Migration Strategies for FDDI-Based Campus Backbones

Enterprise Technical Marketing

Executive Summary

In a world of Pentium PCs, high-speed servers, and an ever-growing slew of bandwidth-hungry networked applications, the importance and demands on campus intranets are intensifying. In the wiring closet, client-side computing is driving a need for better per-user bandwidth management, fueling a migration from shared-media devices to Ethernet switches. This migration in turn has fueled a need to provide greater performance, scalability, and fault tolerance in the campus backbone.

Looking more closely at the campus backbone, we are witnessing shifts in both technology and design. First and foremost is the change in technology. In the past, Fiber Distributed Data Interface (FDDI) was the backbone technology of choice. Today, however, Asynchronous Transfer Mode (ATM) and Fast Ethernet are quickly gaining hold over the campus backbone. And with these new technologies, new designs are being introduced to achieve optimum performance and fault tolerance.

This paper provides an in-depth discussion of campus backbone design. Using FDDI as a starting point for the discussion, it addresses the technical objectives for reengineering the campus backbone. This paper also presents a set of design options based on the available high-speed networking technologies. It concludes with a practical discussion focusing on implementation issues.

To this end, the document follows the outline detailed as follows:

- Introduction: Today’s Campus Backbone—FDDI
- Reengineering the Campus Backbone
- FDDI Solutions
- ATM Solutions
- Fast Ethernet Solutions
- Conclusion

Cisco Systems

Introduction: Today's Campus Backbone—FDDI

FDDI Overview

FDDI was introduced in the mid-1980s, and until recently has been embraced as the premier campus backbone technology. FDDI is a 100-megabit per second (Mbps), dual-ring, token-passing technology. The dual-ring concept refers to FDDI's underlying Layer 1 (physical layer) technology. FDDI employs two distinct rings to which devices can be connected. The dual-ring design is intended to provide low-level fault tolerance. If the primary ring fails, the secondary ring is employed as a safeguard mechanism. Figure 1 depicts FDDI's dual-ring topology.

Figure 1 FDDI Dual-Ring Topology

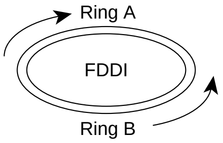

Depending on the physical connections available, devices can connect either to the primary ring or to both the primary and secondary rings. Access to the FDDI network itself is arbitrated using a token that travels around the FDDI network (hence the notion of an FDDI "ring"). A device can access the network only after it has received the token. FDDI operation is similar in concept to Token Ring (albeit faster) and contrasts with Ethernet technologies where devices can access the network at any time. Like Token Ring, FDDI is often referred to as "deterministic," because network access is consistently regulated by means of the token. Figure 2 illustrates the FDDI token passing operation compared with Ethernet technologies.

Figure 2 Network Access—FDDI vs. Ethernet

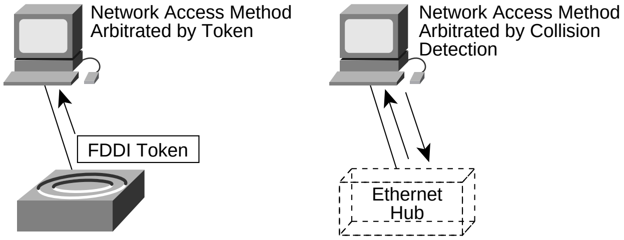
Network Access Method Arbitrated by Token
FDDI Token
Network Access Method Arbitrated by Collision Detection
Ethernet Hub

As mentioned previously, devices can connect either to the primary ring or to both the primary and secondary rings. These two connection options are commonly referred to as single attachment stations (SASs) or dual attachment stations (DASs). As their names imply, SAS connections connect only to the primary ring, while DAS stations connect to both rings. Clearly, to take advantage of FDDI's fault-tolerant capabilities, DAS connections are more desirable. Figure 3 shows SAS versus DAS connection options.

Figure 3 Single-Attached Station versus Dual-Attached Station

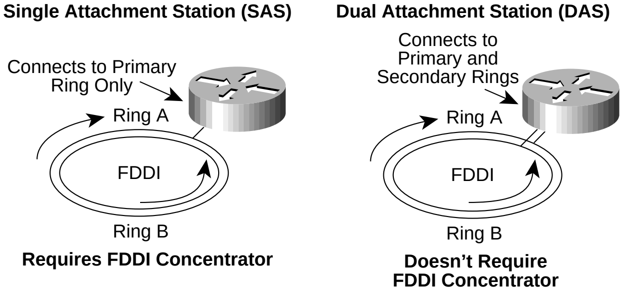
Single Attachment Station (SAS)
Connects to Primary Ring Only
Ring A
FDDI
Ring B
Requires FDDI Concentrator
Dual Attachment Station (DAS)
Connects to Primary and Secondary Rings
Ring A
FDDI
Ring B
Doesn't Require FDDI Concentrator

Typically, routers, which require fault-tolerant connectivity, use DAS connections to connect to an FDDI network, while other devices (end stations, perhaps) use less expensive SAS connections. From a network design standpoint, we will assume DAS connections.

FDDI-Based Campus Backbone Design

Based on the above discussion, there are two principal ways to design FDDI-based campus backbones. The first design is based solely on DAS-equipped routers. Each DAS has two ports, an A port and a B port. In this design, the routers that form the campus backbone are simply connected in an "A-to-B" fashion. The end result is a ring topology in which all routers are connected together. Figure 4 illustrates this design.

Figure 4 DAS-Equipped Router-Only Design

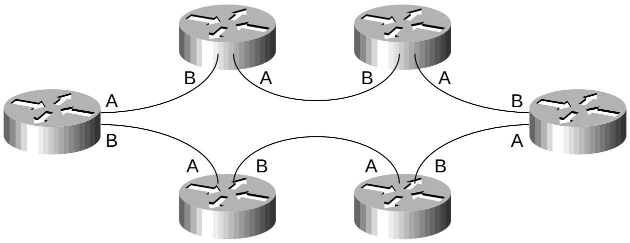

As Figure 4 illustrates, this FDDI implementation is very simple. This design yields high-speed connectivity and offers some degree of fault tolerance. Fault tolerance is qualified, because under certain circumstances this design can be problematic. Consider Figure 5, where six DAS-equipped routers are connected. If router A fails, the FDDI ring must compensate for the failed device and reestablish the FDDI ring with the remaining functional devices. The underlying operation that occurs is called "ring-wrap." The two routers adjacent to the downed router perform a ring-wrap so that the FDDI ring can reconverge based on the revised set of functional devices. Figure 6 details the ring-wrap operation.

Figure 5 Single Router Failure

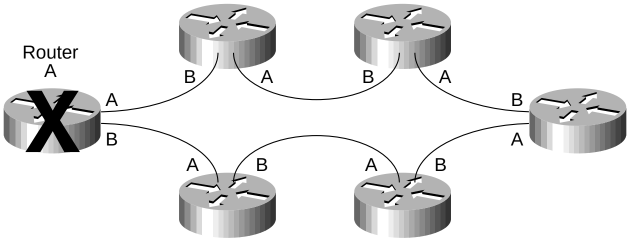

Figure 6 FDDI Recovery—Ring-Wrap Operation

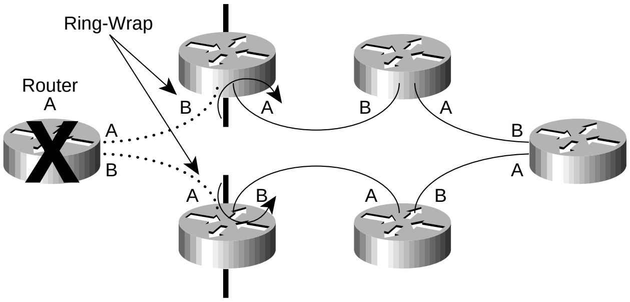

As this sequence highlights, some level of fault tolerance can be achieved in this design based on ring-wrap operation. There are some situations, however, where ring-wrap operation can be problematic. In particular, ring-wrap can result in distance problems with multimode fiber, depending on the physical topology.

And in other situations, ring-wrap operation alone is simply not enough to deliver fault tolerance. Consider the scenario in Figure 7, where routers A and C both fail. As illustrated, the network experiences nonconsecutive router failure.

Figure 7 Nonconsecutive Router Failure

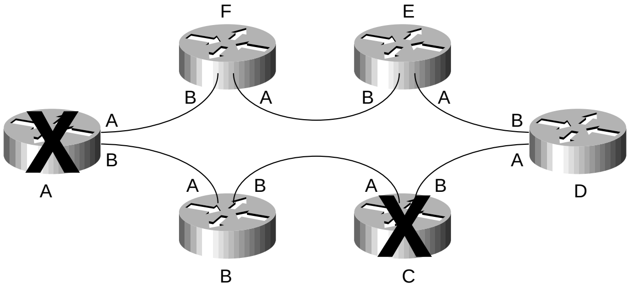

In this situation, Routers B, D, and F perform a ring-wrap operation to try to reestablish the FDDI ring. Unlike Figure 7, in this case the ring-wrap causes Router B to become isolated from the rest of the router infrastructure (see Figure 8).

Figure 8 Router Isolation following Ring-Wrap

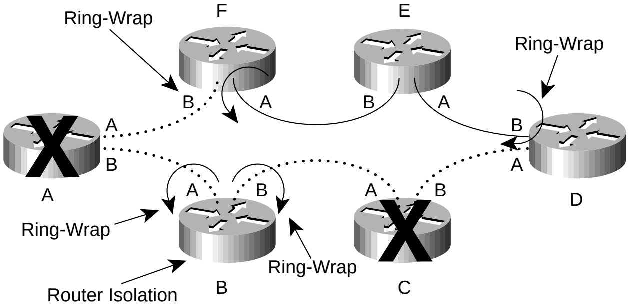

Clearly, this is a suboptimal fault-tolerant strategy. To address this shortcoming, we recommend a new design introducing the use of an FDDI concentrator, a device that incorporates a set of both A and B ports as well as numerous M ports. With a pair of FDDI concentrators in conjunction with DAS-equipped routers, a "dual-homed" design optimized for fault tolerance can be deployed. With this design, the two concentrators are connected using the A and B ports. This setup forms the core of the FDDI ring. The routers are subsequently dual-homed to both concentrators using the A and B ports on the router and two M ports on the concentrators. Figure 9 details a dual-homed design.

Figure 9 Dual-Homed FDDI Backbone Design

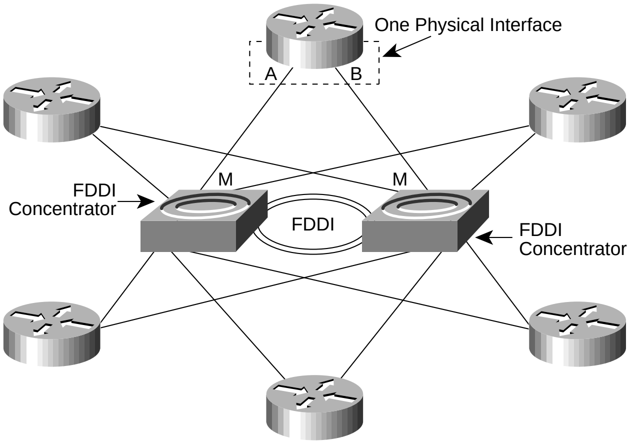

The primary benefit of this design is superior fault tolerance. In a dual-homed design, the network can withstand either concentrator failure or router failure and still be operational with the remaining devices. In Figure 10, concentrator B and router C are both down, but notice how connectivity is still preserved.

Figure 10 Device Failure in Dual-Homed Design

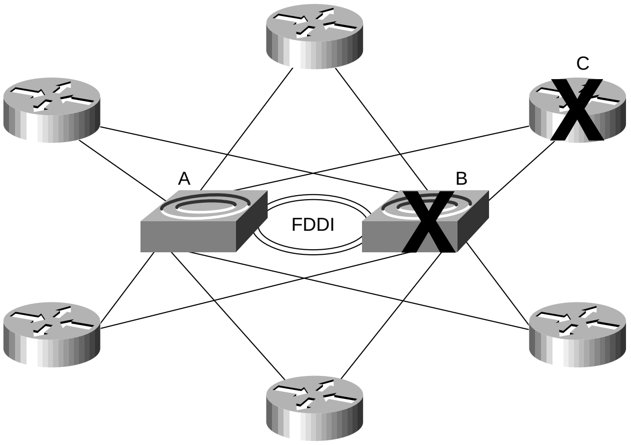

It is evident that the dual-homed design is optimal for the campus backbone from both performance and fault-tolerance standpoints. Hence, we will use the dual-homed FDDI backbone illustrated in Figure 11 as our starting point for a discussion of reengineering and migration options.

