as Novell's Multi Protocol Router (MPR) software
- Server Fault Tolerant (SFT) III-compliant links, such as Novell's Mirrored Server Link (MSL)

Because the on-board CPU does much of the network processing, intelligent NICs usually result in great throughput at a low CPU utilization. Intel's EtherExpress PRO/100 Intelligent Web Server Adapter, shown in Figure 9.8, is one such NIC. Because of their added cost, however, intelligent NICs are not very beneficial outside these four niche areas.

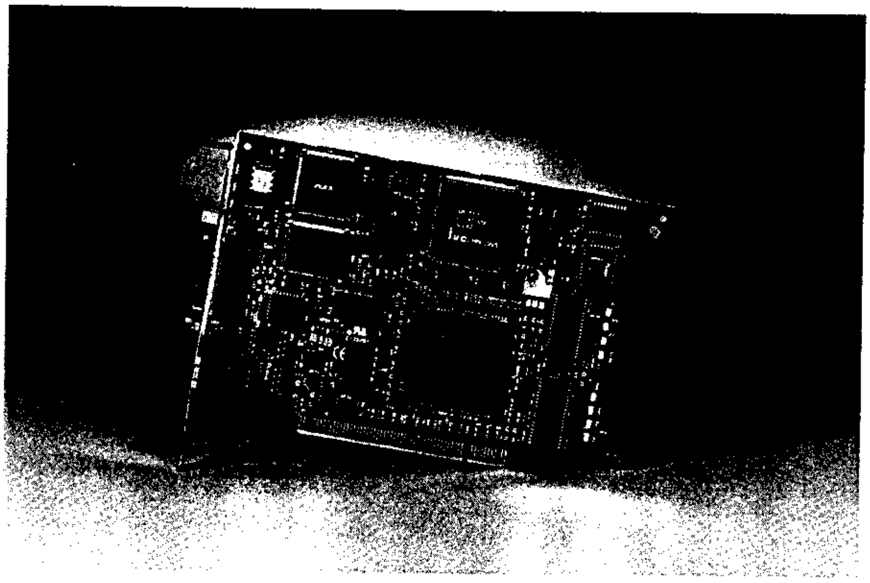

FIGURE 9.8 *Intel's EtherExpress Intelligent Web Server Adapter.*

Multiport NICs combine several NICs into one by providing up to four UTP ports on one card. Multiport NICs are useful in systems where expansion slots are scarce. Many PCI systems have only two or three PCI slots available for expansion cards. As server PCI slots start becoming more plentiful, multiport NICs will slowly become less popular.

Multiport NICs, or even multiple single-port NICs, can be aggregated to provide more bandwidth in and out of the server. As we've discussed, schemes like EtherChannel and 802.3ad port aggregation work well in these scenarios. Using one of these schemes, a server can have up to four Fast Ethernet NICs installed, delivering up to 800Mbps of data (400Mbps each way).

The Patch Panel Approach

Many large companies have network closets that house the connection from each office to a particular hub. This is typically accomplished with a patch panel (this concept is explored in Chapter 6). A patch panel allows maximum configurability of the network. For instance, the network administrator can rearchitect the network layout from the closet by switching the connections on the patch panel. A patch panel approach is very effective for converting from 10Mbps to 100Mbps. Suppose a network has several users with 10/100Mbps NICs operating at 10Mbps. If those systems are known, a quick restructuring of the patch panel wiring will condense these users onto one hub (later, we'll discuss how to automate this process with virtual LANs). When any given 10Mbps hub has all 10/100Mbps-enabled end stations attached, it is a prime candidate for step 3, adding a Fast Ethernet hub.

Deployment Step 3: Converting Workgroups to Fast Ethernet

The next logical step after installing 10/100Mbps NICs is to find a way to convert them from 10Mbps operation to 100Mbps operation. You can best do this by first converting workgroups from 10Mbps to 100Mbps. You can do this by using 10/100Mbps switches or 100Mbps repeaters, depending on the situation. In some cases, you will need to make trade-offs to merge the new Fast Ethernet LAN with the existing 10Mbps environment.

Four main types of workgroups are eligible for conversion to Fast Ethernet. Each requires a slightly different approach to achieve the best performance gain:

- *Existing workgroup*—10Mbps-only NICs; clients only (no local servers).
- *Existing workgroup with local servers 10Mbps-only NICs*—many clients and some local servers.
- *Newer workgroup 10/100Mbps NICs*—clients and/or local servers.
- *Newer power workgroup 10/100Mbps NICs*—clients and/or local servers.

If implemented correctly, each Fast Ethernet upgrade provides increased performance with minimal interruption to the network.

Existing Workgroups with no Local Servers

Replacing every NIC in your network is a seemingly insurmountable task. Therefore, the best Fast Ethernet workgroup solution for existing networks is one that allows you to leverage the installed base of 10BASE-T NICs. Most existing workgroups can achieve a significant benefit from the addition of a 10Mbps switch with 100Mbps uplinks. For existing workgroups consisting entirely of clients, a switch with many 10Mbps ports and a few 100Mbps ports is desirable (for connections to other workgroups and the backbone). Consider the leftmost workgroup in Figure 9.9, which contains only 10Mbps clients and a 10Mbps repeater.

In step 1, we saw how a 10Mbps switch could be deployed in this situation to enhance workgroup performance. A 10Mbps switch with 100Mbps uplinks can enhance performance even more and without the inconvenience of multiple full-duplex uplinks to the backbone. Figure 9.9 shows how a standard 10Mbps shared workgroup can be upgraded with a 10Mbps switch plus 100Mbps uplinks.

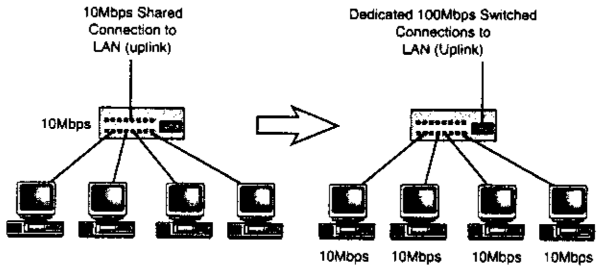

FIGURE 9.9 *Upgrading an existing workgroup to Fast Ethernet without replacing the 10Mbps NICs in the desktops.*

Now what do you do with the newly acquired 100Mbps uplink? Because Fast Ethernet won't be deployed in your backbone until step 4, you may have to live with a 10Mbps uplink for a little while. Performance won't approach 10 times current levels until the Fast Ethernet backbone connection is completed. A good way to approach this problem is by investing in a 10Mbps switch with an upgradable option for Fast Ethernet. One such switch is the Cisco Catalyst 2800, which can be bought as a 10Mbps-only switch with space for two Fast Ethernet modules. You can purchase the Fast Ethernet modules when the rest of your Fast Ethernet network—specifically, the backbone—is in place. Figure 9.10 shows this type of switch deployment.

Note

For an interesting practical example, have a look at `www.cisco.com/partner/786/4.html`*. This shows how the Hard Rock Cafe in Cleveland, Ohio, has implemented step 3.*

Note that this 10Mbps switch upgrade applies to existing workgroups that already have 10Mbps NICs installed, upgrading existing systems with 10/100Mbps NICs would not make sense. A workgroup with mostly ISA-based clients is a classic example of where to deploy 10Mbps switches and 100Mbps uplinks in this manner. Another good example is a workgroup with Category 3 UTP cabling.

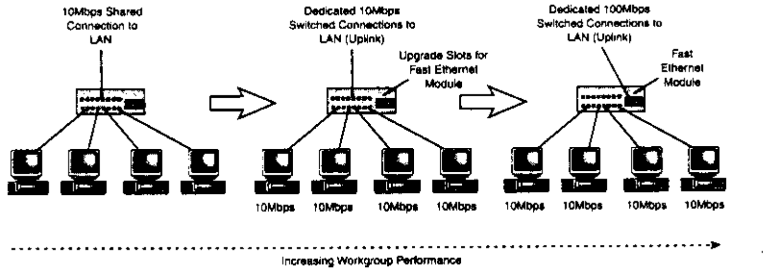

FIGURE 9.10 *Leveraging 10Mbps switches in a workgroup. On the left is the original 10Mbps shared workgroup. First it is upgraded with a 10Mbps switch (center). Finally, a Fast Ethernet uplink modules is added to the switch (right).*

Existing Workgroups with Local Servers

A slight variation on the previous scenario includes the addition of local servers in the workgroup. In some workgroups with local servers, the majority of the traffic (more than 80%) may be between those servers and clients in the workgroup. This is often referred to as a standalone or isolated workgroup. 10Mbps switches with more 100Mbps ports are preferable in these situations. Each local server can be connected to a 100Mbps pipe, and each client gets its own dedicated 10Mbps pipe. This is also referred to as personal Ethernet because each user has his or her own dedicated Ethernet line. Figure 9.11 shows this. Be careful when assessing the characteristics of a standalone workgroup. If many of the users are accessing the Internet through a centralized gateway, it is unlikely that more than 80% of the network traffic is local to that workgroup. Use a network analyzer or an RMON probe(discussed in Chapter 12, "Troubleshooting") if you are unsure.

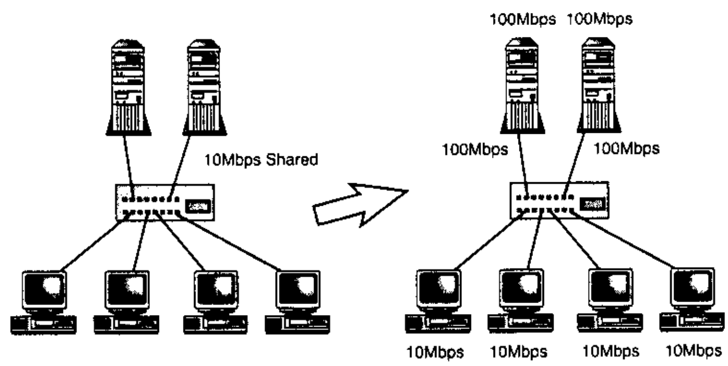

FIGURE 9.11 *A 10/100Mbps Mbps switch upgrade in a workgroup with local servers. Clients get dedicated 10Mbps. Servers get dedicated 100Mbps.*

Two types of 10/100Mbps switches are available to deploy in an isolated workgroup. The first and most obvious type is a switch with many 10Mbps ports and more than a few 100Mbps ports. A good example is the Bay Networks 350T, which has a total of 16 ports that can be configured for 10Mbps or 100Mbps. This type of switch will give excellent performance and configuring capability, but the expense may be a little high for broad workgroup deployment. Another way to achieve similar workgroup performance is to use architecture like that of the Cisco Catalyst 2820. The Catalyst 2820 has 24 switched 10Mbps ports and upgrade slots for one or more switched 100Mbps ports or eight 100Mbps shared ports. This allows several servers to share the same 100Mbps bandwidth and still provides each client with 10Mbps of bandwidth. The single switched 100Mbps port connects the workgroup to the LAN backbone. This architecture, diagrammed in Figure 9.12, is less expensive than switching 10/100Mbps on a per-port basis.