Figure 11 The FDDI-Based Campus Backbone

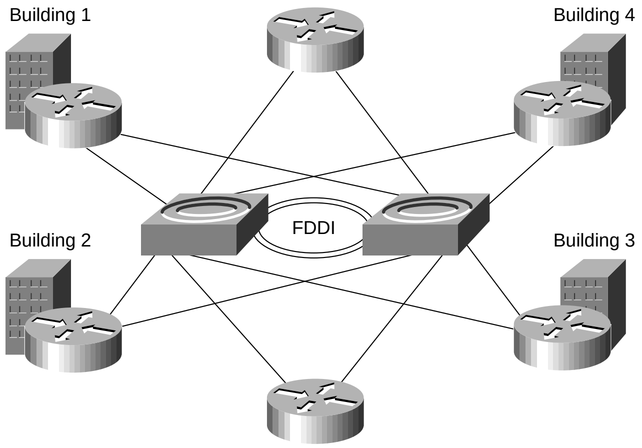

Reengineering the Campus Backbone

Technology Options—FDDI, ATM, Fast Ethernet

Few would argue against the success of FDDI in campus backbones. It provides high-speed communication, it has favorable distance limitations (2 km for multimode fiber, 100 km for single-mode fiber), and it has withstood the test of time.

Yet looking into the future, FDDI is met with stiff competition, mainly from ATM technologies and the more affordable Fast Ethernet technologies. Against these competitors, FDDI is criticized for being a shared-media technology, for being expensive, and for not offering a migration path beyond 100 Mbps. The designs discussed in the previous section implement FDDI in a shared-media environment—there is only one ring with one token that each connected device must contend for. Compare this scenario with switched Fast Ethernet, for example, where there is no contention for directly connected devices. FDDI is staging a minor comeback with FDDI switching, but as we will discuss later, it has its limitations.

The shared-versus-switched argument aside, two other forces hinder FDDI's ultimate success—cost and gigabit migration. Put simply, FDDI is expensive—very expensive. A PCI-based DAS FDDI adapter costs approximately $1995, while a PCI-based fiber Fast Ethernet adapter costs $700 (comparison is between 3Com's FDDILink and Adaptec's ANA-6910/FX-SC). And of course if unshielded twisted-pair (UTP) cabling is available, the cost is even more favorable for Fast Ethernet. A typical PCI-based Fast Ethernet adapter costs $149.

Lastly, there is the gigabit migration discussion. Both ATM, with Optical Carrier (OC)-48 technology, and Fast Ethernet, with gigabit Ethernet, have very clear goals in the area of gigabit and higher-speed networking. FDDI simply does not. For many, this fact alone has been the primary impetus for migrating to Fast Ethernet or ATM. Customers simply want technologies that can grow or scale with their demands. Both Fast Ethernet and ATM are delivering on that front while FDDI has not.

Drivers for Reengineering

Now let's discuss today's campus backbone—the dual-homed FDDI backbone. At present, the current infrastructure may be sufficient, but the future seems to be pointing to more bandwidth-intensive operations, operations that ultimately will tax the campus backbone considerably more. Cross-functional computing, client/server computing, and networked multimedia are all appearing on the IS horizon as network services (see Figure 12). Considering this fact, the campus backbone—the core of the network—is under scrutiny and its capabilities questioned. For many, the conclusion is obvious—the current dual-homed FDDI backbone will not scale with the growth of network-based computing. This scenario in turn leads to design modifications based on either the existing FDDI infrastructure or a migration to ATM or Fast Ethernet technologies.

Figure 12 Drivers for Reengineering

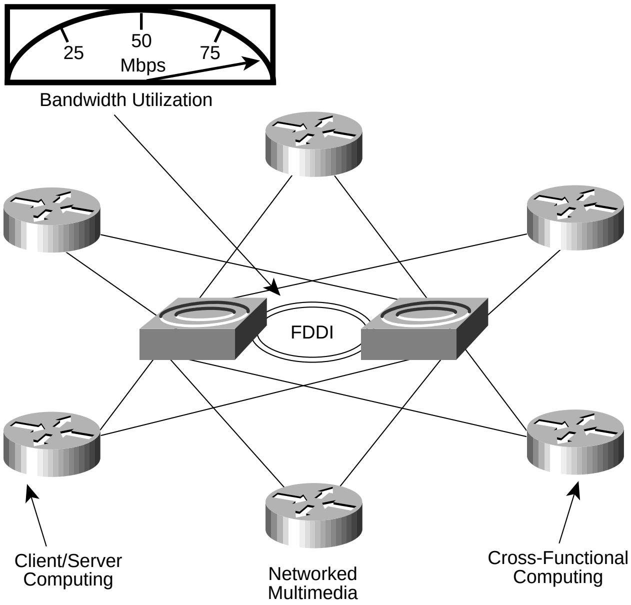

Technical Objectives for Reengineering

Essentially there are three solution sets for reengineering the campus backbone, each based on a specific technology—FDDI, ATM, or Fast Ethernet. Evaluation of the efficacy of each solution set is based on its ability to execute on the principal technical objectives of backbone reengineering, which include:

- Alleviate backbone congestion
- Ensure a fault-tolerant infrastructure
- Provide a migration path to higher-speed networking

The first point, alleviating backbone congestion, is a relatively straightforward concept and is best solved with affordable high-performance Layer 2 switching technologies. The third point is also relatively straightforward. In order for the backbone to endure, it must be able to scale its bandwidth capabilities and offer a migration path to higher-speed networking.

It is worth calling attention to the second point, however, of ensuring a fault-tolerant infrastructure. Ultimately fault tolerance is defined by uninterrupted end-to-end communication. To achieve this goal, fault tolerance ideally should be addressed at Layer 3 with high-performance routing algorithms (Multiprotocol Enhanced IGRP®, Open Shortest Path First [OSPF], and so forth). By solving fault tolerance at Layer 3, we leverage network-layer intelligence to ensure multiple paths and dynamic convergence in the event of network faults. Figure 13 details Layer 3 routing operation and fault tolerance.

Figure 13 Routing Operation: Layer 3 Fault Tolerance

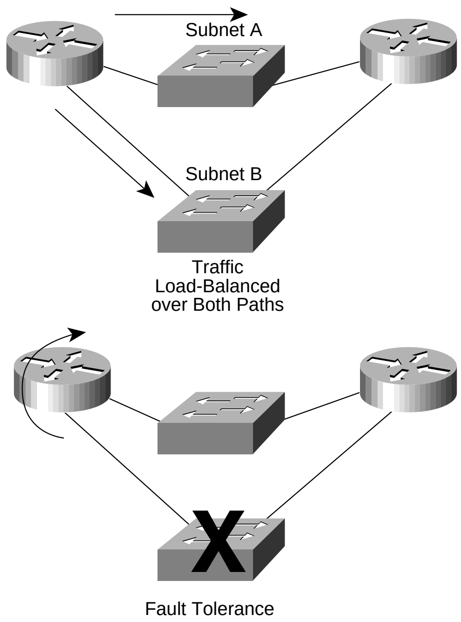

As illustrated in Figure 13, leveraging Layer 3 technologies serves to boost backbone performance while also enabling robust fault-tolerant mechanisms. This sharply contrasts the dual-homed FDDI design where fault tolerance is defined at Layer 1. In this design there is no load-balancing, and fault tolerance targets physical connectivity rather than Layer 3 connectivity.

Based on the combination of Layer 2 switching technologies and Layer 3 intelligence, Figure 14 provides a generic schematic of the new campus backbone.

Figure 14 The New Campus Backbone—Distributed Layer 2/Layer 3 Design

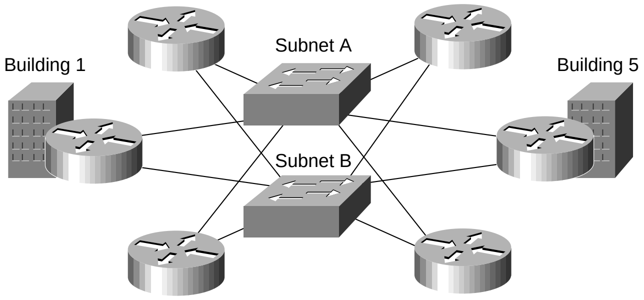

Campus Backbone Design Discussion

It is important to note that in the process of migrating from the dual-homed FDDI design to the combined Layer 2/Layer 3 architecture depicted in Figure 14, the network hierarchy was preserved. Network hierarchy is achieved at Layer 3 based on the network layer addressing scheme and the routed infrastructure. Specifically, in the design there are three distinct components in the network: Layer 2 switched access, intelligent Layer 3 packet classification/distribution, and high-speed Layer 2 core switching (see Figure 15).

Figure 15 Wiring Closet Connectivity to Campus Backbone

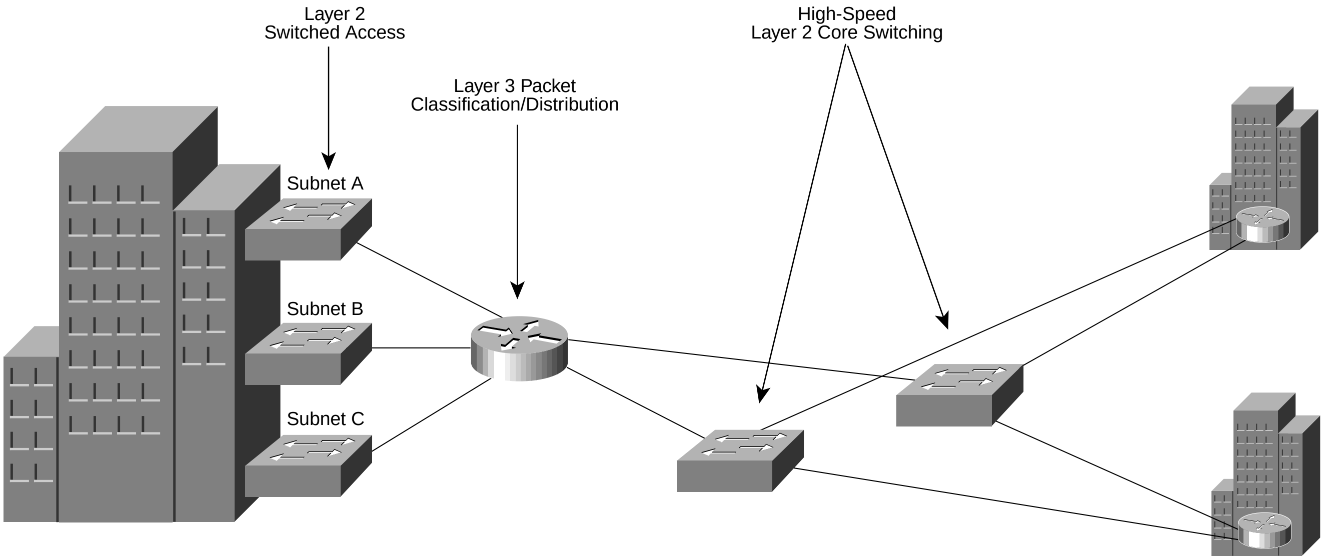

From the wiring closet perspective, the router serves to segment Layer 2 broadcast/multicast domains and provides an intelligent, high-speed Layer 3 ingress point to the campus backbone. Some may argue against routers at ingress points from a performance standpoint. But consider Cisco 7500s at the ingress points. With NetFlow™ switching, these routers can deliver approximately 270,000 packets per second (pps). And with distributed switching and Virtual Interface Processor (VIP)-2 line cards, performance hits the million pps mark (see Figure 16).

Figure 16 Role of Routing in the Campus

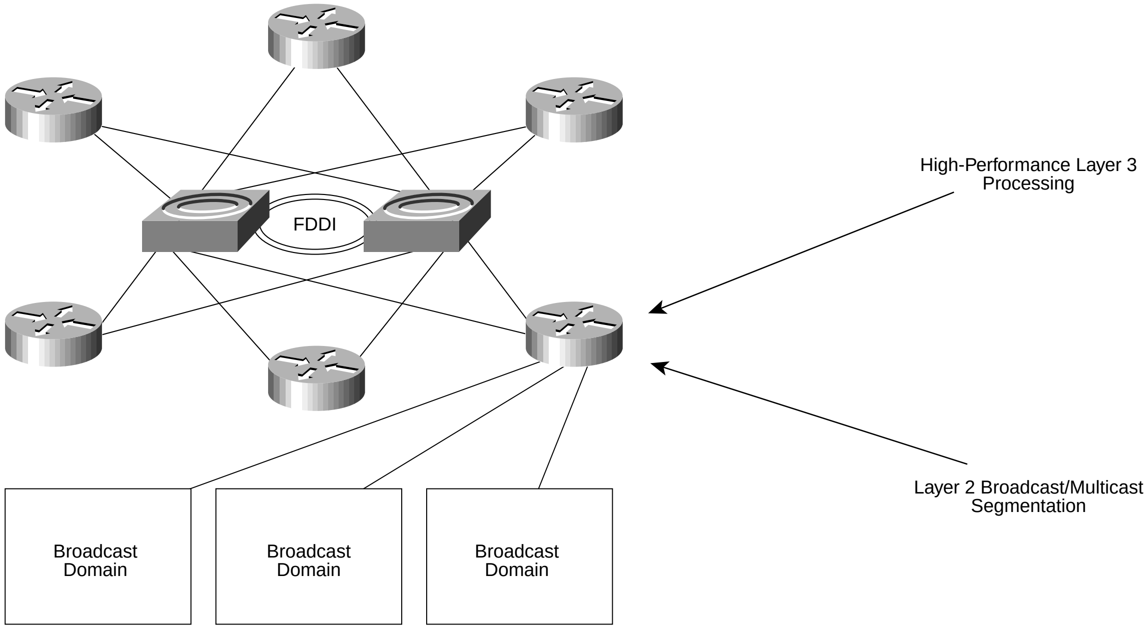

As illustrated in Figure 16, the routers play a crucial role in connecting wiring closets to the campus backbone. While this design is recommended, another design is often considered, based predominantly on Layer 2 switching (see Figure 17).

Figure 17 Layer 2 Design

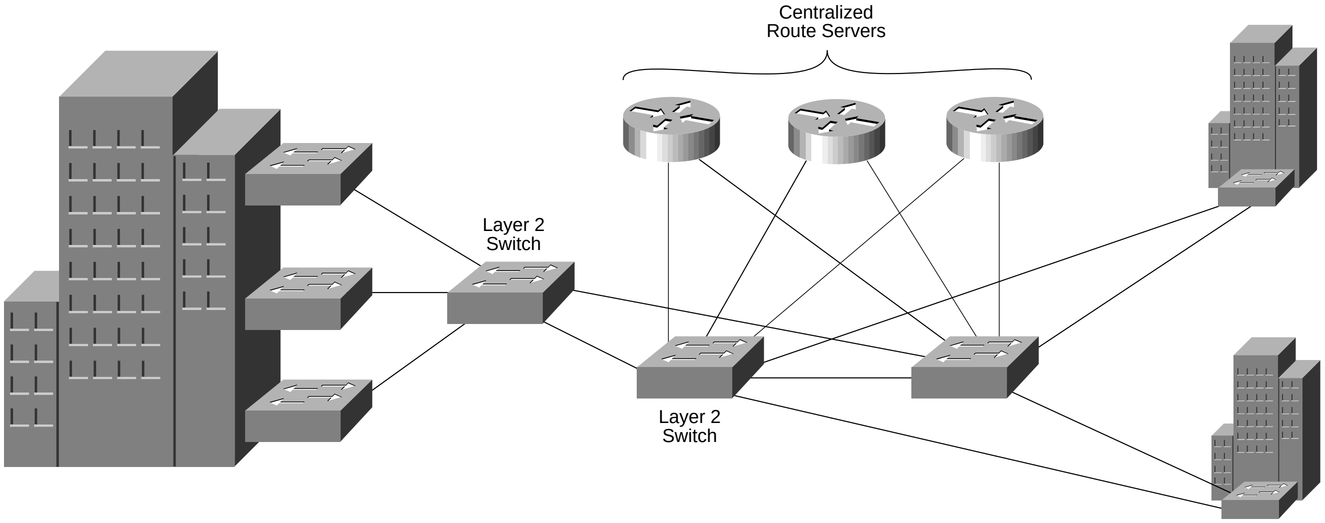

In this design, based on Layer 2 switching, the wiring closets are connected together, and to the backbone, by an upstream Layer 2 switch. Migrating in this fashion flattens out the entire campus. Without VLANs and VLAN routing, there is no Layer 2 broadcast/multicast segmentation, and virtually all Layer 3 intelligence is removed. As a result, Layer 2 broadcast or multicast traffic is transmitted to all switch ports within the Layer 2 fabric (see Figure 18).

Figure 18 Layer 2 Design

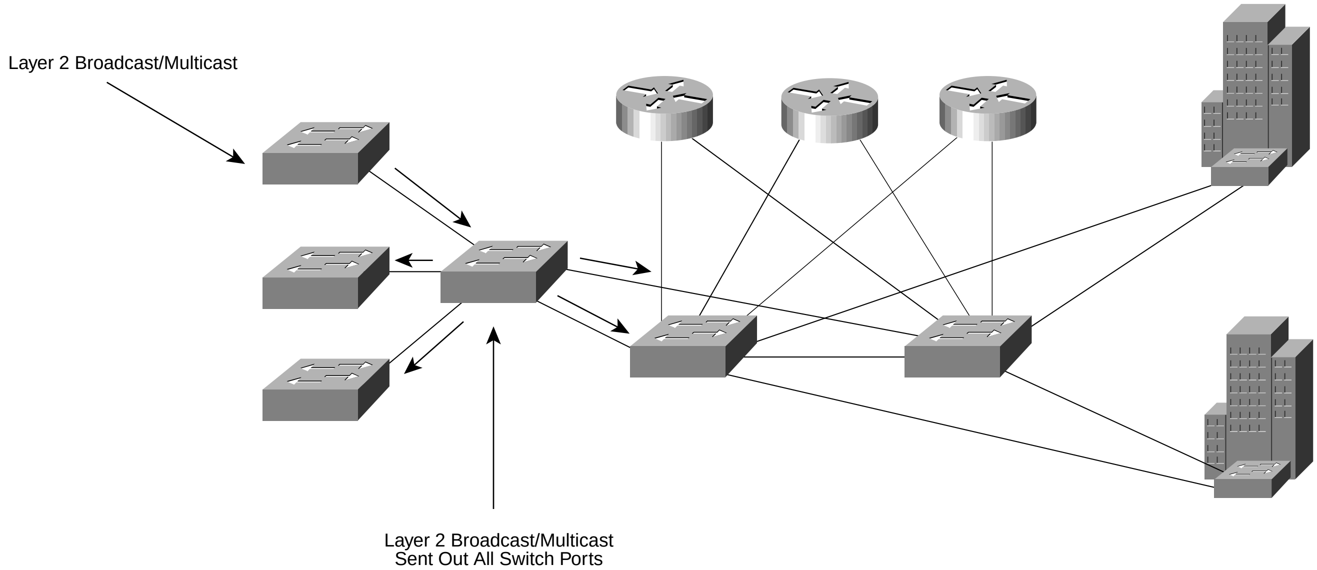
Layer 2 Broadcast/Multicast
Layer 2 Broadcast/Multicast
Sent Out All Switch Ports

As Figure 18 illustrates, replacing the backbone routers with switches directly impacts the amount of traffic on the campus backbone. With a router at the ingress point (see Figure 19) to the campus backbone, traffic is policed at Layer 3 prior to traversing the campus backbone. A Layer 2 switch, however, lacks such intelligent mechanisms and, consequently, offers less control over what traffic reaches the campus backbone.

Figure 19 Layer 2 Design

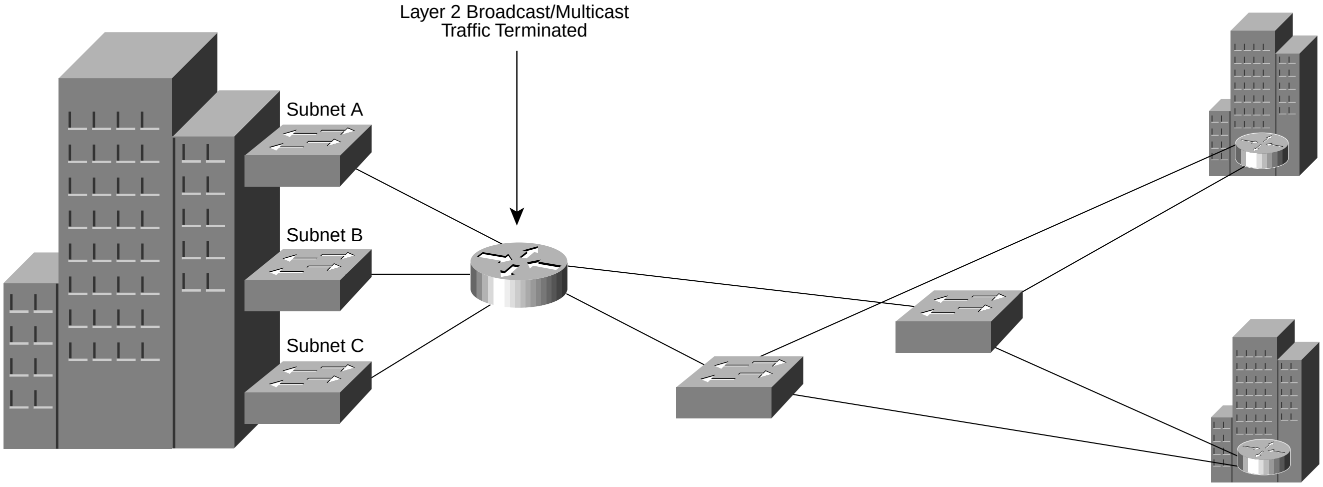

Virtual LANs (VLANs) can address these issues of traffic management in large Layer 2 switched fabrics. With VLANs, multiple broadcast domains can be defined within the Layer 2 fabric. This design, in turn (see Figure 20), will help to contain broadcasts and provide scalability mechanisms for the switched fabric.

Figure 20 Layer 2 Design

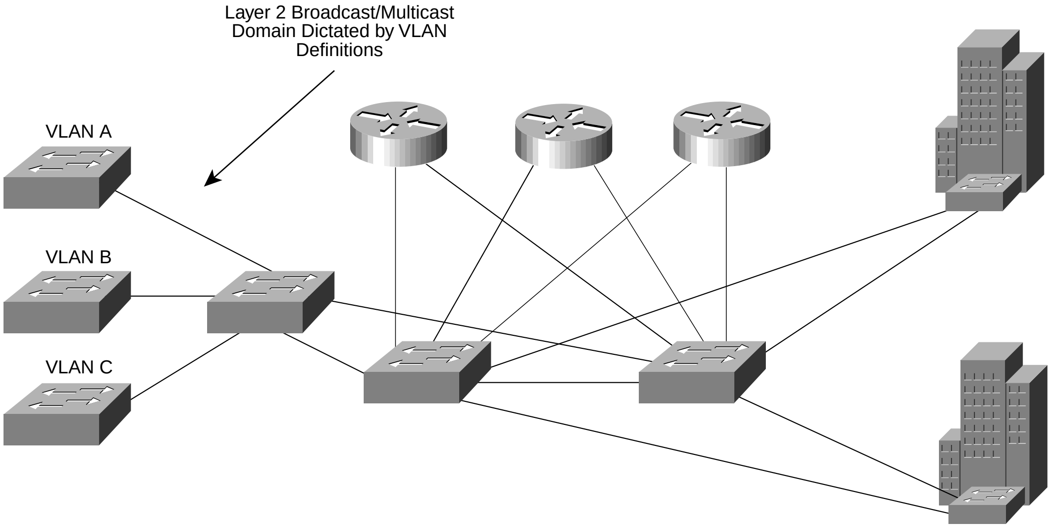

As illustrated in Figure 20, VLANs allow for multiple Layer 2 networks to be defined within the campus, thereby decreasing the scope of Layer 2 broadcasts and multicasts.

While the Layer 2 VLAN design appears an attractive solution, there are disadvantages from the perspective of campus backbone design that must be addressed. Most notable is the performance impact on the campus backbone, the very part of the network where we are trying to alleviate congestion. The distributed Layer 2/Layer 3 design (Figure 21) defines the backbone as two core switches with backbone routers on the periphery. In this design, wiring closet access to the campus backbone is established through the routers. The Layer 2 VLAN network, on the other hand, defines the campus backbone as two core switches and centralized route servers. In this design, the wiring closets (VLANs) connect to the campus backbone via the core switches rather than the routers (see Figure 22).