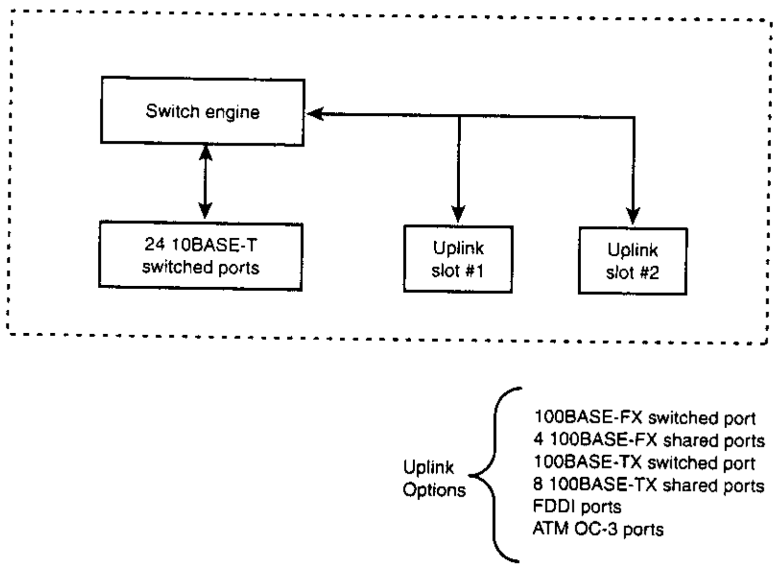

FIGURE 9.12 *Cisco's 2820 FastSwitch ES is a good example of a 10Mbps switch architecture with several uplink options.*

For the 10/100Mbps switch to be of any use in this type of environment, the servers must be outfitted with 10/100Mbps-Mbps NICs. This means that the servers must be capable of supporting 100Mbps data rates. If the local servers are ISA based, a 10/100Mbps-Mbps switch may not dramatically increase performance. A 10Mbps switch may be more appropriate. If the local servers do have high-speed bus types, such as PCI, EISA, or S-Bus, and 100Mbps-capable NICs, the performance increase can be much greater.

New Workgroups with no Local Servers

Up to this point, we have discussed upgrading existing workgroups where the clients are restricted to 10Mbps NICs. Next, we will discuss how to implement Fast Ethernet in new workgroups. (The term new workgroup implies that the systems in that workgroup were recently added to the network.) In step 2, we discussed how to make sure each new system is installed with a 10/100Mbps NIC. The combination of many new systems being added to the network and the inclusion of 10/100Mbps NICs makes upgrading to Fast Ethernet very easy and inexpensive. The most cost-effective way to upgrade this kind of workgroup is with a 100Mbps standalone or stackable hub. The users will instantly jump from sharing 10Mbps to sharing 100 Mbps, which provides a tenfold performance improvement. The following sections describe how you can do this.

Small Workgroups

A small workgroup can be classified as 20 users or fewer. Depending on how you plan to enlarge that small workgroup, you may elect to connect it with a 100Mbps standalone repeater or a 100Mbps stackable hub. As described in Chapter 8, "Network Components," the standalone repeater will be less expensive in the short term, but the stackable hub offers greater flexibility for future growth. Let's look at an example of how a new workgroup is installed using a 10/100Mbps-Mbps standalone repeater. In this example, a new group of users will be added to a small network. We have already seen how existing systems can be connected using a 10/100Mbps switch. Now those same systems can be connected using a 10/100Mbps-Mbps repeater. Part A of Figure 9.13 shows a small network with just a few 10Mbps repeaters, clients, and servers. Part B shows the new network after the addition of new users.

The newer clients share 100Mbps of bandwidth with the newly 100Mbps-enabled servers. The existing clients get 10Mbps dedicated connections. The number of users on the network has doubled, the available bandwidth has gone from 10Mbps to 100 Mbps, and no desktops were opened unnecessarily. The fact that the standalone repeater can operate each port at 10Mbps or 100Mbps makes it very flexible for this type of deployment. No matter what speed a client is running, it can be connected to the repeater. The cost of adding the 10/100 repeater is low, therefore providing cost-effective migration from 10Mbps to 100Mbps for this small workgroup.

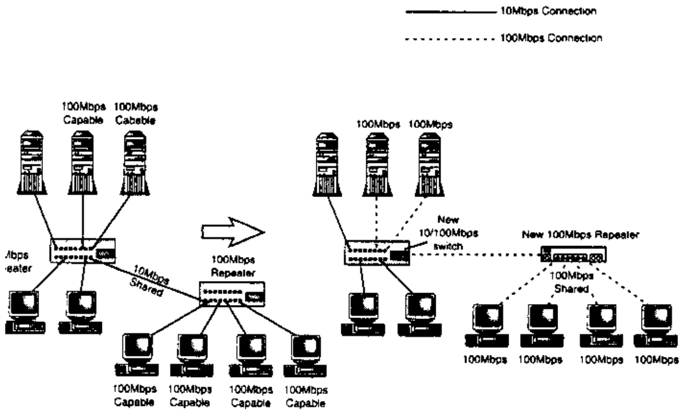

FIGURE 9.13 *Network before (a) and after (b) adding a new Fast Ethernet workgroup. Older clients with 10Mbps NICs are connected to a 10/100Mbps workgroup switch. New clients with 10/100Mbps NICs are connected to a 10/100Mbps repeater.*

Large Workgroups

When deploying new users in a large network, 10/100Mbps stackable hubs or chassis hubs are the right choice. Large networks are typically upgraded from the network closet, where new users are connected through patch panels to rack-mounted stackable hubs or chassis hubs. Figure 9.14 shows a typical large 10Mbps Ethernet installation of several racks of stackable hubs and a chassis hub.

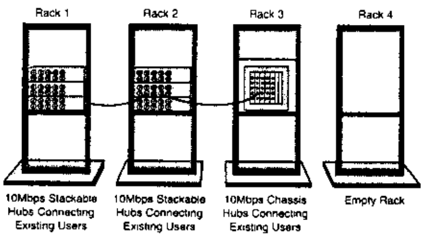

FIGURE 9.14 *A typical large 10Mbps wiring closet installation including stackable hubs and chassis hubs.*

You can add new users to this network by installing 10/100Mbps-Mbps stackable hubs and connecting the new users to them. You can connect the new users to the existing 10Mbps network via a 10/100Mbps switch module in the chassis. You could also connect additional new Fast Ethernet users via one of the ports on the 10/100Mbps stackable hub because they all run at 10 or 100Mbps speeds. Figure 9.15 shows the upgraded network. In this example, a 10/100Mbps switching module, like the Bay Networks System 5500 XXX, connects the 10Mbps and

100Mbps networks. You could also bridge 10 and 100Mbps by adding a 10/100Mbps Mbps-switching module to one of the stackable hubs. In either case, some of the network users are running at 100Mbps and some at 10Mbps.

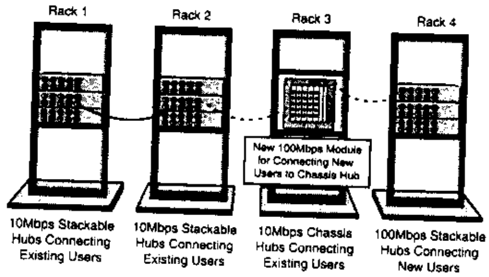

FIGURE 9.15 *Adding new Fast Ethernet users to a large 10Mbps wiring closet.*

New Power Workgroups with Local Servers

A small but growing percentage of new workgroups fall into the classification of power workgroups. This classification indicates that the workgroup puts so much local traffic on the network that even a 100Mbps repeater cannot handle the workload. A good way to recognize a power workgroup is to measure the amount of data that workgroup transmits on the wire, also called the workgroup network utilization. If the workgroup network utilization is above 30% of the total wire bandwidth available and over 80% of that traffic is local (meaning it is destined for another station or server within that workgroup), it may qualify as a power workgroup. Some examples of power workgroups are CAD workstation clusters, groups of multimedia systems, and desktop publishing centers.

You can connect clients and servers in power workgroups with per-port 10/100Mbps switches, such as the Bay Networks 28200. Each client in the workgroup is given a dedicated 100Mbps pipe. Servers are also given 100Mbps connections, or in some cases, aggregated 100BASE-TX connections. The connection to the backbone can also be multiple, aggregated 100Mbps connections, or a single Gigabit Ethernet connection. Therefore, look for 10/100Mbps switches that support Gbps Ethernet unlinks. Figure 9.16 shows a Fast Ethernet power workgroup deployed with a 100Mbps switch.

The Low-Cost 10Mbps/100Mbps Desktop Switch Factor

The rapid decrease in the cost of 10/100Mbps per-port switches has made 100Mbps switching to the desktop worth considering. Today for less than $150 per port—this price will have undoubtedly changed even before this book goes to print—you can purchase a 24-port 10/100Mbps switch from such big names as 3Com, Cisco, Intel, Bay Networks, and Hewlett-Packard. Although this price is far from the $30 to $50 per port cost of today's 10Mbps stackable hub, it is close enough to draw attention. At some point, the network architects will have to ask themselves if their network is ready for 100Mbps desktop switching.

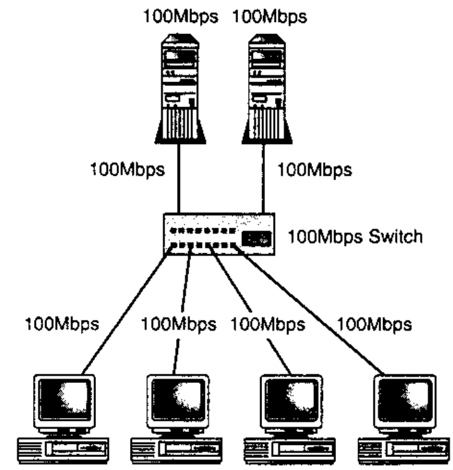

FIGURE 9.16 *A power workgroup connected via a Fast Ethernet switch. Each client and server gets 100Mbps of dedicated bandwidth. Half- or full-duplex modes can be supported.*

Deciding Between 10Mbps Switches and 100 Hubs

You may be thinking, "But I don't have any power workgroups on my network. My network utilization is about 30% and I only need a small performance boost." In this case, now may not be the time for Fast Ethernet deployment in your network. Step 1 (10Mbps switching) may be the answer to your current LAN bandwidth needs. In step 3, 100Mbps workgroups are advocated. Which gives better performance in a given situation? Vicious industry discussions have centered around which is better, shared 100Mbps or dedicated 10 Mbps. The simple answer is that both solutions measure about the same in terms of performance. 10Mbps switches offer a collision-free environment, whereas 100Mbps hubs, although collision-prone, offer 100Mbps transfer rates all the way to the desktop. In the end, both solutions are good for about 95% of today's workgroups.

Our advice is this: If at all possible, run a desktop at 100Mbps. If this means 100Mbps stackable hubs are your choice, fine. If this means your budget allows for 10Mbps/100Mbps switched connections to the desktop (power workgroup), even better. New PCs will eventually outgrow 10Mbps-only switched connections.

Still, it helps to know when Fast Ethernet to the desktop is really necessary. Here are a few hints on how to determine whether 10Mbps switching

continues

continued

is enough power for a given workgroup.

If the following are true, 10Mbps switching may be the final step needed in your high-speed deploment strategy:

- The total number of users in a given workgroup is under 40.
- The workgroup's primary network usage is basic file transfers and network printing.
- The current network utilization in the workgroup is under 30% (less than 3Mbps).
- End users do not complain about network performance or response time.