Figure 21 Campus Backbone Definitions: Distributed Layer 2/Layer 3

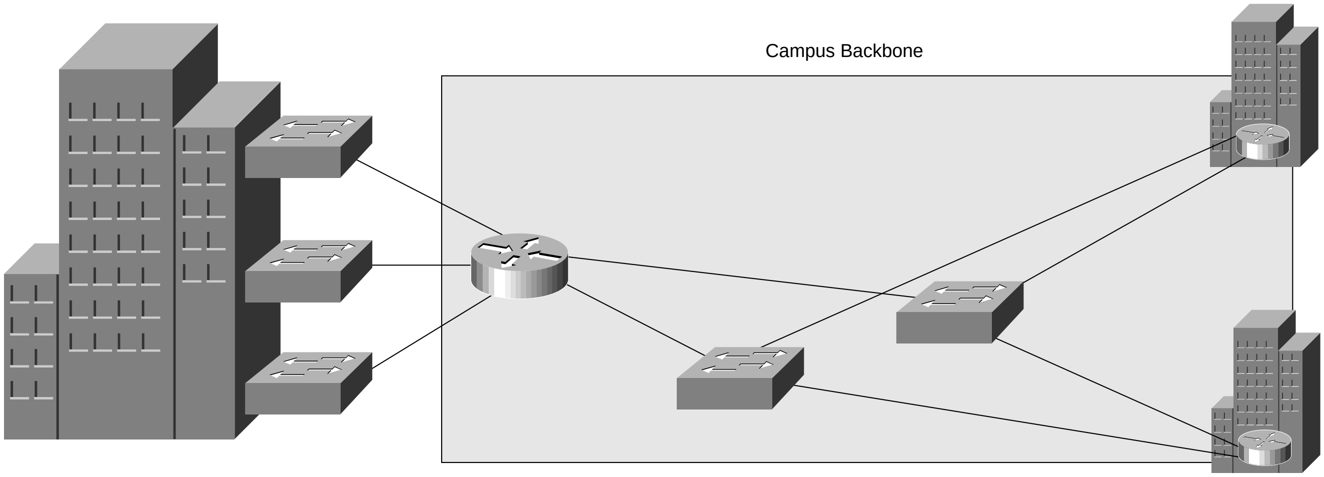

Figure 22 Campus Backbone Definitions: Layer 2 VLAN Designs

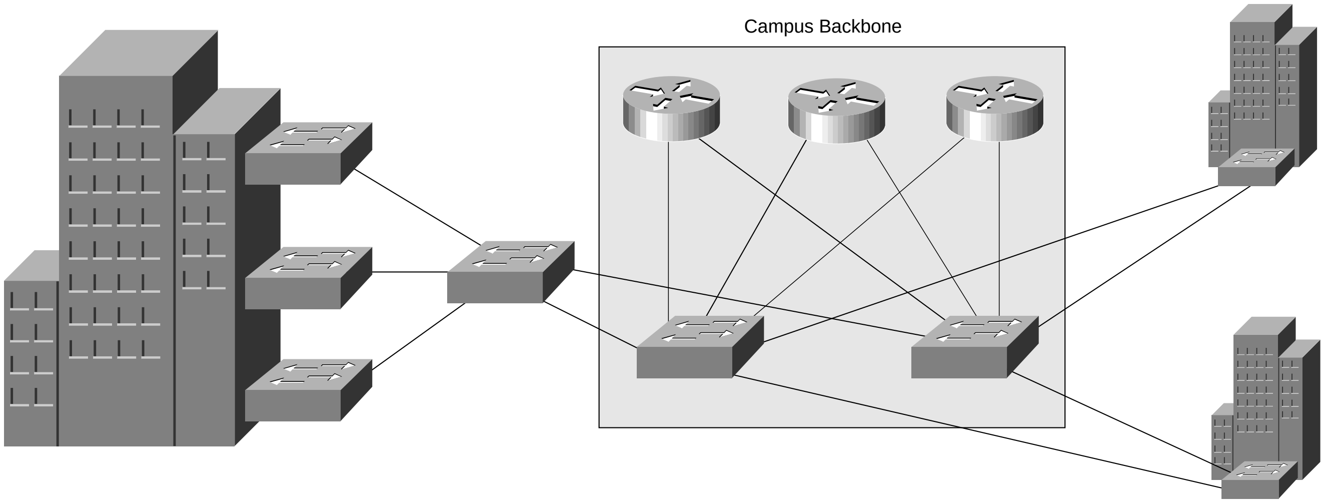

It follows from the preceding definitions and designs that backbone bandwidth utilization will differ. Consider Figures 23 and 24. In Figure 23, intersubnet traffic is routed at router A in the distributed Layer 2/Layer 3 design and does not traverse the campus backbone. In the Layer 2 VLAN design, however, intersubnet (or inter-VLAN) traffic does traverse the campus backbone and is subsequently routed by one of the centralized route servers.

Figure 23   Intersubnet Traffic Overhead in Layer 2/3 vs. Layer 2 Environments

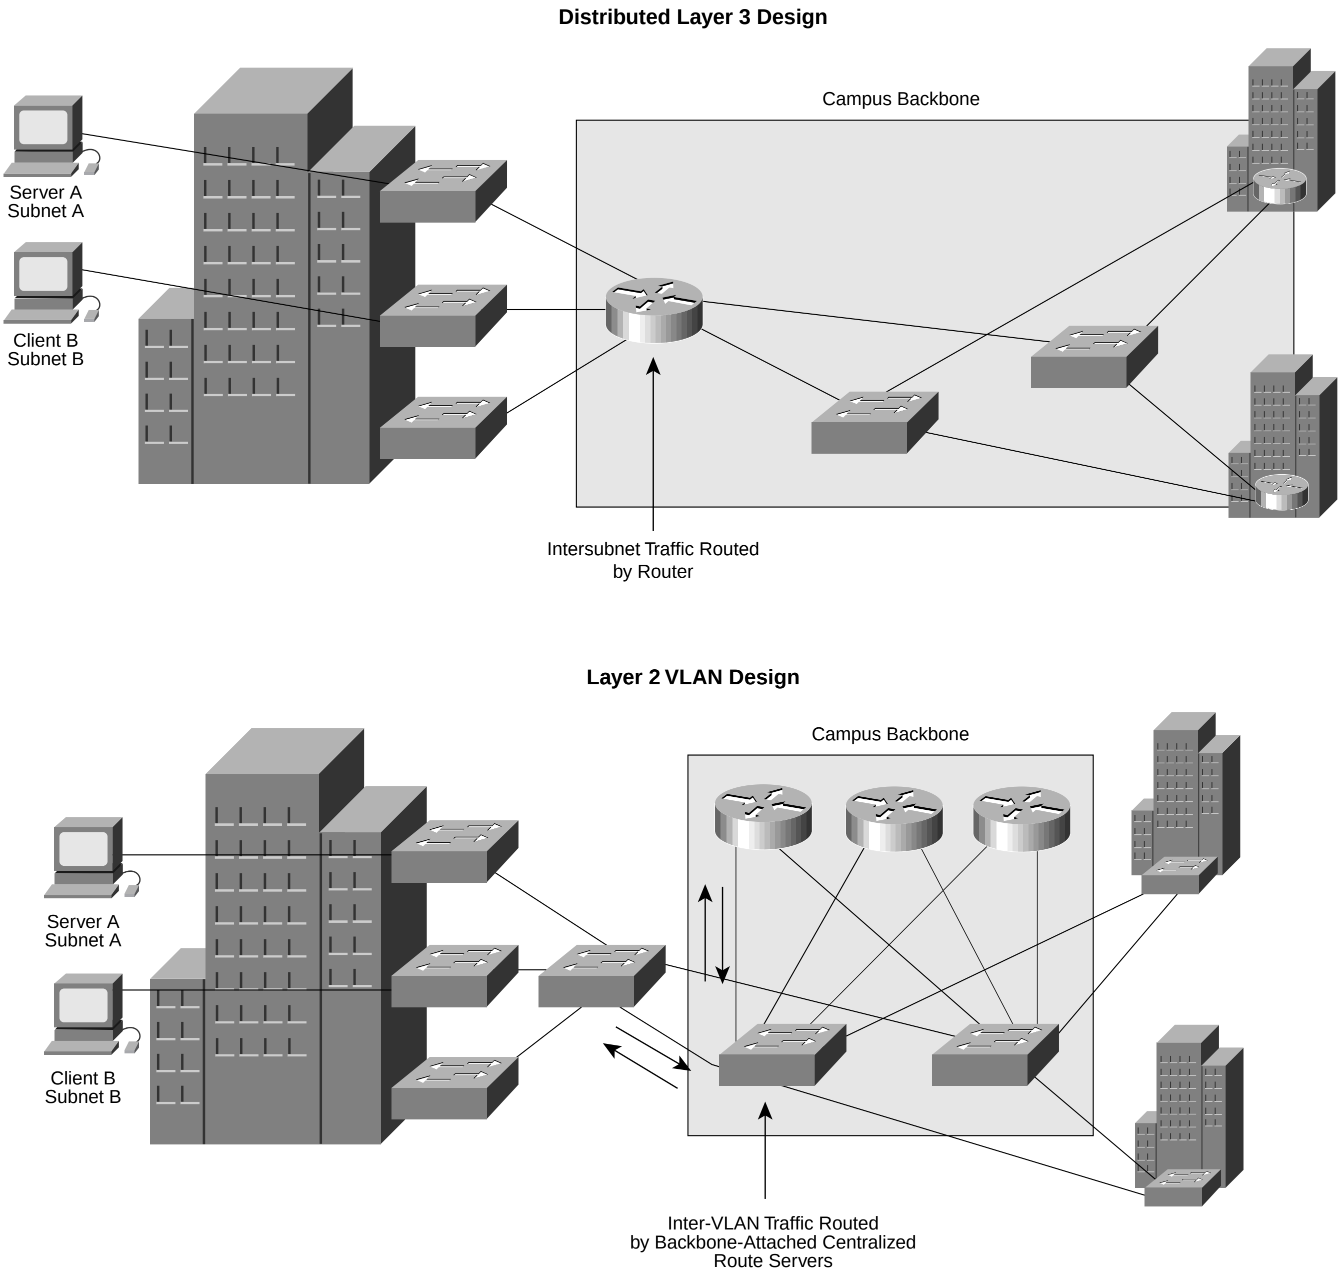

And Figure 24 details intrasubnet traffic. Here again, traffic never traverses the campus backbone in the distributed Layer 2/3 design but it can potentially in the VLAN-based Layer 2 design, depending on whether or not VLANs are defined across the campus.

Figure 24 Intrasubnet Traffic Overhead in Layer 2/3 vs. Layer 2 VLAN Environments

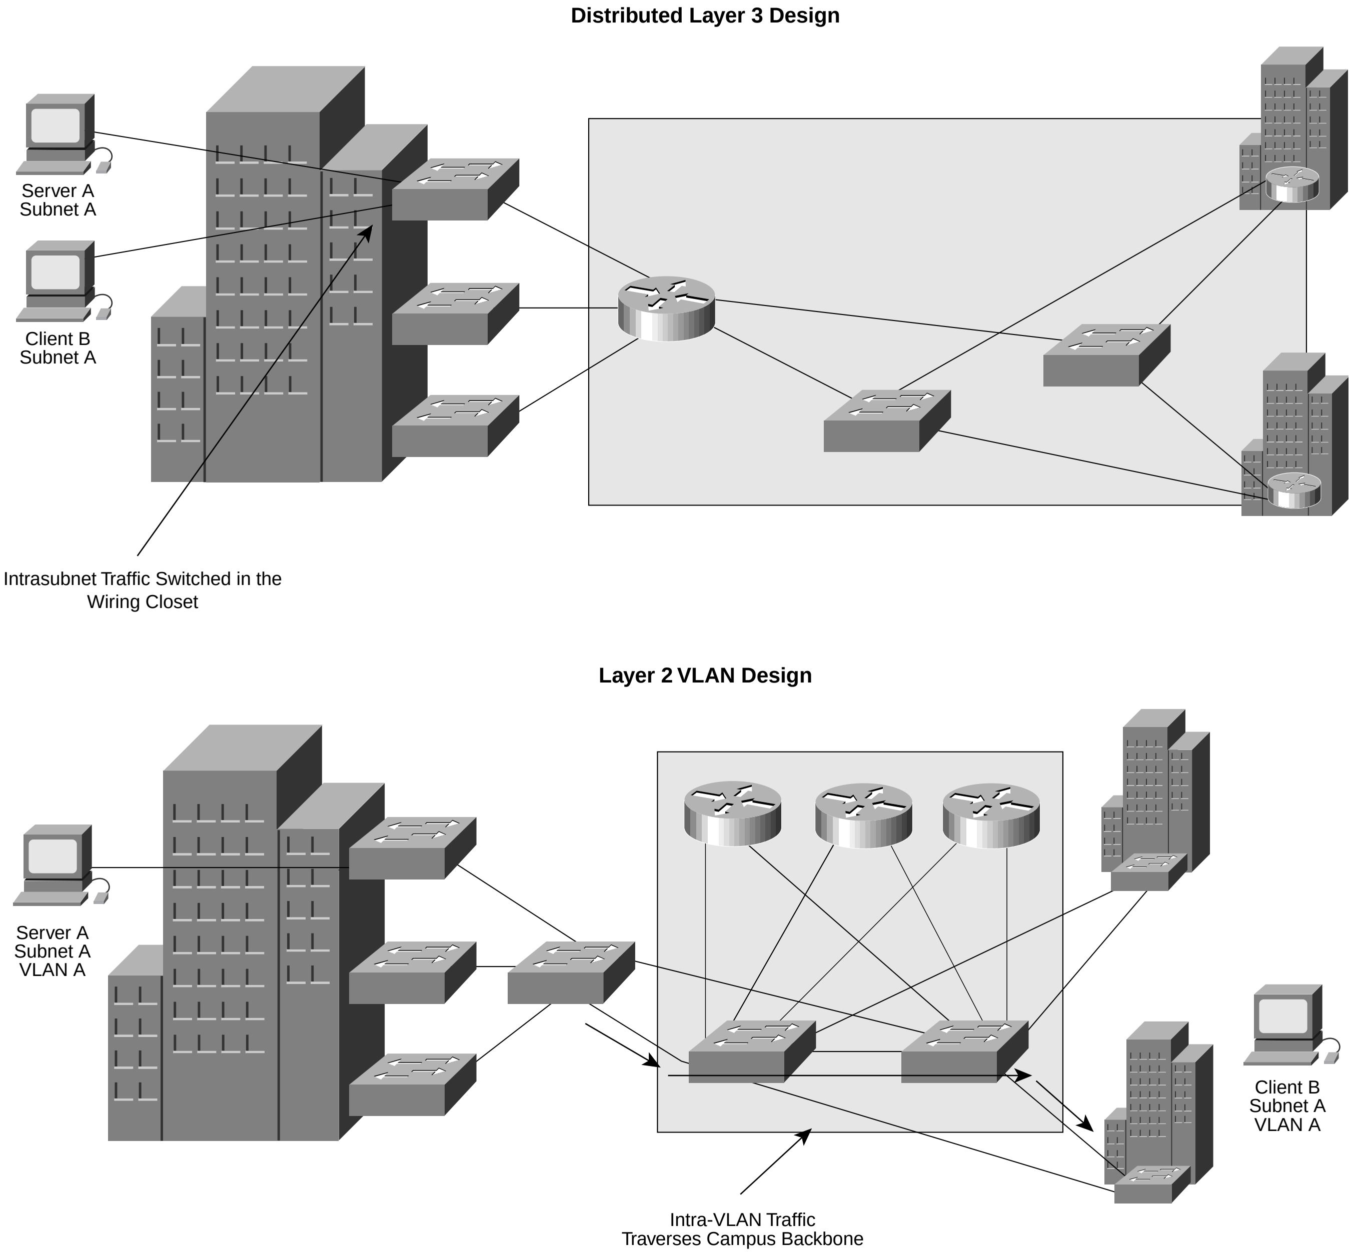

Distributed Layer 3 Design

Server A
Subnet A
Client B
Subnet A
Intrasubnet Traffic Switched in the Wiring Closet

Layer 2 VLAN Design

Server A
Subnet A
VLAN A
Client B
Subnet A
VLAN A
Intra-VLAN Traffic Traverses Campus Backbone

It should be clear from the preceding discussion that migrating the dual-homed FDDI backbone to a Layer 2 VLAN design can be problematic. In particular, the Layer 2 VLAN design introduces a new traffic model within the campus that could potentially undermine backbone scalability. Depending on how routing is implemented in the design, more traffic may actually have to go to the backbone in order for it to be routed. Also, if cross-campus VLANs are deployed, even more traffic must traverse the backbone. The distributed Layer 2/Layer 3 design, on the other hand, establishes a more uniform traffic model and offers greater control over backbone traffic patterns. This scenario, in turn, allows for greater backbone scalability and capacity.

Reengineering the Campus Backbone—Summary

When evaluating technologies and considering reengineering the campus backbone, consider the following technical objectives:

- Alleviate backbone congestion
- Ensure a fault-tolerant infrastructure
- Provide a migration path to higher-speed networking

Consider also the network hierarchy and the critical role that routing plays in traffic management across the network. Do not simply solve one problem by creating another one elsewhere. Lastly, consider all this with economics in mind. Ultimately reengineering will involve an investment in technology and a cash outlay. As a result, it is important that the reengineering effort not only be technically sound but also cost effective.

FDDI Solutions

Design Options

Based on the proceeding discussion, there are three principal solutions for reengineering the dual-homed campus backbone using FDDI technologies.

- FDDI ring segmentation
- FDDI ring redundancy
- FDDI switching

Layer 3 FDDI Ring Segmentation

FDDI ring segmentation is based on segmenting the dual-homed FDDI backbone into two separate Layer 3 segments, or essentially creating two separate dual-homed backbones connected together with a router dual-homed to both segments (see Figure 25).

Figure 25 Layer 3 FDDI Ring Segmentation—Logical Design

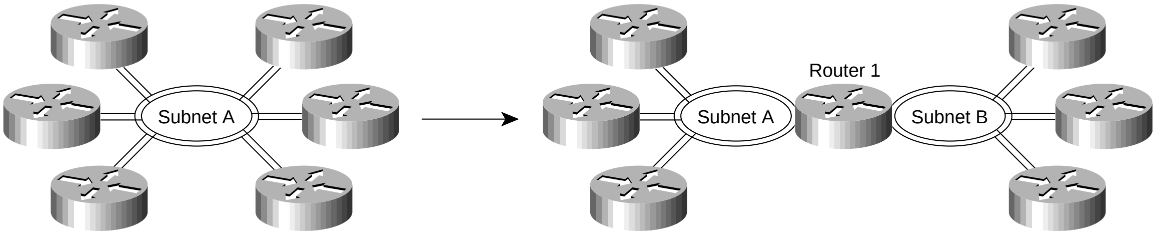

As illustrated, ring segmentation has reduced the number of devices on any given FDDI ring and has thus helped to alleviate congestion. Depending on how segmentation has occurred, traffic can be better managed across the backbone. Implicit in this statement is a sufficient understanding of backbone traffic patterns to know where segmentation should occur. This level of segmentation can be a challenge, especially when there are servers on the backbone. Improper backbone segmentation may reduce congestion but may introduce other inefficiencies (for example, extra unnecessary router hops).

The primary disadvantage from a design standpoint is that there is a single point of failure in the backbone with Router 1. If Router 1 goes down, backbone connectivity will be partially disrupted. To address this problem, a second router could be deployed to connect the two segments together (see Figure 26). The additional router not only provides hardware-level fault-tolerance, it also provides Layer 3 load-balancing and fault tolerance. With two routers connecting the rings, traffic can be load balanced across the two available Layer 3 paths. And if one path becomes unavailable, the Layer 3 routing technologies dynamically reconverge based on the available routes (see Figures 26 and 27).

Figure 26 Layer 3 Load Balancing

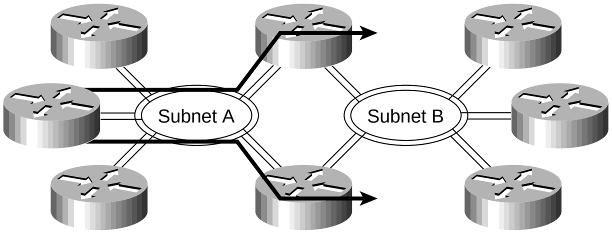

Figure 27 Layer 3 Fault Tolerance

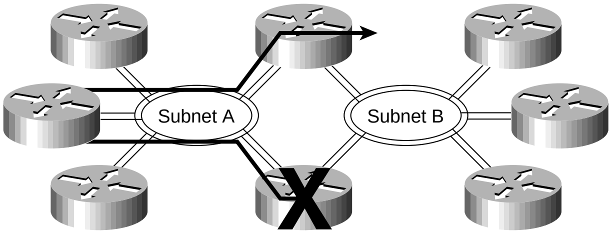

If done properly, FDDI ring segmentation can successfully provide better bandwidth management and fault tolerance compared with the dual-homed design. This design has two principal disadvantages: the cost and no migration path to higher-speed networking. This design will most likely require two new FDDI-equipped routers, a considerable investment. Moreover, it is unclear to what extent segmentation will solve backbone congestion problems. Eventually it will come down to the fact that 100 Mbps simply is not enough. And at that point, what action should be taken? FDDI appears to have no future beyond 100 Mbps, leaving an opening for ATM or high-speed Ethernet offerings.

Layer 3 FDDI Ring Redundancy

Another option for reengineering the campus backbone with FDDI technologies involves ring redundancy. In this design, a second parallel FDDI network is implemented to provide dual connections for the backbone routers. Unlike the previous design where the dual-home ring was segmented into two separate Layer 3-defined segments, this design contains two parallel Layer 3 paths for each router (see Figure 28).

Figure 28 Layer 3 FDDI Ring Redundancy

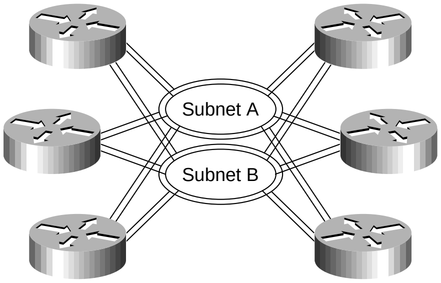

Based on intelligent Layer 3 router algorithms, each router can load-balance traffic across the two FDDI rings (see Figure 29). Additionally, the Layer 3 intelligence provides for fast convergence in the event of route failure (see Figure 30).

Figure 29 Layer 3 Load Balancing

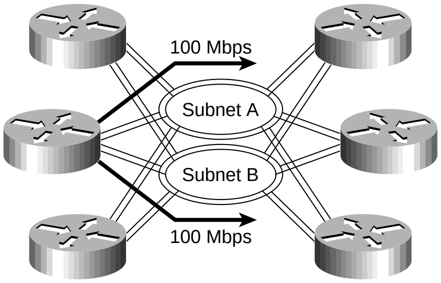

Figure 30 Layer 3 Fault Tolerance

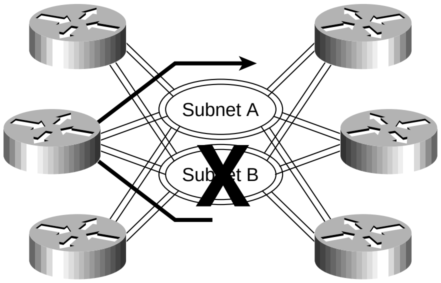

The Layer 3 FDDI ring redundant design clearly adds more bandwidth (a second FDDI ring to be precise) to the backbone and provides an extremely resilient infrastructure. The primary disadvantage is cost. All routers would need to be outfitted with an additional FDDI interface, and some additional investment in FDDI concentrators may be necessary.

FDDI Switching
The third option for reengineering with FDDI is FDDI switching, a replacement technology for FDDI concentrators, similar to what Ethernet switches are to Ethernet repeaters. The primary goal of the switching technology is to provide more efficient network access. From that perspective, FDDI switching is a very effective solution.

The weak point with FDDI switching is translational bridging. Translational bridging is the technology used when two devices connected to different media, yet in the same bridge domain, communicate. Consider Figure 31. Server A is connected to the network via FDDI, while client A is connected to the network via Ethernet.

Figure 31 Mixed-Media Bridge Domains—Ethernet and FDDI

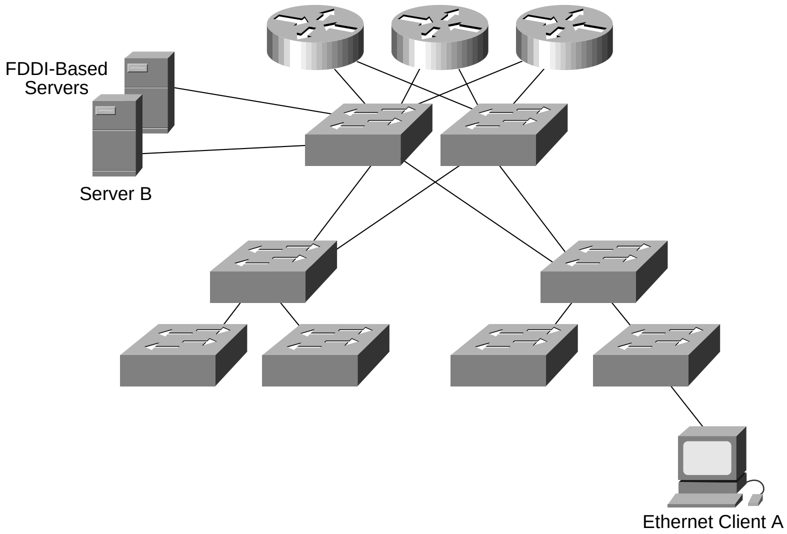

The challenge here is for the switch to successfully translate the Ethernet frames from the client to FDDI frames for the server and vice versa. This challenge can be particularly problematic. Consider a Novell NetWare environment (see Figure 32), where four possible Ethernet encapsulations exist (Ethernet_II, Ethernet_802.2, Ethernet_802.3, and Ethernet_SNAP) and two FDDI encapsulations (FDDI_II and FDDI_SNAP). A poor translational bridging implementation can result in encapsulation errors. FDDI-bound Ethenet_SNAP frames may end up being Ethernet_II frames after being translationally bridged back to Ethernet.