Also, as discussed previously, your desktop PC architecture or cabling infrastructure may dictate that 10Mbps switching is the highest performance solution available.

No matter what you deploy, though, always be on the lookout for increasing network traffic and plan your 10Mbps architecture in such a way that you can easily migrate to Fast Ethernet (step 3 and beyond) if the need arises. This includes deploying Fast Ethernet-ready 10Mbps switches and 10Mbps/100Mbps NICs in all new desktops and servers.

Luckily (or unluckily depending on your risk-taking tendencies), it will take some time for the industry to drive the price down. This is your chance to start evaluating 10/100Mbps desktop switches and start determining how you will connect 10/100Mbps switched workgroups with Gigabit Ethernet backbones (step 6). Worried? Well, read on.

Future Workgroup Considerations

As workgroups grow, changes will have to be made to the architecture of those workgroups. Most importantly, the location and type of local server should be considered. This will determine how much network traffic stays local to the workgroup and how much leaves the workgroup, destined for the backbone or WAN. Consider the following issues when planning future workgroup deployment:

- *Multimedia support*—Multimedia traffic differs from regular network traffic in that it is constant and not bursty in nature. A television-quality, full-screen MPEG video will consume 1.5 to 3.0Mbps of network bandwidth. If four or five such transmissions are occurring simultaneously in a workgroup, you can easily see how 10Mbps won't be able to handle the load. Transferring voice and data together also causes problems because long latencies on voice packets can break up a live discussion. In the future, switches will support features such as RSVP bandwidth reservation protocol, 802.1p frame prioritization, and virtual LANs (VLANs) to overcome these issues, but this is still a few years away. In the meantime, use the best weapon you have to approach these applications: bandwidth. You'd be surprised what 100Mbps of dedicated bandwidth to a desktop will allow you to do.

- *Layer 3 filtering in the workgroup*—Layer 3 support is recommended primarily in the backbone in this book, but in the future, Layer 3 capabilities could be driven down into the workgroup. In theory, the earlier a packet is identified and filtered, the better it can traverse the network. Today, Layer 3 capabilities in a workgroup switch will significantly add to the cost. In the future, this capability may become standard.
- *Gigabit Ethernet and ATM support*—Many 10 or 100Mbps Ethernet switches already support Gigabit Ethernet and 155Mbps ATM uplinks. This is important if Gigabit Ethernet or ATM are being considered for the backbone. You should remember, however, that by mixing ATM and Ethernet, ATM becomes nothing more than a fat pipe. The QoS features inherent in ATM must be implemented with RSVP and 802.1Q/p tagging in an Ethernet workgroup. Therefore, consider ATM very carefully before combining it with an existing Ethernet network. You're better off using Gigabit Ethernet to uplink a 100Mbps workgroup switch.

Workgroup Options

The flexibility of switches that automatically configure to 10Mbps or 100Mbps on a per-port basis will help in the deployment of high bandwidth to your existing client and server base. These switches are jacks-of-all-trades because they can provide 100Mbps dedicated connections for backbones, repeaters, servers, and power users or 10Mbps dedicated connections for older clients and 10Mbps repeaters. Also, as described in Chapter 11, standard SNMP applications should be able to manage workgroup switches.

As an alternative to switches, 100Mbps stackable hubs are a cost-effective way to deploy Fast Ethernet to users with new systems and 10/100Mbps NICs. Table 9.7 illustrates the types of workgroups and the appropriate Fast Ethernet workgroup solution.

TABLE 9.7 WORKGROUP SWITCH DEPLOYMENT RECOMMENDATIONS FOR FOUR BASIC TYPES OF WORKGROUPS

Workgroup Type	Defining Characteristics	Recommended Fast Ethernet Upgrade Path (and Examples)
Existing workgroup	All nodes are clients	10Mbps workgroup switch with 100Mbps uplinks
Existing workgroup with local servers	Clients are legacy systems with 10Mbps-only NICs (usually ISA)	10/100Mbps workgroup switch
	Some local servers may be in the workgroup	
New workgroup	Mostly client nodes	10/100Mbps standalone repeater for small workgroups
	Local servers may be attached to the workgroup	10/100Mbps stackable hub for large workgroups
	Clients are new systems with 10/100Mbps NICs	10/100Mbps workgroup switches for large workgroups when budget allows
	Servers are new systems with 10/100Mbps NICs	
New power workgroup	Clients and servers are new high-end systems or workstations	10/100Mbps Workgroup Switch
	Clients and servers have 10/100Mbps NICs	
	Workgroup traffic level is above 30%	
	Most traffic is local (more than 80%)	

Deployment Step 4: Switched and Fast Ethernet in the Backbone

The next step after implementing Fast Ethernet in the workgroup is either to create a new backbone or convert your existing backbone to Fast Ethernet. First, we will discuss what defines a backbone and what types of backbones exist. Then we will discuss how to deploy Switched and Fast Ethernet in those backbones. If you have implemented the previous steps correctly, the backbone conversion should be straightforward. Also, by upgrading the backbone, your entire local area network will have increased performance instead of just the workgroup. Later, in step 6, we'll examine where Gigabit Ethernet fits in these backbones.

You can upgrade many different backbone implementations to Fast Ethernet from 10Mbps Ethernet. Two of the more common types are distributed backbones and collapsed backbones.

A distributed backbone is one that couples major sub-LANs via a chaining technique. A *sub-LAN* is defined as a floor, site, or other physical collection of workgroups. For instance, in a 10-story building, the LAN may incorporate all ten floors, plus the basement where the servers are kept. Each floor of the building is a sub-LAN, and each sub-LAN consists of several workgroups. As usual, the workgroups are concentrations of clients and local servers. Figure 9.17 shows a diagram of a distributed backbone. Note the resemblance to an actual human backbone, which is where the name originates. Distributed backbones include FDDI rings, Token Rings, and some switched Ethernet backbones.

Collapsed backbones are typically deployed when delays through the various switches or routers in a distributed backbone become too great. This occurs by deploying one high-end router in the basement and connecting each major sub-LAN to it directly. In a collapsed backbone, a packet from one sub-LAN must go through only the high-end router to reach any other sub-LAN. Collapsed backbones are also employed with FDDI Rings, Token Rings, 10Mbps switches, and high-end routers. A collapsed backbone is also shown in Figure 9.17.

The Location of Servers in a Backbone

Before we delve into the intricacies of distributed and collapsed backbones, we must spend a moment discussing server location. As discussed in the last chapter, servers represent information and clients represent users who create or use that information. Servers are located everywhere: in the workgroup (local), on the backbone, or across the wide area network. Because servers are included in almost every network transaction, you should understand how server location affects network traffic before embarking on a full-scale backbone upgrade.

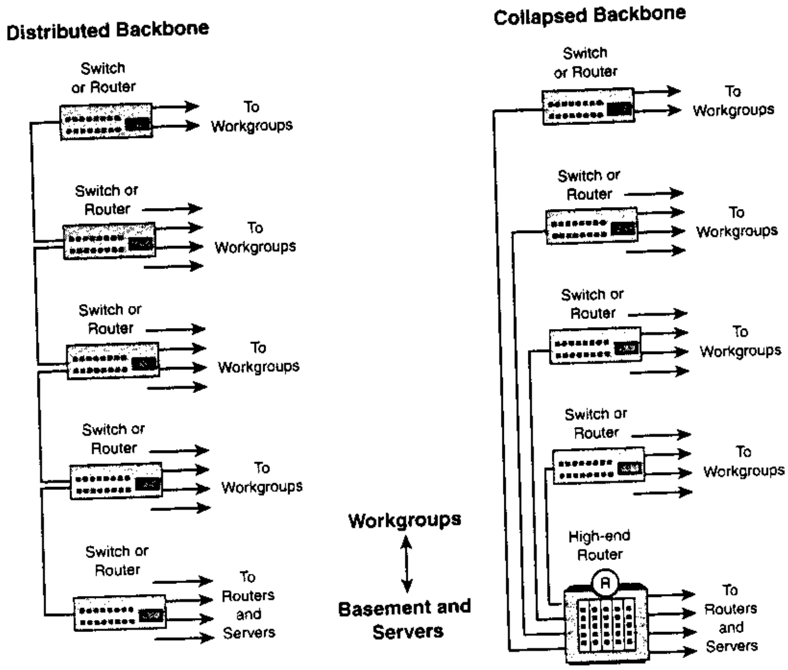

FIGURE 9.17 *The two types of backbones: distributed and collapsed.*

Consider Figure 9.18. Server A is a local file server. It contains spreadsheets, text documents, and other files that are of interest mainly to the local workgroup. Server B is actually a print server. Print servers, like Intel's NetPort, allow clients to print to many different printers. The clients in the workgroup will access server A and server B in almost a random fashion, causing the network traffic to burst from near zero to near 100 percent for short periods of time. This traffic, however, will not leave the workgroup.

Server C is an email server located on the network backbone. Access to the email server is typically periodic, with heavy access early in the morning and around lunchtime. Clients who access the email server do so across the backbone, so this traffic is not considered local. Server D is an Internet server that allows access to the Internet and WWW. Clients accessing this server will be asking it to be an Internet proxy. In other words, to keep the LAN secure, this special server will make requests of the Internet on a client's behalf. Traffic to and from the Internet server is typically bursty, or irregular, and can vary greatly in volume.

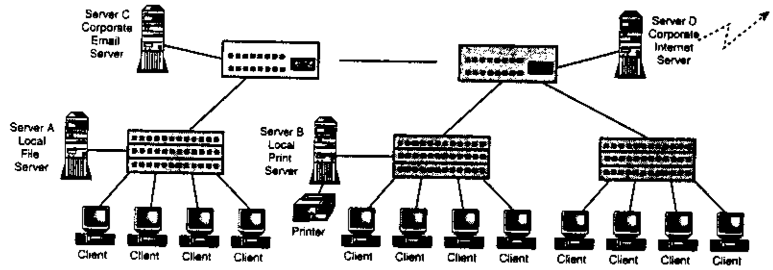

FIGURE 9.18 *Location of server resources drives the flow of network traffic.*

When traffic from all four servers is mixed, one can start to understand the traffic patterns that occur on today's networks. If servers A and B are used more often than C and D, the 80/20 rule applies, and the backbone only needs to support 20% of the traffic generated by each workgroup. If email and Internet access dominate the use of this network, the backbone must handle up to 100% of traffic generated by each workgroup. These two scenarios are completely different and require varying degrees of backbone power.

We intended this exercise to show what server placement can do to network traffic, but it illustrates another point as well. Unless you can accurately determine the location and type of your *future servers,* you can't really predict the requirements of your *future network.* For this reason we recommend overdesign of your network backbone to handle the possibility of most workgroup traffic propagating to the backbone. It may be overkill for today, but it will make your network more than ready for tomorrow.

Distributed Backbones

The distributed backbone deployment of today can incorporate a wide variety of solutions. On the performance scale, it can range anywhere from 10BASE-T to 1000BASE-T. In a distributed backbone, a high-end router is not necessary because each individual sub-LAN is connected via a particular switch or router. Therefore, the cost of a distributed backbone is usually lower than that of a collapsed backbone. The performance of a distributed backbone may be worse, however, due to high transmission latencies. In the example from Figure 9.17, a packet originating from the basement must pass through four backbone switches before reaching a destination on the top floor. In an application where latency matters, such as video-conferencing, this large latency can ruin performance.