Figure 32 Translational Bridging: Novell NetWare Environments

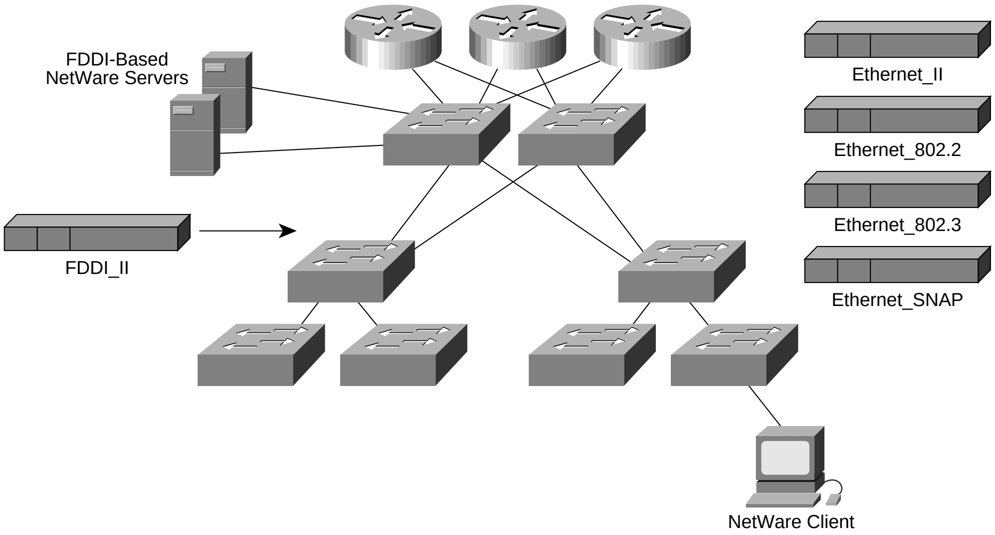

Besides the relative complexity of translational bridging and the possibility of error, translational bridging is a process-intensive operation. Within the context of switching, which is intended to provide optimized performance, translational bridging can undermine overall performance. For these reasons, it is recommended to address the mixed-media communication at Layer 3, thereby circumventing Layer 2 pitfalls.

This said, FDDI switching can be used to replace FDDI concentrators to build a distributed Layer 2/Layer 3 design. The only assumption is that neither servers nor downstream Ethernet switches equipped with FDDI uplinks will be connected to the FDDI switch. In Figure 33, the two FDDI concentrators have been removed, and an FDDI switch is put in their place.

Figure 33 FDDI Switching

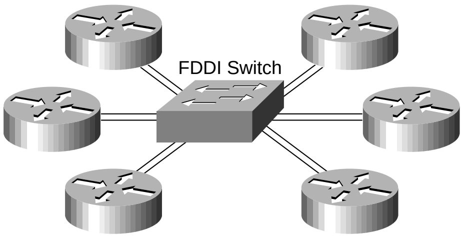

This design will succeed in providing better per-device access to the FDDI network. The primary drawback is that a certain level of fault tolerance has been sacrificed. The FDDI switch now represents a single point of failure in the network. If the FDDI switch goes down, connectivity across the backbone is lost.

To address this problem, there are two options. The first option involves buying a second FDDI switch and essentially dual-homing each router to the two FDDI switches. However, while dual-homing is straightforward with concentrators, it can be problematic with switches. With switches, forwarding decisions are based on entries in the switching table. If a router is connected to two different switches, each switch has an entry in its switching table for the router's FDDI Media Access Control (MAC) address. Assuming that one path is defined as the primary path and the other path is in standby mode, the challenge is for two switches and all connected devices to be able to withstand fault and communicate.

To illustrate this challenge, consider a frame going from router A to router B across the primary path (see Figure 34). The frame leaves router A and goes to switch 1. Switch 1 checks its switching table, sees an entry for router B, and attempts to forward the frame out through the appropriate port. Now suppose the physical connection to router B breaks. Switch 1 needs to detect the link failure and potentially delete the entry in the switching table. In addition, switch 2 needs to know that it is the primary path for communication to and from router B. It follows that all connected devices also need to know that they should use their secondary paths (via switch B) for communications to router B.

Figure 34 Layer 2 MAC Address Management

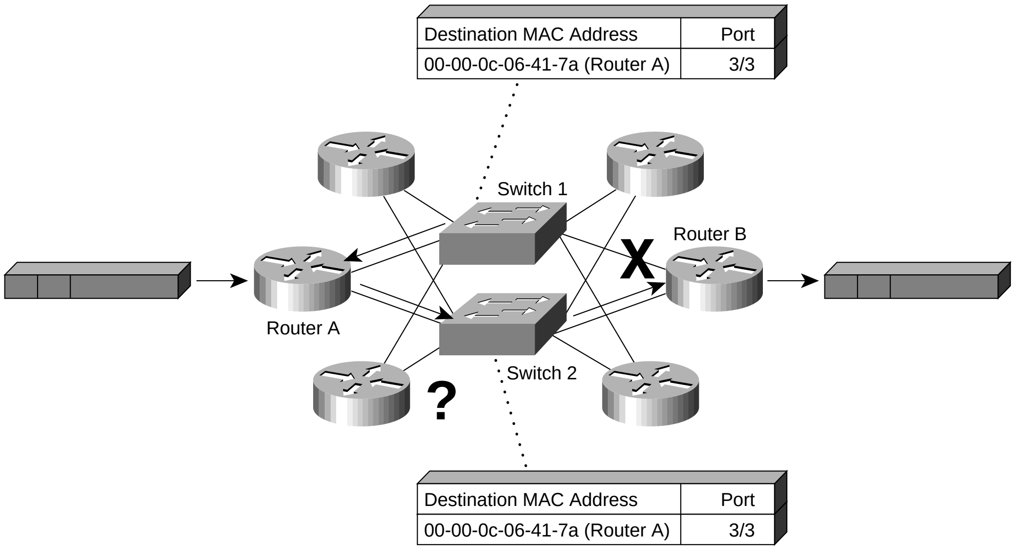

Destination MAC Address	Port
00-00-0c-06-41-7a (Router A)	3/3

Destination MAC Address	Port
00-00-0c-06-41-7a (Router A)	3/3

One workaround for this problem is to connect the two FDDI switches together and run spanning tree between the switches and the dual-homed routers. Of course the drawback is that the backbone is running both Layer 3 routing processes as well as spanning tree.

The safest way to solve this problem is to add a second FDDI interface on the each router. This solution achieves fault tolerance using Layer 3 routing technologies and also enables load balancing over both links (see Figure 34).

Figure 35 Layer 3 FDDI Switching

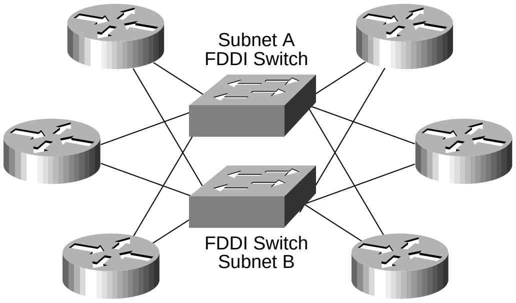

Note that in Figure 35, each of the FDDI interfaces on the ingress routers can be either be single-attached or dual-attached to the respective FDDI switches. Clearly, the dual-attached connection provides an added degree of physical-layer redundancy.

The principal disadvantages of this design are cost, and again, a limited migration path to higher-speed networking. Yes, this design will provide a 200Mbps fabric but compared with a similar design using Ethernet technologies, the two FDDI switch/two FDDI interface processor design is extremely costly.

FDDI Solutions—Summary

The three solutions presented—ring segmentation, ring redundancy, and FDDI switching—can all help to better manage bandwidth and traffic across the campus backbone. And each can deliver acceptable fault tolerance. The common disadvantages center on cost and migration capabilities. Relative to other available technologies, FDDI is extremely expensive, especially when trying to deliver fault tolerance. In addition, it is unclear how long the backbone can remain at 100 Mbps. If the answer lies somewhere in the near future, additional investments in FDDI may not be the most cost effective. In these cases migrating either to ATM or Fast Ethernet would be more prudent.

ATM Solutions

Technology Overview/Design Issues

Despite its slow-moving standards body, ATM has managed to deliver a set of standards-based implementations for data networking—permanent virtual circuits (PVCs), switched virtual circuits (SVCs), LAN Emulation (LANE), and Classical IP over ATM (RFC 1577). Each of these implementations permits traditional LAN data (such as IP and IPX) to run on top of an ATM fabric. Their primary benefit is to enable high-speed connectivity between traditional LAN devices. The typical ATM connection, for example, runs at 155 Mbps (OC-3). Higher connections are also available, albeit in limited fashion—OC-12 (622 Mbps) and OC-48 (2.488 Gbps) are the two principal offerings today.

But bandwidth is not necessarily ATM's strongest suit. ATM is based on a connection-oriented/circuit-switched technology. For two devices to communicate across an ATM fabric, a connection must first be established. The connection can be permanent (such as PVCs), or it can be set up dynamically using SVCs. This scenario leads to one of ATM's greatest selling feature—Quality of Service (QOS). With ATM's QOS features, devices can request varying degrees of quality (expressed by numerous different transfer attributes such as cell transfer rate and minimum delay) from the ATM fabric. The vision is to offer a network service that offers the greatest flexibility and integration compared with the needs of connected devices.

QOS is undoubtedly a major contributor to ATM's lure. However, while QOS is part of the ATM technology, it is absent from current ATM implementations for data networking. Consider an IP/IPX backbone using ATM as the underlying technology. There are three principal ways of implementing ATM for this design—PVCs, SVCs, or LANE (see Figures 32, 33, and 34). Classical IP is not an option in this case, because the backbone carries multiprotocol traffic.

Figure 36 Implementing ATM for Data Networking Using PVCs

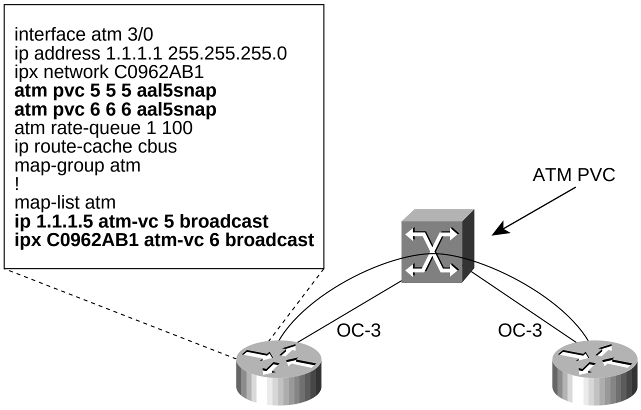

```
interface atm 3/0
ip address 1.1.1.1 255.255.255.0
ipx network C0962AB1
atm pvc 5 5 5 aal5snap
atm pvc 6 6 6 aal5snap
atm rate-queue 1 100
ip route-cache cbus
map-group atm
!
map-list atm
ip 1.1.1.5 atm-vc 5 broadcast
ipx C0962AB1 atm-vc 6 broadcast
```

Figure 37 Implementing ATM for Data Networking Using SVCs

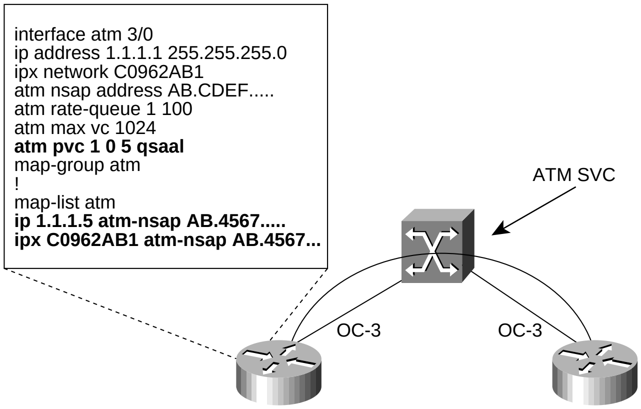

```
interface atm 3/0
ip address 1.1.1.1 255.255.255.0
ipx network C0962AB1
atm nsap address AB.CDEF.....
atm rate-queue 1 100
atm max vc 1024
atm pvc 1 0 5 qsaal
map-group atm
!
map-list atm
ip 1.1.1.5 atm-nsap AB.4567.....
ipx C0962AB1 atm-nsap AB.4567...
```

Figure 38 Implementing ATM for Data Networking Using LANE

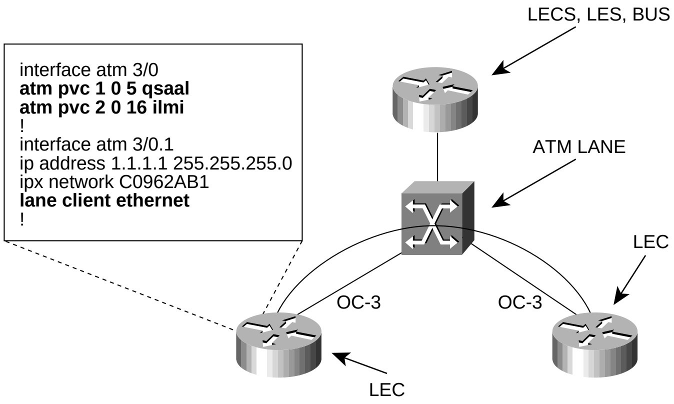

```
interface atm 3/0
atm pvc 1 0 5 qsaal
atm pvc 2 0 16 ilmi
!
interface atm 3/0.1
ip address 1.1.1.1 255.255.255.0
ipx network C0962AB1
lane client ethernet
!
```

In each of these cases, QOS mechanisms through the ATM fabric are limited. With PVCs, the connection is static and thus offers no QOS mechanisms. With SVCs, QOS can be negotiated at setup, but because SVCs map to Layer 3 routes, QOS is based on network-layer information, not application-layer information. Resource Reservation Protocol (RSVP)-to-ATM mappings offer tighter integration, albeit only for IP-based applications. And lastly, LANE also has no current notion of QOS. With LANE 1.0, data direct virtual circuits (VCs) between LANE clients cannot specify QOS preferences at call setup. LANE 2.0 will provide LAN Emulation Client (LEC)-driven QOS, but it is a future offering.