This section will focus on how to upgrade an existing 10Mbps distributed backbone to Fast Ethernet and how to prepare for future backbone additions. Figure 9.19 shows how a 10Mbps distributed backbone would progress through steps 1, 2, 3, and finally step 4 of this chapter. By implementing Fast Ethernet in the backbone, step 4 marks the deployment of Fast Ethernet throughout the LAN.

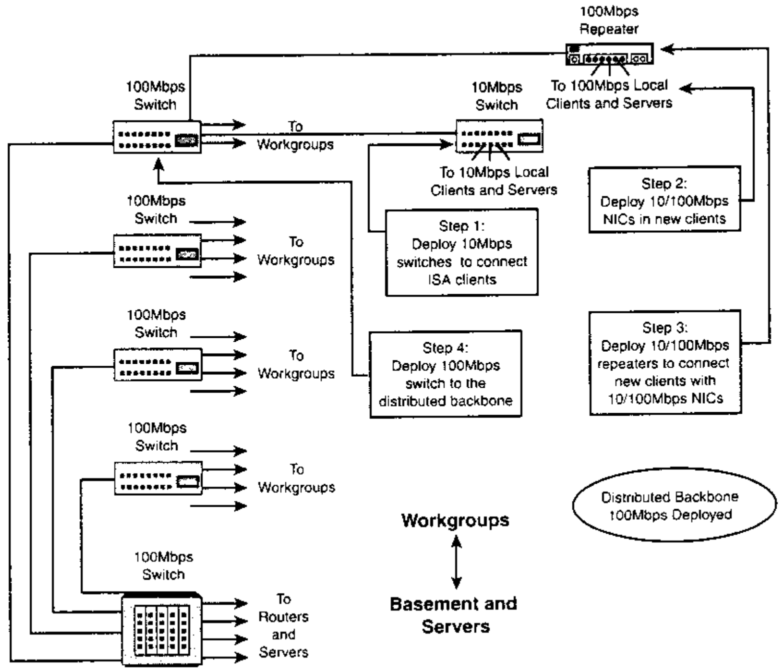

FIGURE 9.19 *Distributed backbone after Fast Ethernet deployment. Note that each step of Switched and Fast Ethernet deployment is shown.*

To complete the transition to Fast Ethernet, 100BASE-T switches have been added to the backbone in Figure 9.19. This nullifies any 100BASE-T network diameter constraints because switches don't suffer from collision domain restrictions. Therefore, a 100BASE-T distributed backbone of any size can be built with 100BASE-T switches. 100BASE-T switches enabled with a port aggregation scheme can provide up to 400Mbps dedicated bandwidth in and out of each backbone switch. The performance improvement over 10Mbps shared can therefore be anywhere from 10 to about 100 times, depending on the loading of the backbone.

The Bay Networks 450T 100BASE-TX switch provides this type of functionality. Each of the 24 ports on the 450T will auto-configure for 100BASE-TX or 10BASE-T operation. In the example from Figure 9.19, one workgroup hub is a 10/100Mbps stackable hub and should be connected at 100Mbps on the backbone switch. The other is a 10Mbps switch and should be connected to a 10Mbps port on the backbone switch (or a 10/100Mbps switch could be used in its place). The Bay 450T also provides modules for Gigabit Ethernet ports for connecting backbone switches. This feature alone can provide a tenfold bandwidth improvement in the backbone.

In this example, UTP wiring is used because the switch-to-switch connections of a multistory building will unlikely be greater than 100m. This doesn't hold true for a physically spread-out distributed backbone. Consider a multiple-building site where each building is 500m from the next. To connect these buildings to the same 100BASE-T backbone, 100BASE-FX fiber connections must be made. 100BASE-FX can transmit over 2km of multiple-mode fiber.

A drawback to deploying a 100BASE-T distributed backbone is the long latencies encountered in hopping from switch to switch. In cases where you'd like to deploy real-time LAN products, such as video conferencing, a 100BASE-T collapsed backbone may be the better choice. The advantage of a distributed backbone, however, is that you don't need a central, high-end router or switch in the basement.

Collapsed Backbones

A collapsed backbone is not much different from a distributed backbone. In fact, the sub-LANs and workgroups are structured the same in either case. The major difference is the types of switches or routers installed in the backbone and the way they are connected. For instance, there is no such thing as a shared-media collapsed backbone. Collapsed backbones are switched or routed by definition.

The issues with deploying Fast Ethernet into a collapsed backbone are of two varieties. First, determine what kind of router you have in your basement. Can you simply upgrade it with a Fast Ethernet module or will you need to purchase a whole new chassis-based router? Second, determine how your various sub-LANs are connected to the router and with what media type. It is highly possible that you can install Fast Ethernet in your collapsed backbone with no rewiring at all.

When upgrading a collapsed backbone to Fast Ethernet, you must replace the basement router and some sub-LAN switches together. Not all sub-LANs need to be upgraded to Fast Ethernet at the same time, however. As soon as the new Fast Ethernet-enabled router is in place, you can replace each sub-LAN switch as the need arises. Figure 9.20 shows a typical 10Mbps collapsed backbone. Figure 9.21 indicates how you can convert part of the backbone to Fast Ethernet by adding one module to the basement router and one 100Mbps switch to the sub-LAN.

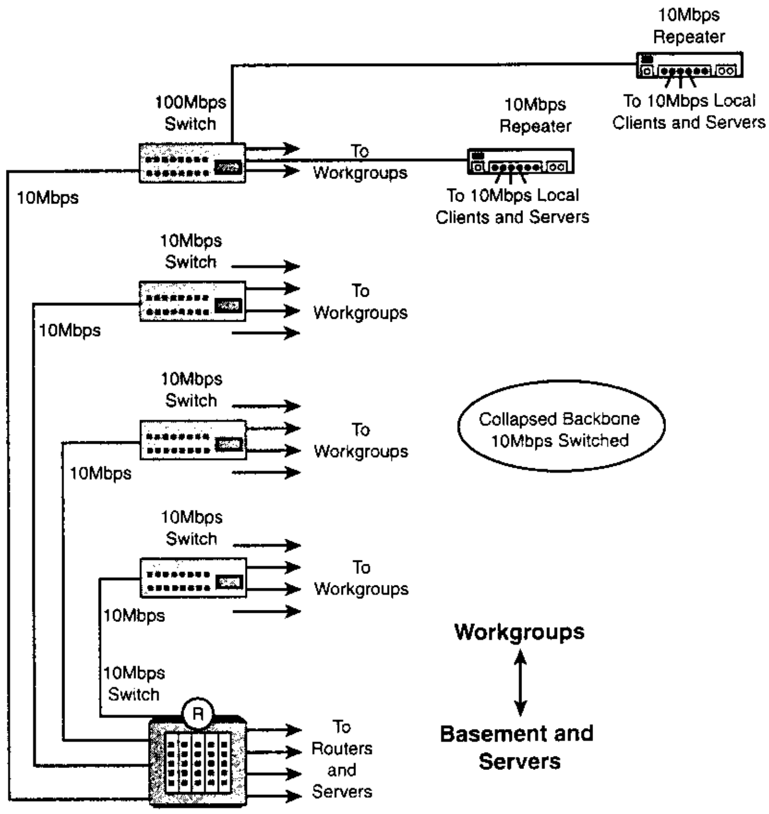

FIGURE 9.20 *A typical 10Mbps collapsed backbone. The high-end chassis router in the basement has many 10Mbps modules.*

By upgrading the basement router with Fast Ethernet modules, you can deploy 100Mbps collapsed backbones one step at a time. Of course, it is critical that the router support Fast Ethernet modules. Three such routers are the Cisco 7500 router, the Bay Networks Backbone Node router, and the 3Com NetBuilder II or CoreBuilder router. If you cannot upgrade the router at the center of your collapsed backbone to Fast Ethernet, you will either have to purchase a new router (very costly) or only deploy Fast Ethernet in the workgroup. Whether you use a new router module or a brand-new router, look for it to support fast and efficient routing between Ethernet and Fast Ethernet segments. In general, whatever features you look for in an Ethernet router, look for those in a Fast Ethernet router as well.

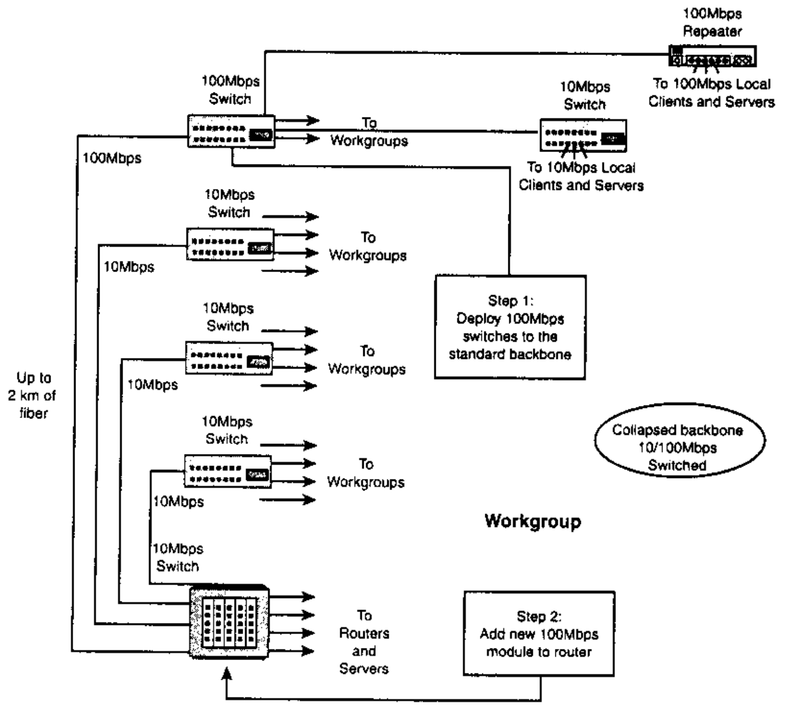

FIGURE 9.21 *The backbone after adding a new Fast Ethernet module to the router and a new Fast Ethernet switch in one sub-LAN. One sub-LAN at a time can be upgraded to Fast Ethernet in a collapsed backbone.*

Collapsed backbones are currently preferred over distributed backbones because of their performance benefits and lower packet latencies. However, a collapsed backbone has typically required a sophisticated piece of routing equipment, such as the Cisco 7500, which may be more expensive (we'll discuss how to get around this in step 5). Another problem that often arises with collapsed backbones is that, because of large physical distances between sub-LANs, the sub-LAN connections commonly require cable lengths greater than 100m. You cannot use UTP in these situations. FDDI has been a common solution until now, but 100BASE-FX provides a more cost-effective alternative.

For instance, if the example in Figure 9.22 requires a sub-LAN to be connected more than 100m away, you can insert a 100BASE-FX module in the basement router and connect it to a 2km fiber run. Be sure to consider that the sub-LAN switch (at the top of Figure 9.22) must also be capable of connecting to 100BASE-FX.

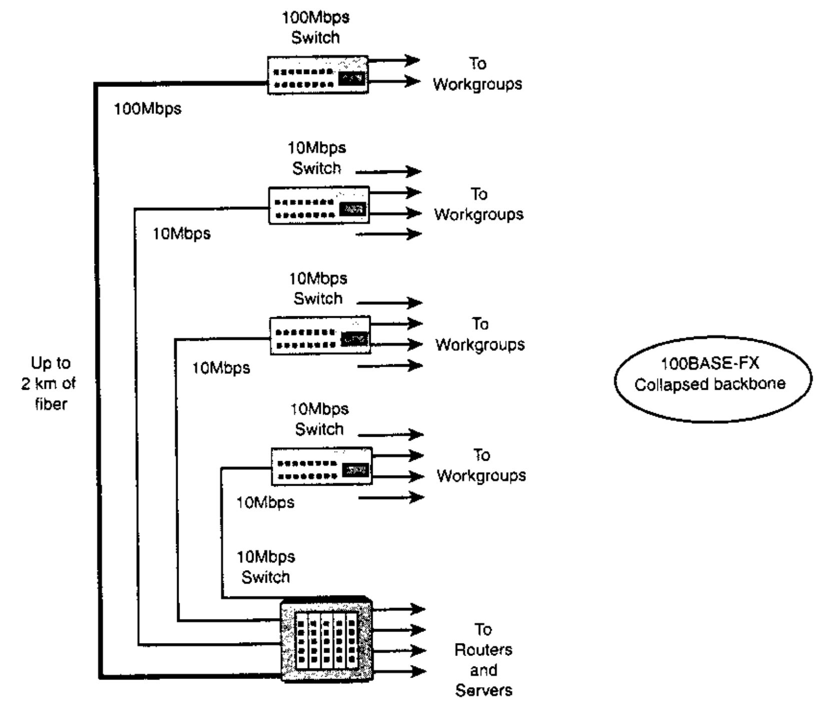

FIGURE 9.22 *Using 100BASE-FX in a collapsed backbone.*

Port Aggregation and Gigabit Ethernet Uplinks

We don't add Gigabit Ethernet to the backbone until step 6, but it's important to be prepared for it by using 100Mbps backbone switches with slots for Gbps Ethernet uplinks. The previously cited examples use Cisco's Fast EtherChannel port aggregation scheme to interconnect switches. This provides an economical way of increasing bandwidth between switches up to 400Mbps in each direction (in and out of the switch). Gigabit Ethernet provides a true high-speed uplink, however, which is capable of aggregating traffic from a 100Mbps backbone switch. In this step, when Fast Ethernet is the interswitch connection of choice, use a port aggregation scheme like Fast EtherChannel to connect switches.

Server Farms

Recall from step 2 that *server farms* are classically defined as a cluster of servers that are connected directly to the backbone. Server farms can also be of many shapes and sizes, some of the more popular versions being file servers, Internet servers, and application servers. File server farms are typically accessed randomly and provide file sharing and directory services to anyone on the LAN. File servers are subject to high-bandwidth spikes and bursty traffic, but their overall average utilization usually remains low. Internet servers are also accessed randomly, but they perform additional duties such as caching Internet data and protecting the LAN from Internet viruses and eavesdroppers. Application server farms are under more constant bandwidth demand because they are providing email, database, and multimedia services. The average traffic generated by application server farms is typically large.

File servers require high-speed disk subsystems, but little else. For instance, a Pentium processor-based file server will likely perform much like a Pentium II processor-based file server if similar hard drive subsystems are used in each. File servers also tend to be used randomly, with high bursts of network utilization when users open or close a file. For this reason, file server farms show large network utilization spikes but remain low in average bandwidth. As we have already seen, repeaters provide an excellent solution for this type of network. The trick is to determine how to distribute the shared bandwidth between servers. Consider the following, where 16 file servers are part of a server farm on a 10Mbps backbone. By measuring the amount of traffic generated by each server, one can determine how to upgrade it. Assume the servers put the following loads on the 10Mbps repeater to which they are connected:

File Server	Load
1	1.0Mbps
2	0.2Mbps
3	0.2Mbps
4	0.4Mbps
5	1.6Mbps

continues

File Server	Load
6	0.4Mbps
7	0.2Mbps
8	0.2Mbps
9	1.0Mbps
10	0.4Mbps
11	0.2Mbps
12	0.4Mbps
13	0.2Mbps
14	0.2Mbps
15	0.2Mbps
16	0.2Mbps
Total	**7.0Mbps**

When upgrading this particular file server farm to Fast Ethernet, you have the opportunity to optimize its configuration. First, replace the 10Mbps NICs in each server with 10/100Mbps NICs (if you haven't already done so as a result of step 2). Next, consider that the best way to segment these file servers is to put the few high traffic servers together, separating them from the others. Each group is connected with a 100Mbps repeater and attached to the network via a 100Mbps switched connection. This results in two server farms: one with servers 1, 5, and 9 and the other with the rest of the servers. Each server farm shares 100Mbps of bandwidth through a Fast Ethernet repeater. An even better solution would be to provide individual 100Mbps switched ports to servers 1, 5, and 9.

80-20 Rule

Internet and intranet servers are typically placed centrally and often connected to the backbone. (When an internet server serves information to internal users only, it is sometimes referred to as an *intranet server*.) These servers have almost single-handedly reversed the 80/20 rule that was common in most workgroups. That is, before the Internet, 80% of workgroup traffic stayed local to that workgroup. Now that so many corporate employees are accessing the Internet, a large portion of the network traffic they create is destined to or from the Internet. This makes Internet servers a constant source of action on a busy network and often causes their location to be on the high-speed backbone, rather than the local workgroup.

Application servers are also implemented as server farms. In this case, the bandwidth demands are usually too high and constant to use a repeater. A

100Mbps backbone switch may be the best solution for an application server farm. The deployment is very similar to that of the file server farm, except each server gets its own dedicated 100Mbps link. This way, when multiple users are accessing a database on an application server, it can respond without regard for bandwidth constraints from other servers. This type of server farm deployment also allows for multiple servers to be accessed simultaneously. Unlike file servers, an application server's performance gets much better as the quality of the NICs, processors, expansion bus speeds, and amount of system memory improves. Pentium processors, the PCI bus, and intelligent Fast Ethernet NICs may be worthwhile investments for application servers.

In addition, application servers (and to some extent file servers) are often critical parts of the network. If users cannot gain access to the information on these servers, productivity will drop, and users will become frustrated. For this reason, it is a good idea to consider redundancy when designing server farms. One of the best ways to provide redundancy is to populate each server with at least two NICs, connecting each of them to a separate 100Mbps switch. In this way, if the NIC, cable, or even the switch goes down, the server will remain connected to the network.

Managing Backbone Switches

Backbone switches are more instrumental to network uptime than workgroup devices. Therefore, you should pay more attention to managing these backbone switches. In a workgroup switch, RMON Groups 1, 2, 3, and 9 are recommended. In a backbone switch, all nine groups of RMON should be supported. Because this is an expensive option, look for switch vendors to provide add-in modules that support varying degrees of RMON support or look for roving RMON support (RMON groups 1 through 9 supported on all ports, but not all ports simultaneously). RMON2 support can also be beneficial where the benefits outweigh the costs of the RMON2 hardware and software. Chapter 11 goes into a more detailed discussion of switch management with RMON and RMON2.

Backbone Options

Deployment of a switched Fast Ethernet backbone will allow your entire LAN to communicate at 100Mbps. By deploying Fast Ethernet in the backbone, the lion's share of the upgrade from 10BASE-T to 100BASE-T will be complete. Access to both local and backbone servers will be at 100Mbps, and the performance gain should be evident in higher network traffic, faster network response time, and lower overall network utilization. Table 9.8 summarizes the important points of backbone deployment in each of these key areas.

TABLE 9.8 BACKBONE DEPLOYMENT ISSUES

Type of Backbone	Fast Ethernet Upgrade	Issues
Distributed backbone	Multiple 100BASE-T switches	A 100BASE-T switch is needed for each sub-LAN (floor, site, and so on) to be upgraded
		Distributed backbone can be implemented a piece at a time
		Packet latencies may be high
		Use 100BASE-FX for multiple building backbones
		Less expensive than upgrading a collapsed backbone
Collapsed backbone	Multiple 100BASE-T switches and a 100BASE-T router module for basement router	A 100BASE-T switch is needed for each sub-LAN (floor, site, and so on) to be upgraded
		A 100BASE-T router module is needed for basement router
		If basement router does not support a 100BASE-T module, a new router may have to be considered
		Collapsed backbone can be implemented a piece at a time
		Packet latencies should be low
		Use 100BASE-FX for multiple building backbones
		More expensive than upgrading a distributed backbone
Backbone Server Farm	100BASE-T repeater or 100BASE-T switch	Upgrade all servers to 10Mbps/100Mbps NICs
		Connect server directly to 100Mbps switched port whenever possible
		When using repeaters, use one for high-usage file servers and another for the rest of the file servers

Deployment Step 5: Routers, Layer 3 Switching, and VLANs in the Backbone

With 100BASE-T implemented in the backbone, the foundation for high-speed local area networking has been laid. So far, the discussion has been concerned with increasing the network bandwidth, but we haven't tackled the issue of network traffic flow. Leaving the network in this state is like adding larger pipes to an outdated water system but not replacing the old valves and filters. The "valves" of a LAN are the routers, Layer 3 switches and VLANs in the network. Routers include WAN (T1, T3, ISDN) and LAN (Token Ring, FDDI) varieties, Layer 3 switches are simplified routers embedded in switches, and VLANs represent ways of segmenting a network based on certain properties. In this section, we'll discuss the routers and Layer 3 switches in the backbone and glance at various VLAN implementations.

Fast Ethernet Routers

In Chapter 8, we discussed three areas where routers are useful: LAN segmentation, connecting disparate LANs, and connecting to the WAN. In this section, we'll examine how each of these applications changes when upgrading to Switched and Fast Ethernet.

Fast Ethernet Segmentation Routing

If your current network is based on a collapsed backbone and a chassis-based router is its center, upgrading to Fast Ethernet is much like what we discussed in Chapter 8. In this application, routers filter broadcast traffic, keeping it only on the subnet where it originated.

A new Fast Ethernet routing module, such as the Cisco Route Switch Processor for the Series 7500 router, would do the trick nicely in this situation.

Of course, now that Layer 3 switches are available, there are sometimes less expensive ways to do this. For now, use the Table 9.9 to determine whether you need a new Fast Ethernet routing module or whether you should consider Layer 3 switches.

TABLE 9.9 WHEN TO USE ROUTERS VERSUS LAYER 3 SWITCHES IN THE BACKBONE

Criterion	Use Fast Ethernet Routing Module	Use Layer 3 Switches
Protocol support	IP, IPX, DecNet, and so on	IP-only (some IPX)
Router connects to	Ethernet, Token Ring, FDDI, ATM, or other technology	Ethernet only

continues

TABLE 9.9 CONTINUED

Criterion	Use Fast Ethernet Routing Module	Use Layer 3 Switches
Router connects to WAN	Chassis-based router also has WAN routing modules	WAN connectivity is through a separate device, such as a standalone router

We discuss the use of Layer 3 switches as segmentation routers more in the following sections.