To summarize, while ATM reads like the network panacea, today's implementations of ATM are not. Ultimately from the perspective of designing a distributed Layer 2/Layer 3 design today, ATM will be best implemented using PVCs or SVCs. Today this scenario will prove the easiest way to run multiprotocol traffic across the backbone and offer the greatest scalability. As for LANE, besides lacking QOS, its usefulness is in question given a backbone architecture consisting of Layer 3 engines at the ingress points. At the core of LANE is a service that maps traditional LAN-based addresses (for example, Ethernet MAC addresses) to ATM addresses. This is essentially a bridging function handled by the combined efforts of the broadcast and unknown server (BUS) and the LAN Emulation Server (LES). To enable LANE to scale to large networks, LANE uses the concept of an emulated LAN (ELAN) to define a group of devices within the domain of a LES/BUS pair. Traffic within an ELAN will be bridged while traffic between ELANs must be routed (see Figure 39).

Figure 39 Intra-ELAN and Inter-ELAN Communication

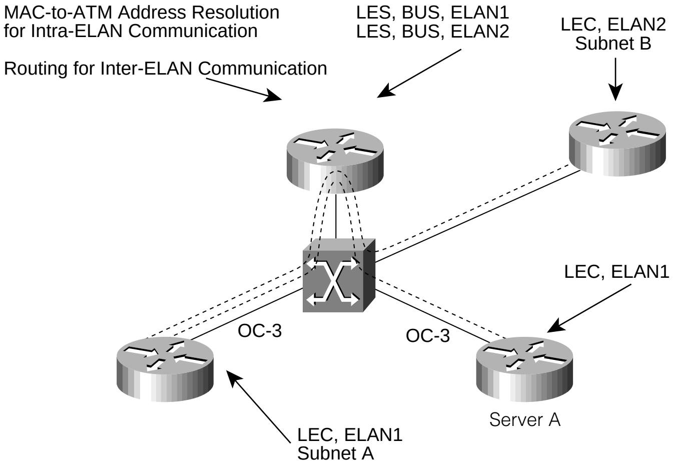
MAC-to-ATM Address Resolution
for Intra-ELAN Communication
Routing for Inter-ELAN Communication
LES, BUS, ELAN1
LES, BUS, ELAN2
LEC, ELAN2
Subnet B
LEC, ELAN1
OC-3
OC-3
Server A
LEC, ELAN1
Subnet A

Based on LANE operation, there is little need for LANE's services in the campus backbone. In fact, especially if the backbone carries multicast traffic (processed by the ELAN's BUS), backbone scalability becomes a product of BUS scalability, which is clearly suboptimal (see Figure 40).

Figure 40 Multicast Processing in LANE-Based Campus Backbone

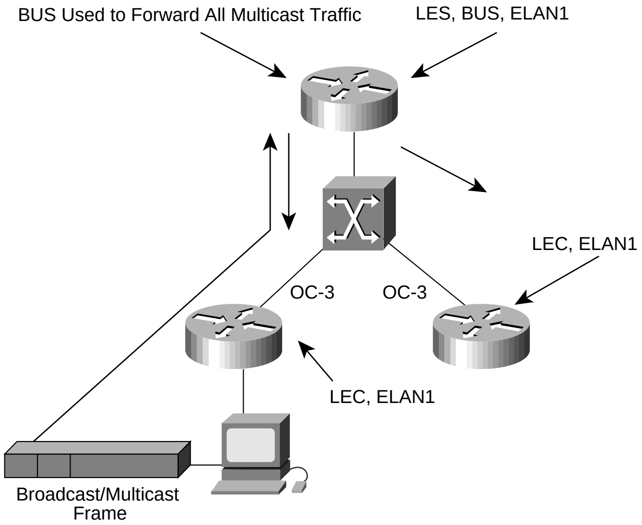
BUS Used to Forward All Multicast Traffic
LES, BUS, ELAN1
LEC, ELAN1
OC-3
OC-3
LEC, ELAN1
Broadcast/Multicast Frame

As illustrated above, in a LANE environment, all multicast frames are processed by the BUS. Additionally, the ELAN1 is sent to all LECs within the ELAN. From a traffic management perspective, such an operation can be undesirable in the campus backbone.

ATM Backbone Designs

Using PVCs or SVCs, the dual-homed FDDI backbone could be converted to the ATM-based design shown in Figure 41.

Figure 41 ATM-Based Campus Backbone

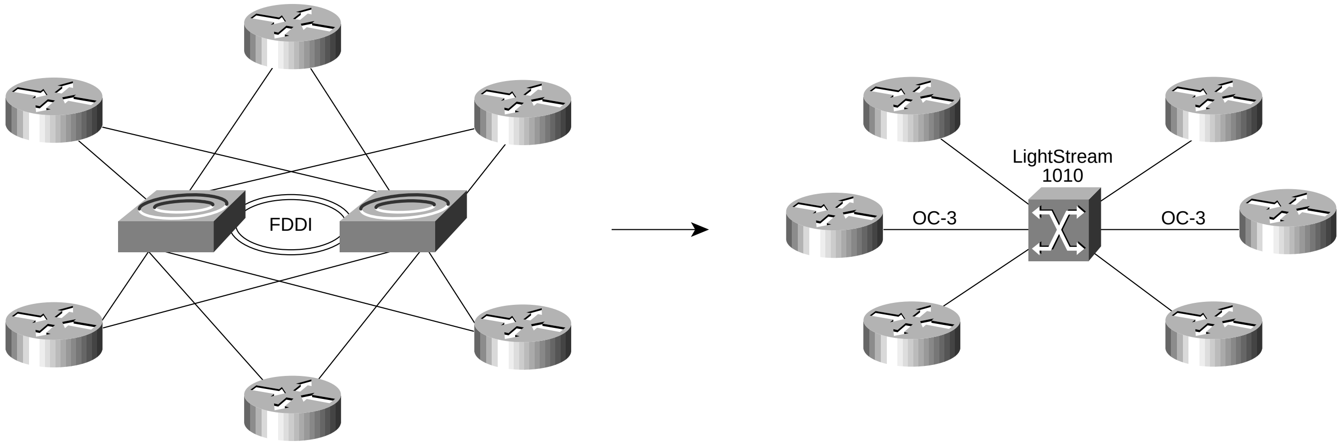

This design is based on replacing the FDDI Interface Processors (FIPs) on the ingress routers with ATM Interface Processors (AIPs) and by replacing the FDDI concentrators with a single LightStream® 1010. From the perspective of alleviating backbone congestion, this design is very effective. ATM's circuit-switched technology coupled with 155-Mbps connections provides a scalable, efficient Layer 2 access method. The primary drawback in this design, of course, is fault tolerance. Here, the LightStream 1010 at the core of the network represents a single point of failure in the backbone. The best way to address this problem and add fault tolerance is to add a redundant Layer 3 path to the backbone. In the case of ATM, this is done by adding a second AIP to the ingress routers (see Figure 42) and adding a second LightStream 1010.

Figure 42 Layer 3 Fault-Tolerant ATM Backbone Design

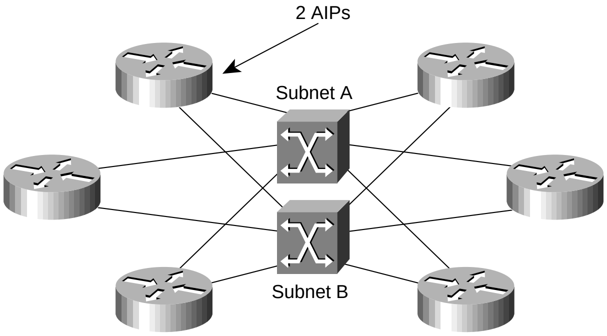

As shown in Figure 42, adding a second LightStream 1010 and additional AIPs will surely deliver effective fault-tolerance. This design will also help to add additional bandwidth to the backbone. Each router has two OC3 connections, which translates to an aggregrate bandwidth of 310 Mbps half duplex (620 full duplex). The central issue in this design, however, is cost. Adding AIPs and a second LightStream 1010 to the backbone requires a considerable monetary investment. At this point, the value of ATM must be questioned seriously.

ATM Backbone Solutions—Summary

Simply put, in many cases, ATM technology is outweighed by its associated costs, especially when architecting fault tolerance. As we will discuss in the next section, a comparable Fast Ethernet design would be considerably more affordable.

In addition to the hardware investments, two other issues weigh against ATM—administrative overhead and performance. ATM introduces a new paradigm in data communication. Besides being a circuit-switch technology, which in itself is a novel concept in LAN-based networking, ATM uses a fixed-length cell for its transmission unit as opposed to the variable-length frame used in FDDI or Ethernet, for example. As a result, there is a legitimate "ramp-up" period from an administrative/management perspective. Network managers and administrators alike will have to undergo an education process in order to truly understand ATM and effectively monitor and troubleshoot network operation. Undoubtedly, this education process will also entail a monetary investment in both educational tools (books, training classes) as well as analysis tools (protocol analyzers, traffic analyzers).

Additionally, there is ATM performance, often referred to as the ATM cell tax, which refers to the processing overhead associated with converting traditional LAN-based frames to ATM cells for transmission through the ATM fabric. In Figure 39, for example, all traffic that traverses the backbone must undergo this process. And typically this conversion results in a hit in forwarding performance—hence, the "cell tax."

Figure 43 The Cell Tax—Frame-to-Cell Conversion

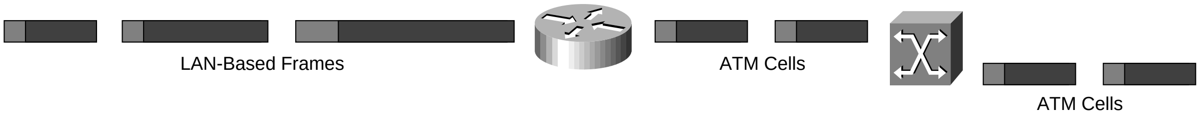

Put simply, migrating to ATM is a nontrivial process. A decision to migrate to ATM should not be reached without fully analyzing all the pertinent issues:

- Implementation options
- Cost
- Performance
- Training and retooling

Fast Ethernet Solutions

Technology Overview/Design Options

The last migration option for the dual-homed FDDI backbone employs Fast Ethernet for the underlying Layer 2 transport mechanism. Fast Ethernet can be implemented in either a shared or a switched fashion. Shared implementations provide a single, shared, Fast Ethernet segment that ports contend for. Switched implementations, on the other hand, provide multiple Fast Ethernet segments, one per port. It follows that Fast Ethernet performance is optimized when switching is employed to eliminate network contention.

As for network transmission, Fast Ethernet can operate in either half-duplex or full-duplex mode. Full-duplex operation is based on point-to-point Fast Ethernet connections and enables two-way, 100-Mbps transmission. Because full-duplex operation grants two separate transmission paths (one for transmitting, one for receiving), it is ideal within the backbone where two-way transactions are common.

Lastly, Fast Ethernet can run over either UTP, single-mode fiber, or multimode fiber cabling. The distance limitations for the respective media are detailed in Table 1.

Table 1 Fast Ethernet Distance Limitations

Cable Type	Distance Limitation	
—	Full-Duplex Mode	Half-Duplex Mode
—	—	—
Unshielded Twisted-Pair	100 meters	100 meters
Multimode Fiber	2 km	400 meters

Table 1 Fast Ethernet Distance Limitations

Cable Type	Distance Limitation	
Single-Mode Fiber	Greater than 2 km	Greater than 2 km

UTP is clearly the most affordable medium and aids in Fast Ethernet's affordability, but it is not necessarily a requirement. In fact, from the perspective of migrating from FDDI, leveraging the existing multimode or single-mode fiber for Fast Ethernet may make more sense.

As a campus backbone technology, switched Fast Ethernet has several advantages:

- Affordability
- Ease of use
- Performance

Affordability and ease of use, although essentially nontechnical in nature, are extremely important benefits. This fact not only makes initial investments manageable, but it permits an affordable solution set for fault tolerance as well as an affordable migration path to higher-speed networking.