Standalone Routers

When connecting to the WAN, standalone routers are often used. This is because chassis-based WAN connectivity is usually very expensive. In contrast, standalone routers cost less but are less flexible. Even so, some standalone routers can be upgraded in performance via Fast Ethernet router modules, but only if they have slots for new modules. A good example of this kind of router is the 3Com NetBuilder III. Figure 9.23 shows the basement of a large networked site. The collapsed backbone router sits in the basement receiving many packets of data and concentrating on forwarding them. This router does not have time to filter and route packets to the WAN, so a standalone router is used for WAN access. Many networks use separate standalone routers in this configuration.

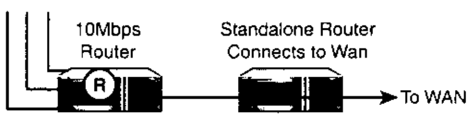

FIGURE 9.23 *Standalone routers may need to be replaced with fast Ethernet versions because they are not inherently upgrade capable.*

Another way to upgrade this network is with new, low-cost, Fast Ethernet access routers. Intel's Express 8100 Access Router has a 10/100Mbps connection built right into the box. Sometimes a new standalone router will also come with other new features like RSVP bandwidth reservation, support for Internet virtual private networking, and encryption software. If WAN access is the primary need, standalone routers are often the best way to connect the WAN to your Fast Ethernet LAN.

PC-Based Routing

An even more cost-effective approach to LAN segmentation or WAN connectivity is PC-based routing. Many networks employ software or PC-based routing because it usually costs less than upgrading the modules in existing routers or adding standalone routers. In addition, you can specify Fast Ethernet and other network types by the type of NIC used, allowing an infinite number of configurations. Figure 9.24 shows an example of a software-based router.

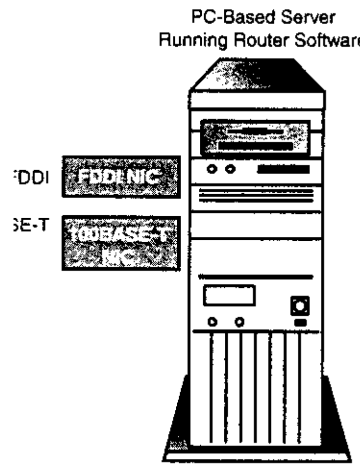

FIGURE 9.24 *Fast Ethernet PC-based routing (FDDI to Fast Ethernet). PC-based routing offers an economical alternative to expensive chassis-based routers, but the performance often leaves a lot to be desired.*

PC routing is limited to the performance of the PC on which it runs. As discussed in Chapter 8, quite a bit of PC horsepower is required to handle multiple Fast Ethernet connections. Therefore, PC-based routing should be limited only to networks where budget restrictions are the most stringent.

Layer 3 Switches in the Backbone

If your network doesn't have a backbone already, you should seriously consider Layer 3 backbone switches. Table 9.9 discusses the criteria for when Layer 3 switches can be considered. Basically, you should use them in networks where:

- Layer 3 protocols are converging to IP (and possibly IPX).
- Connectivity to FDDI, Token Ring, and other non-Ethernet technologies is not required.

- Connectivity to the WAN is done through a separate device.
- RIP, RIP2, and OSPF are the routing protocols used on the network. Most Layer 3 switches don't support other protocols, such as BGP, IGRP, EIGRP, and the like.

As the bottom line, if your network is consolidating around IP and Ethernet, Layer 3 switching is for you.

Where to Put Layer 3 Switches in the Backbone

Layer 3 switching is usually designed into a network in one of two configurations: *centralized* or *distributed*. Alert readers might notice that this is another way of saying *collapsed* or *distributed*, but the terminology used here insinuates different properties about the architectures. Distributed Layer 3 switching, unlike a distributed backbone, requires upgrades on all major sub-LANs (*subnets* if you're using TCP/IP or *floors* using the example from step 4). Centralized Layer 3 switching can be implemented a sub-LAN at a time, but it requires a central Layer 3 switch. Because centralized Layer 3 switching is more flexible, it is usually the preferred choice in existing networks.

Centralized Layer 3 Switching

A centralized Layer 3 switched architecture closely resembles a routed backbone from step 4. The only difference is that the central device is a Layer 3 switch and not a high-end router. Remember that a Layer 3 switch, when using a RIP or OSPF-based routing scheme, acts much like a router.

Adding a new Layer 3 chassis switch to the network is fine when a brand new network is being designed, but you're never quite so lucky. In most cases, a central router already exists in the collapsed backbone. Are we advocating throwing away this costly item and replacing it with a Layer 3 switch? Absolutely not (well, maybe if your network supplier has a liberal return policy). Instead, you can easily add Layer 3 switches to existing routed backbones, leveraging the existing router(s). Figure 9.25 shows how a Fast Ethernet Layer 3 switch is used to add two new IP subnets to an existing collapsed backbone.

In this example, the network is using only the TCP/IP protocol, and because the Layer 3 switch supports this protocol, it acts like another router. The Layer 3 switch communicates with the existing router via RIP, RIP2, or OSPF protocols, transferring the necessary routing tables. In this way, the Layer 3 switch resembles the routing modules in the existing router.

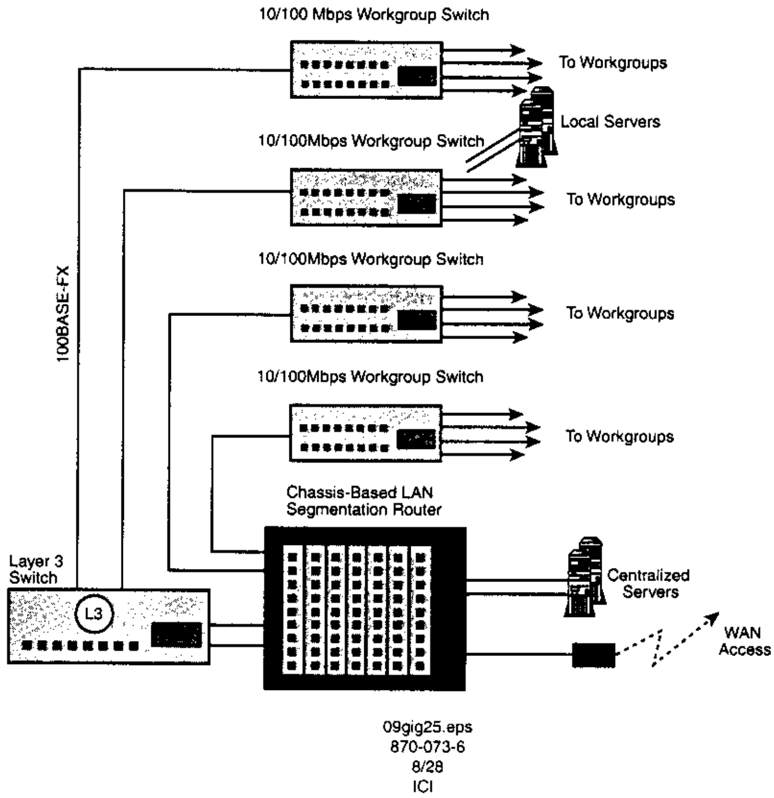

FIGURE 9.25 *The addition of a Layer 3 backbone switch to a routed, collapsed backbone.*

In the centralized scheme, you may want to consider adding Layer 3 capability to certain switches on each floor. For instance, on the third floor, the local servers may create a lot of broadcast traffic that doesn't need to be propagated beyond that floor. In this case, a Layer 3 switch in the third floor wiring closet will help alleviate the routing requirements of the central Layer 3 switch.

Be careful to consider all aspects of compatibility between your existing router and the new Layer 3 switch before connecting them. The following are some of the questions you should ask yourself before upgrading your network in this way:

- Does the Layer 3 switch support all required protocols (IP, IPX, and so on)?
- Does the Layer 3 switch support the required inter-router protocols (RIP, RIP2, OSPF, and so on)?
- Is port aggregation supported (Fast EtherChannel or 802.3ad)?
- Does the switch support enough ports? If not, can the port density be upgraded?
- Does the switch support fiber? Pay attention to the fiber connectors supported (see Chapter 6). You don't want to end up with a switch with SC connectors and a fiber run with ST connectors.
- Does the switch support uplink modules (ATM, Gigabit Ethernet, and so on)? Do those modules support Layer 3 switching?

Also, note that most router vendors provide Fast Ethernet Layer 3 switches on chassis modules. These modules are less expensive than full-fledged router modules but are often more expensive than buying separate Layer 3 switches. Consider all your options before selecting a solution.

Distributed Layer 3 Switching

Distributed Layer 3 switching is less flexible than the centralized model because it requires that all sub-LANs have a Layer 3 switch, which are interconnected by a high-speed, Layer 2 switch. While you are scratching your head, let us explain.

If you look at the example of a distributed Layer 3 switched architecture in Figure 9.26, the first thing you notice is that each floor has a Layer 3 switch in its wiring closet. The logic to this goes as follows. Because an end-station (either client or server) originates all the data on the network, that data will have to pass through a Layer 3 switch before it reaches the backbone. If each of the floors is equipped with a Layer 3 switch, the switch on each floor will perform all broadcast traffic filtering (and other routing functions). Thus, a broadcast packet originated by a client on the second floor will never propagate to the backbone, and therefore, to any other floor.

This is a more efficient way to filter the traffic because the central backbone switch does not deal with routed packets. If a packet needs to be routed from one floor to another, it is sent through the backbone switch, which forwards it according to Layer 2 (MAC address) information only. The Layer 3 switches on each floor need to update each other's routing tables (via RIP2 or OSPF), and these updates will also be sent across the backbone. Data traffic is reduced on the backbone, but additional router overhead traffic is increased, so it is hard to determine the net effect of the distributed architecture on any given network. It really depends, as in the centralized case, on the amount of data that needs to be routed.

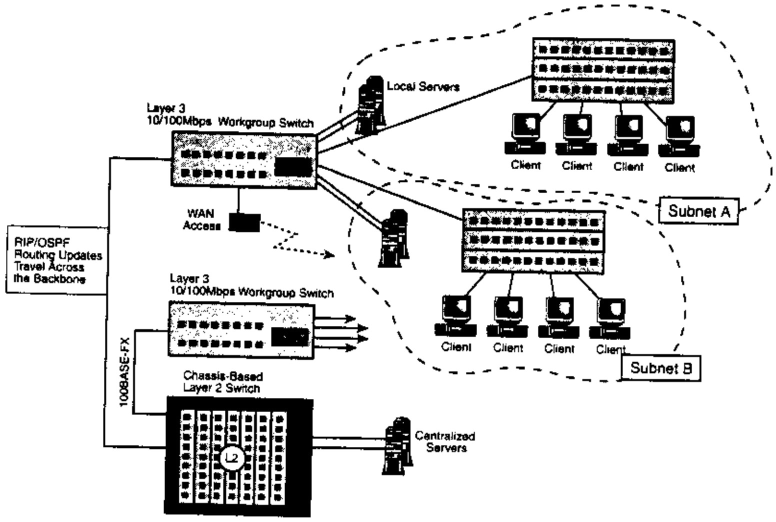

FIGURE 9.26 *Distributed Layer 3 switching in the workgroup. This architecture takes advantage of an existing Layer 2 switch in the core of the network.*

If you are going to deploy Layer 3 switches in each sub-LAN (or on each floor), you should also deploy a central switch with Layer 3 capabilities. This will cost more, but it gives you added flexibility in how you roll out future workgroups. A central Layer 3 switch signifies a centralized architecture, so our recommendation is this: if you have determined that you need Layer 3 switching, deploy a centralized Layer 3 switching architecture.

Performance and Cost Analysis

Distributed and centralized Layer 3 switched solutions cost approximately the same. With the centralized solution, the cost may be more in the central switch, whereas the distributed solution requires more, lower cost Layer 3 switches for each sub-LAN. Table 9.10 highlights the pros and cons of each architecture.

TABLE 9.10 A COMPARISON OF LAYER 3 SWITCH ARCHITECTURES

Issue	Centralized	Distributed
Layer 3 forwarding	Done in the core of the backbone and/or in the wiring closet	Done in the wiring closet
Backbone traffic	Higher than Distributed	Lower than Centralized
Requirements in workgroup	None	Layer 3 switching supported

Again, our recommendation is to use centralized Layer 3 switches and to use them in conjunction with your existing routers.

Configuring Layer 3 Switches

In the preceding discussion, Fast Ethernet was the inferred technology for the Layer 3 switches used, but Layer 3 capabilities can also be designed into a 10Mbps switch or into a Gigabit Ethernet switch. The speed of the central Layer 3 switch depends on the amount of traffic expected on the backbone. In a centralized architecture, because broadcast traffic is forwarded to the backbone, the central switch needs to be higher performance than the connected workgroup switches. If the workgroup switch is Fast Ethernet, the backbone switch should be Gigabit Ethernet, or at least multiple Fast EtherChannel aggregated 100Mbps connections. We'll discuss this more in Step 6.

Subnets

One of the primary benefits of a Layer 3 switch is its capability to filter broadcast traffic. This depends on how the subnets (sub-LANs, in some of our examples) are configured. *Subnets* refer to groups of IP addresses, which are often clustered according to physical location. A Layer 3 switch is most efficient when it is configured to map each port to an existing subnet, thus resembling a router. Figure 9.27 shows an example of this.

This configuration can be either manual or automatic, depending on the design of the switch. Switches that automatically determine the IP subnet on each port usually do so by monitoring the IP addresses and requests that are generated by the workgroup connected to that port. This is a clever system, but it can be easily corrupted by changing desktops or a faulty NIC. You're usually much safer directly assigning the subnet structure for the network and configure the Layer 3 switch to mirror that structure.

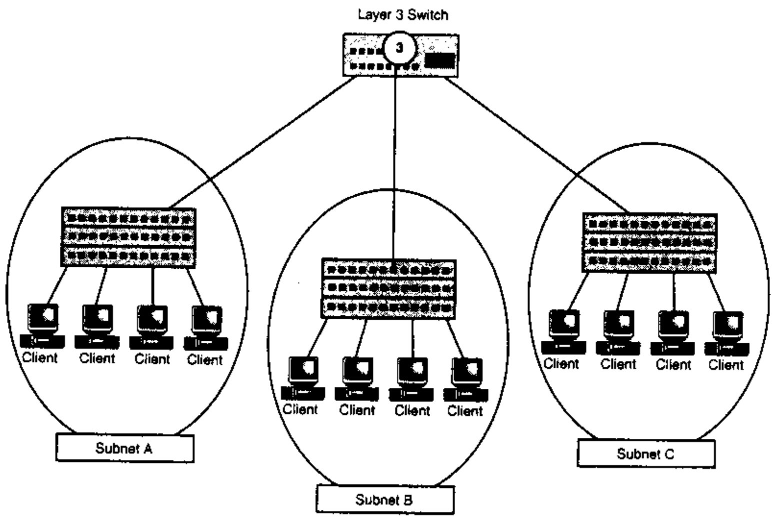

FIGURE 9.27 *Configuring a Layer 3 switch to match physical IP subnets.*

VLANs

Chapter 11 discusses VLANs in more detail, so this section is restricted to how you can use virtual LANs to aid in the deployment of Switched and Fast Ethernet. Virtual LANs are an important tool that, if used properly, can maximize the performance and minimize the maintenance of your network. VLAN implementations vary quite dramatically from vendor to vendor, but generally fall into five categories: port-based, protocol-based, MAC-based, IP multicast-based, or tagged.

Port-Based VLANs

Port-based VLANs were the earliest instantiation of the technology. They basically allow certain ports on a switch to be isolated from other ports on the switch. The two VLANs can be connected via an internal router (Layer 3 switch) or an external router. Some switches even support an interswitch communication scheme that allows ports on switch A to be on the same VLAN as ports on switch B. These interswitch communications are, however, proprietary and limited to single-vendor implementations.

Port-based VLANs can be helpful in isolating workgroups that use non-routable protocols, such as NetBEUI, from the rest of the network.

Protocol-Based VLANs

VLAN identification by protocol is perhaps the best way to use VLANs in an existing routed network. Switches that identify Layer 3 information, such as protocol type or IP address, can then assign that frame to a given VLAN. If an existing interswitch VLAN communication scheme exists, this VLAN identification can be transmitted to other switches on the network.

Most Layer 3 switches actually use protocol-based VLANs to segment traffic internal to the switch and then use RIP2 or OSPF to communicate the information to other switches. In this way, the switch is completely compatible with existing routers but still uses VLANs to segment traffic within the switch. Figure 9.28 illustrates this concept.

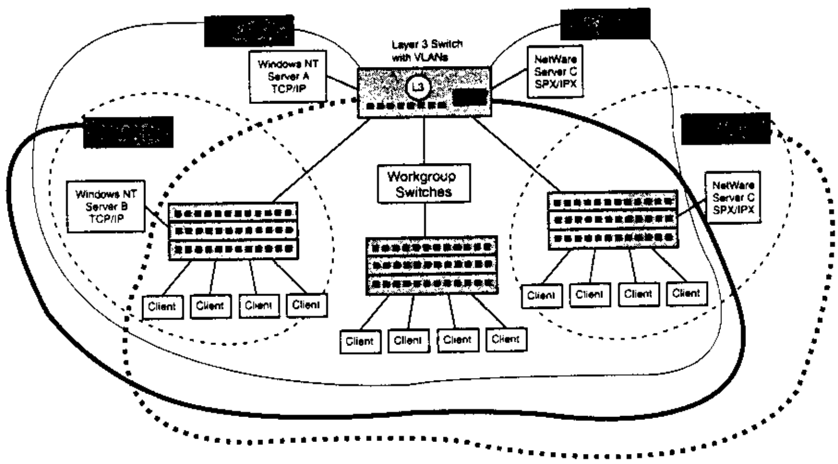

FIGURE 9.28 *Using VLANs to map servers and clients through a switched network. VLANs A and C span all clients but serve different protocols. VLANs B and D are workgroup specific.*

MAC Address-Based VLANs

A list of MAC addresses can be assigned to a particular VLAN, which means that a given end-station, no matter where it is on a network, can be assigned to a particular VLAN. This type of scheme works well in networks where users take their laptops with them to empty cubicles or conference rooms. Because the laptop has a given MAC address, the network can identify it no matter where it is attached.

The downside to MAC-based VLANs is that they need constant maintenance and monitoring. A laptop may contain a PCMCIA card with one MAC address and a docking station with another. Also, PCMCIA cards with new MAC addresses are installed all the time. Additionally, security is poor with the MAC-based method, as most NICs allow a software version of the MAC address to be used.

IP Multicast Address-Based VLANs

VLANs can also be identified with an IP multicast address. An IP multicast address is a proxy address for a larger group of IP addresses. If a frame needs to go to this group of IP addresses, it is sent first to the proxy IP address and then forwarded to the whole group. Membership in the group is voluntary, meaning that each desktop can determine whether its IP address should be included in the group. This kind of VLAN identification is useful in networks where video or audio data is being broadcast on the network and only a select few users are allowed to view or listen to the information.

802.1Q Tagged VLANs

The preceding VLAN implementations are nice, but if you want to implement VLANs in your production network, you might as well start getting used to the concept of a tagged frame (discussed in Chapters 3 and 8). The IEEE 802.1Q tagged VLAN standard is where most of the industry is headed, and where you should be as well. A tagged frame has a field called the *VLAN ID*, which indicates to which VLAN the frame belongs. All switches in the network would know this VLAN ID because switches are responsible for communicating VLAN membership amongst themselves.

What identifies which VLAN ID belongs in each frame? In theory, the end-station will identify the VLAN ID when the frame is originally created. If a frame comes into a switch with no VLAN ID, a switch can opt to insert a VLAN ID and forward the frame. Figure 9.29 shows the results of a forwarded frame depending on the final destination.

Should I Use VLANs in Step 5?

Through 1998, VLANs were still not standardized among major switch vendors and, therefore, were still proprietary. The 802.1Q standard will likely change this dramatically and become the de facto VLAN standard in 1999. Therefore, if you are considering VLANs, you should be considering 802.1Q support in your network components.

Although VLANs show a lot of promise, you can most likely meet the performance and segmentation needs of your network with simple Layer 3 switch deployment. Layer 3 switch deployment leverages existing routers and uses standard routing protocols. VLAN deployment means a new set of management tools and a proprietary interswitch communications scheme. With this in mind, you're best off using VLANs only in a limited fashion in workgroups where you can justify the need. Table 9.11 highlights some common examples of logical VLAN uses.

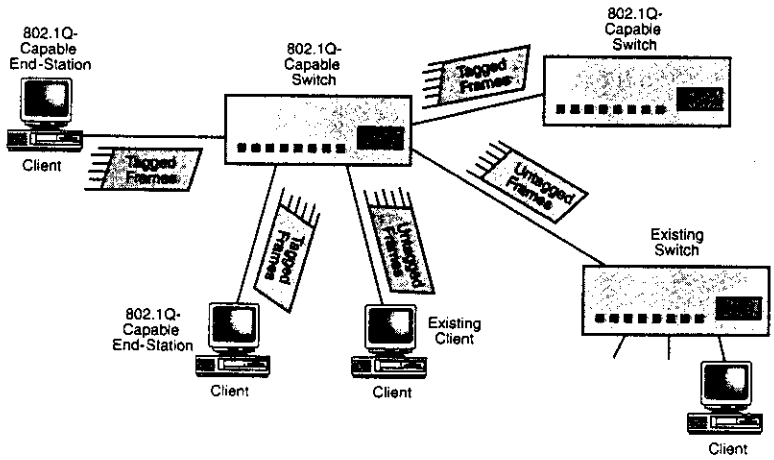

FIGURE 9.29 *802.1Q tagging and how tags are forwarded through a switched network.*

TABLE 9.11 PRACTICAL VLAN USES

VLAN Type	Complexity	Example of When to Use
Port-based	Low	All stations on port 10 are in VLAN 2.
MAC-based	High	Station `00-00-81-01-01-01` is in VLAN 3.
Layer 3 address	Med	IP subnet traffic `256.256.127.100` is on VLAN 1.
Protocol-sensitive	Low	IPX traffic on port 10 is in VLAN 4.
IP Multicast-based	High	IP Multicast address X is associated with VLAN Y.
Tagged	High	Standards-based 802.1Q/p tagged or proprietary tagging scheme.

Note

3Com has a great white paper on its Web site (`www.3com.com`) that gives more detail on VLANs.

Deployment Step 6: Gigabit Backbone Switches

Of all the recent developments in local area networking, Gigabit Ethernet has to be one of the most compelling. Not only does Gigabit Ethernet address the bandwidth requirements of today's backbones, but it has also shifted the bottleneck from the network to the server for the first time in nearly a decade.

This section identifies where Gigabit Ethernet should be deployed in your Ethernet network.

Connecting Multiple Workgroups

The obvious place where Gigabit Ethernet is needed first is the backbone switch. In step 4, we identified how to connect a collapsed backbone using aggregated Fast Ethernet switches. In step 6, we'll replace the central switch or router with a Gigabit Ethernet backbone switch.

Figure 9.30 shows the configuration of a Gigabit Ethernet collapsed backbone. The Fast Ethernet switches on each floor need to be upgraded with a Gigabit Ethernet uplink module in order to be connected to the central switch.

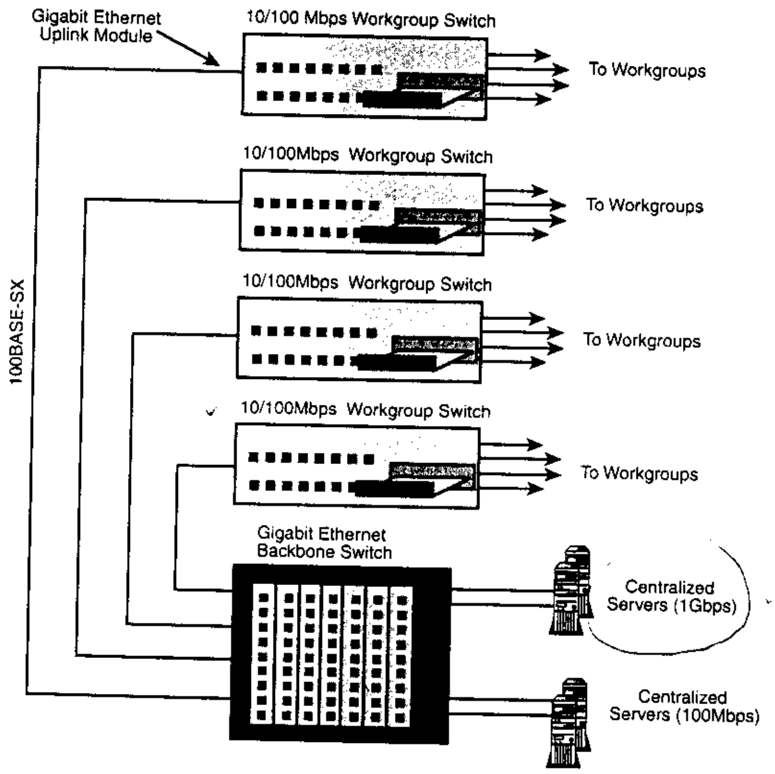

FIGURE 9.30 *Building a Gigabit Ethernet Backbone.*

In the previous section, we discussed the benefits of centralized Layer 3 switching, so the Gigabit Ethernet switch deployed should also be capable of Layer 3 switching. The packet per second routing performance of this switch should depend upon the model for backbone traffic in your network.

Because 1000BASE-LX Gigabit Ethernet over multiple-mode fiber is limited to 500m, you may need 100BASE-SX single-mode Gigabit Ethernet for especially long links. A 100BASE-LX Gigabit Ethernet switch with the capability to add a few 1000BASE-SX connections is the proper solution for this example.

This model for Gigabit Ethernet deployment allows for each desktop to connect at 100Mbps switched speeds.

Scaling Gigabit Ethernet Switches

What happens if more than one Gigabit Ethernet switch is needed in the backbone? There are really two options for this scenario. The first is to install a chassis-based solution with a multiple-Gbps backplane. A good example of this is the Extreme Networks Black Diamond 6800 Switch. Gigabit Ethernet modules can be added as the network grows, allowing for a flexible deployment. Of course, 48 ports of Gigabit Ethernet and 64Gbps of bandwidth doesn't come cheap: Extreme costs about $2,750 MSRP per Gigabit Ethernet port at the time of this book's publication.

The second option is to buy a fixed-configuration Gigabit Ethernet switch. A fixed-configuration switch will come with, say, eight ports. In order to enlarge this network beyond eight ports, an additional Gigabit Ethernet switch is needed. How would you connect these two switches? Our familiar port aggregation schemes (EtherChannel or 802.1ad) scale to Gbps speeds, allowing you to trunk up to four Gigabit Ethernet ports together. Figure 9.31 illustrates the Gigabit port aggregation scheme, creating a 12 port Gigabit Ethernet switched backbone. Because the interswitch connection is full-duplex, the transfer rate is 4Gbps to 2Gbps each way.

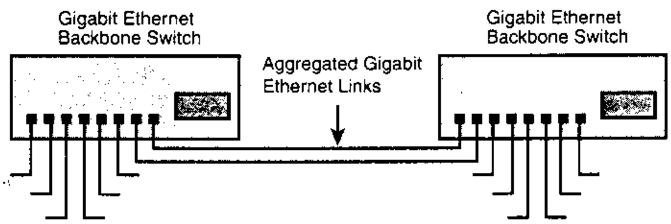

FIGURE 9.31 *Connecting two Gigabit Ethernet backbone Switches with port aggregation. The configuration shown here provides 4Gbps of bandwidth between the two switches: two in each direction.*

Connecting Gigabit Ethernet Switches to Gigabit Ethernet Servers

Another key aspect of Gigabit Ethernet deployment is how to connect backbone servers to the Gigabit Ethernet backbone switch. As shown in Figure 9.32, you can do this in one of three ways.

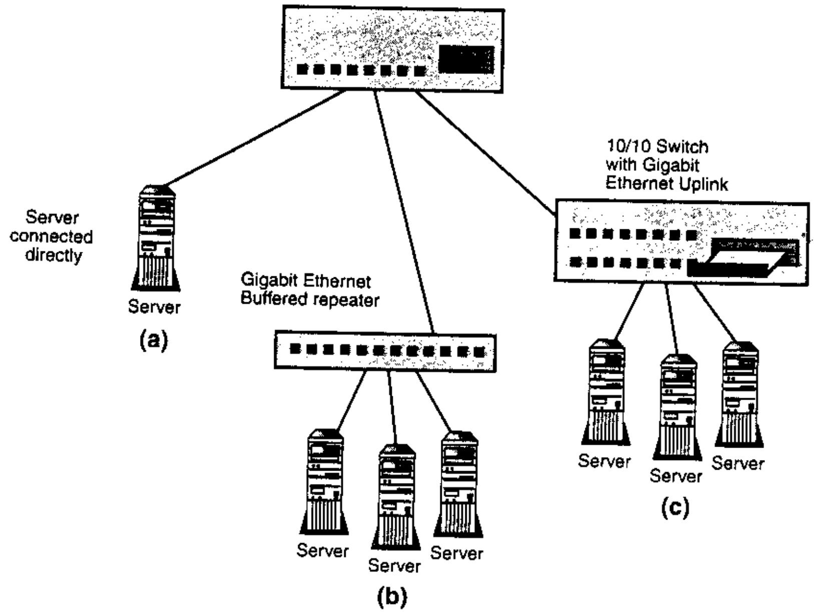

FIGURE 9.32 *Three different ways to connect servers to a Gigabit Ethernet backbone switch: (a) directly to a switched port, (b) through a Gigabit Ethernet buffered repeater, and (c) through a 10Mbps/100Mbps switch with a Gigabit Ethernet uplink.*

You should consider only the highest performance servers for a direct connection to a Gigabit Ethernet backbone switch. As discussed in Chapter 8, very few servers will be able to fill a Gigabit Ethernet pipe, so a direct connection may not be the most efficient method. Servers of the future, however, will be able to communicate at these speeds, so one should plan for a certain amount of direct Gigabit Ethernet server connections. This is shown in Figure 9.32 (a).

You can still connect servers via Gigabit Ethernet, but you can aggregate them to form Gigabit Ethernet server farms. A Gigabit Ethernet buffered repeater is used for this purpose. Because most servers can transfer data faster than 100Mbps but slower than 1Gbps, sharing a Gbps of bandwidth seems like an efficient option. A Gigabit Ethernet buffered repeater, used in this fashion, will be a good stopgap solution for a Gigabit Ethernet server farm. This is shown in Figure 9.32 (b).

The most common backbone connection will still be Fast Ethernet. Gigabit Ethernet switches will connect to these servers via a Fast Ethernet switch and a Gigabit Ethernet uplink module. This is shown in Figure 9.32 (c).

Summary

Table 9.12 outlines the six steps to Switched, Fast, and Gigabit Ethernet deployment. The right column provides examples of an actual company's history of moving through these steps. There were no shortcuts and no dramatic fix-all-your-problems revolutions.

The six-step deployment plan has built the company a network that has grown with its needs, been utterly upgrade capable, and still has plenty of room for future growth. And by the way, in June 1994, most of the company was FDDI and Token Ring. Today, it is mostly Ethernet.

Table 9.12 The Six Steps to Deploying Switched, Fast, and Gigabit Ethernet and How They Map to a Large Company's Actual Deployment

Step	Description	Timeline	Justification
1	Adding 10Mbps Switching	June 1994	Power workgroups in Engineering Department required more bandwidth.
2	Adding 10/100Mbps	June 1995	New PCs cost $2,000 each NICs, with $30 extra per PC for 10/100Mbps future-proofing was approved.
3	Converting Workgroups to Fast Ethernet	Feb 1996	New workgroups created from new PCs with 10Mbps/100Mbps NICs required hubs. After pricing a few 10Mbps hubs, MIS realized that 100Mbps hubs were not much more expensive.
4	Upgrading the backbone to Switched and Fast Ethernet	March 1997	This was the hardest step to get approved. When the Cabletron chassis-based FDDI router was overloading about every week, budget was approved for a new core device. A chassis switch with 100BASE- FX routing modules was approved and each department put up the funds for a new 10/100Mbps workgroup switch and FX uplink module.

Step	Description	Timeline	Justification
5	Adding Layer 3 switches and VLANs	March 1998	A cost comparison between 100BASE-FX routing modules and fixed configuration 100BASE-FX Layer 3 switches showed that Layer 3 switching was about 25% less expensive.
6	Adding Gigabit Ethernet switching to the backbone	TBD	Traffic on the 100BASE-FX backbone hasn't yet reached alarming levels, but the increase due to the Internet/intranet is causing backbone traffic to rise every month. Currently considering Gigabit Ethernet as a future-proofing option.

Many LANs may differ from the examples discussed in this chapter; therefore, you should think of this six-step plan as a guide rather than a strict set of instructions. The steps are outlined in such a way that you can plan their execution over a long period of time, with steady performance gains. This will result in an interesting phenomenon: the more users that are added to the network, the more Fast and Gigabit Ethernet will be deployed. The performance of the newer users will be dramatically better than the existing users will remember. When the existing users are upgraded with 10/100Mbps switches, they too will enjoy the increased performance of Switched and Fast Ethernet. Gigabit Ethernet and Layer 3 switching complete the picture.

In Chapter 10, we apply the lessons of this chapter to some specific deployment examples, including building brand new 100Mbps networks, attaching Fast Ethernet networks to FDDI, and using a real Gigabit Ethernet backbone.