Also working in Fast Ethernet's favor is ease of use. Particularly in environments that currently employ 10-Mbps Ethernet technologies, introducing Fast Ethernet to the campus is a relatively trivial operation. The only difference between the two technologies is speed—framing is the same and the underlying transmission properties remain the same. Most likely additional investment will be required for Fast Ethernet protocol analysis tools; today such tools are both available and affordable. A Fast Ethernet protocol analyzer tool is, on average, 50 percent cheaper than a comparable ATM protocol analysis tool.

The last major benefit of Fast Ethernet is performance. Fast Ethernet supports full-duplex operation on point-to-point links. This capability in turn permits an aggregate bandwidth of 200 Mbps (100 Mbps each way) across a Fast Ethernet backbone connection. Figure 44 details full-duplex operation and its advantage in the campus backbone.

Figure 44 Fast Ethernet Full-Duplex Operation in the Campus Backbone

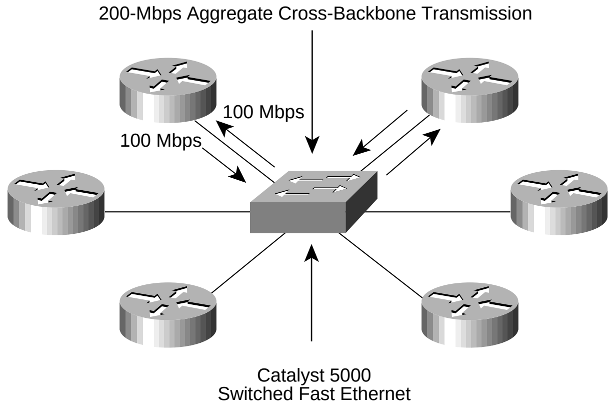

As we will illustrate in the designs that follow, this feature greatly enhances backbone bandwidth management. Also, unlike ATM, there is no associated penalty for moving from Ethernet to Fast Ethernet. The connecting device needs only to perform a speed-matching operation, unlike ATM, which must also perform a frame-to-cell conversion process.

Fast Ethernet Backbone Designs

Based on the proceeding discussion, Figure 45 details a Fast Ethernet-based migration solution for a dual-homed FDDI backbone.

Figure 45 Fast Ethernet Backbone Design

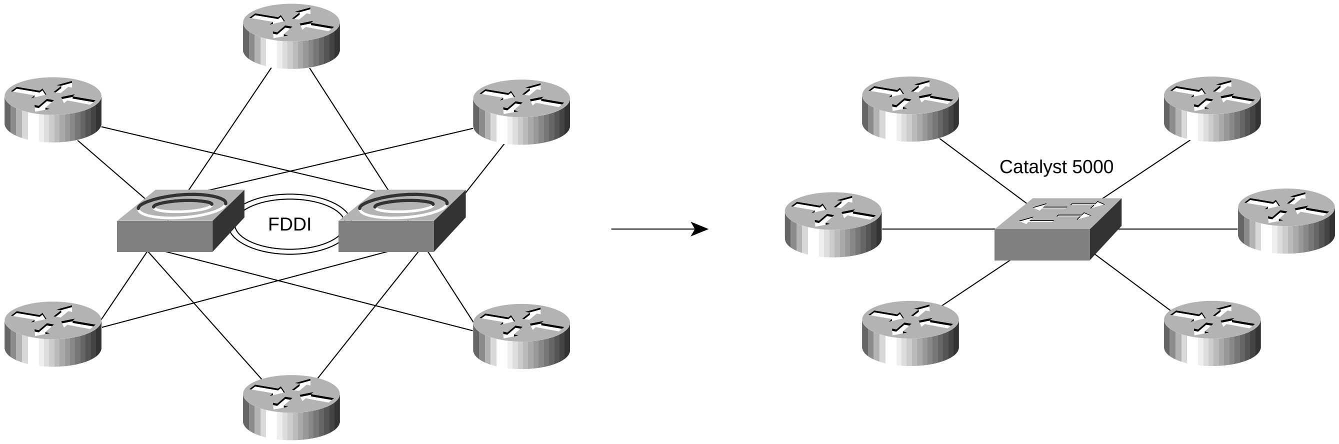

This design is based on replacing the FIPs on the ingress routers with Fast Ethernet Interface Processors (FEIPs) or VIP-2 Fast Ethernet line cards and replacing the FDDI concentrators with a single Catalyst™ 5000. From the perspective of alleviating backbone congestion, this design will prove effective. The switched Fast Ethernet logic in the Catalyst 5000 coupled with full-duplex connections to the ingress routers provides a scalable, efficient Layer 2 switched fabric. From a cabling perspective, this migration could include a migration to UTP cabling or could leverage the cabling infrastructure from the FDDI network to deliver a fiber-based implementation.

The primary drawback in this particular design is fault tolerance. Here, the Catalyst 5000 at the core of the network represents a single point of failure in the backbone. Like the FDDI switching and ATM designs, the best way to add fault tolerance is by adding a redundant Layer 3 path to the backbone. In the case of Fast Ethernet, this is done by adding a second FEIP or a second line card to the VIP-2 paddle on the ingress routers (see Figure 46) and adding a second Catalyst 5000 to the network.

Figure 46 Layer 3 Fault-Tolerant Fast Ethernet Backbone Design

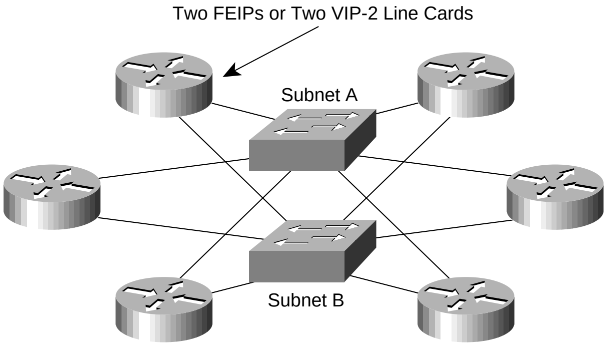

This design closely resembles the fault-tolerant design proposed for ATM; the principal difference is cost. Adding Fast Ethernet interfaces and a second Catalyst 5000 is a more affordable undertaking than the design proposed with comparable ATM equipment. In fact Fast Ethernet's affordability permits backbone designs to grow beyond a simple two-Catalyst 5000 design. Consider the backbone design in Figure 47. This backbone design comprises four parallel Fast Ethernet segments.

Figure 47 Scaling Fast Ethernet in the Backbone

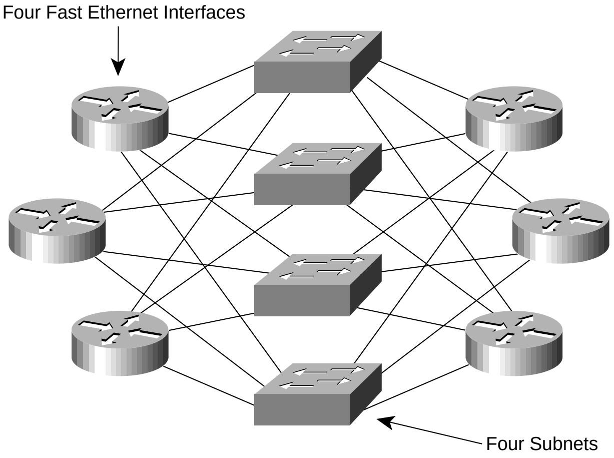

As illustrated, the Layer 2/Layer 3 design shown in Figure 47 offers four parallel Fast Ethernet paths for connected devices. Based on Layer 3 routing technologies, traffic can be dynamically load-balanced across the parallel paths. From a performance perspective, this setup delivers aggregate backbone throughput capabilities of 800 Mbps (four parallel full-duplex connections). This design also delivers exceptional fault tolerance. Besides having extensive hardware redundancy, the Layer 3 routing technologies provide fast convergence and rerouting mechanisms in the event of route failure.

Fast Ethernet Backbone Solutions—Summary

Weighed against the technical objectives of backbone reengineering, Fast Ethernet is in many ways the superior technology. Properly architected Fast Ethernet can deliver exceptional performance, fault tolerance, and scalability. And all these features can be achieved at an affordable price point and with minimal retooling and training. Environments already using Ethernet will also opt for Fast Ethernet in the backbone to instill campus-wide uniformity from a technology perspective. The wiring closets are Ethernet-based, the uplinks to data centers are Ethernet-based, and the backbone is Ethernet-based. The speed of the Ethernet feeds may be different in different parts of the network, but there still is only one underlying technology—Ethernet (see Figure 48).

Figure 48 Ethernet—The Unifying Technology

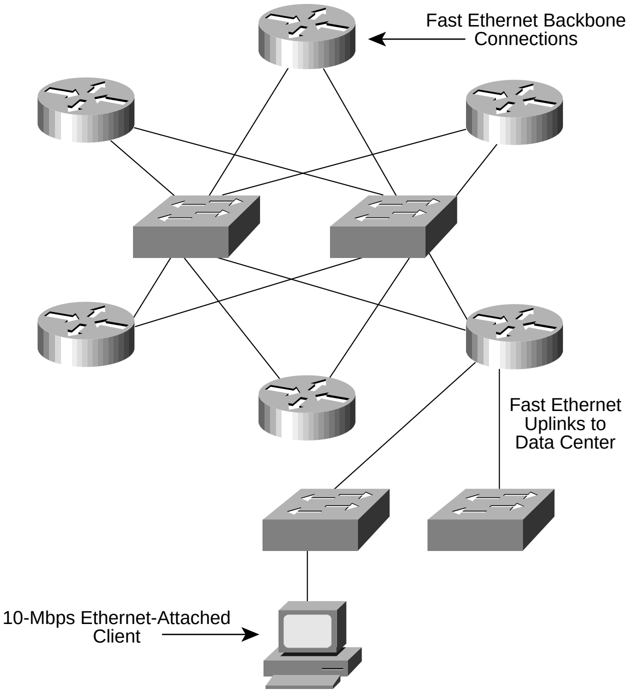
Fast Ethernet Backbone Connections
Fast Ethernet Uplinks to Data Center
10-Mbps Ethernet-Attached Client

Another important feature is the inherent migration path that Fast Ethernet has to higher-speed networking. In particular, the Fast Ethernet backbone design provides a logical migration path to Gigabit Ethernet technologies when they become available. Coupled with intelligent Layer 3 technologies, this will yield a high-performance multigigabit backbone (see Figure 49).

Figure 49 Futures: Gigabit Ethernet Migration

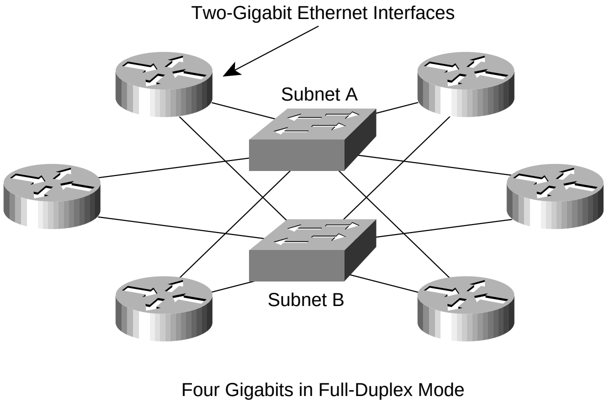
Two-Gigabit Ethernet Interfaces
Subnet A
Subnet B
Four Gigabits in Full-Duplex Mode

Conclusion

In the final analysis, all three technologies—FDDI, ATM, and Fast Ethernet—can deliver on the principal technical objective of campus backbone reengineering. In the case of FDDI, Layer 3 ring segmentation or redundancy will undoubtedly help to alleviate backbone congestion, and if done properly, will also ensure backbone fault tolerance. FDDI switching can also achieve similar ends. The downside of these strategies, though, are twofold—high cost and migration path to higher-speed networking. All these solutions are expensive and are predicated on further investment in FDDI technologies. In light of emerging high-speed standards with ATM and Fast Ethernet that push the gigabit mark, such an investment may not be wise.

As for ATM, the proposed designs will certainly provide a robust and fault-tolerant campus backbone.

But as stated earlier, migrating to ATM is a nontrivial process that first and foremost will be expensive. Add to that issues such as frame-to-cell conversion, which may impede forwarding performance, and the overhead associated with retooling and training and justifying ATM in the campus becomes a difficult prospect.

This leads us to what should be considered the most cost-effective, scalable solution of the lot—Fast Ethernet. The combination of Fast Ethernet-equipped Cisco routers and Catalyst 5000 backbone switches will deliver both exceptional bandwidth management and fault-tolerant capabilities, all at a price point that neither FDDI nor ATM can compete with.

Appendix A: Implementation Issues

The following section addresses some of the more practical issues of migrating to either ATM or Fast Ethernet. Clearly migrating from the dual-homed FDDI design to either ATM or Fast Ethernet requires careful thought and planning. To ease the process, we recommend a transition period during which the new technology is phased into production. Figure 50 details the first step in a phased strategy for Fast Ethernet backbone migration.

Figure 50 Phased Migration to Fast Ethernet—Step 1

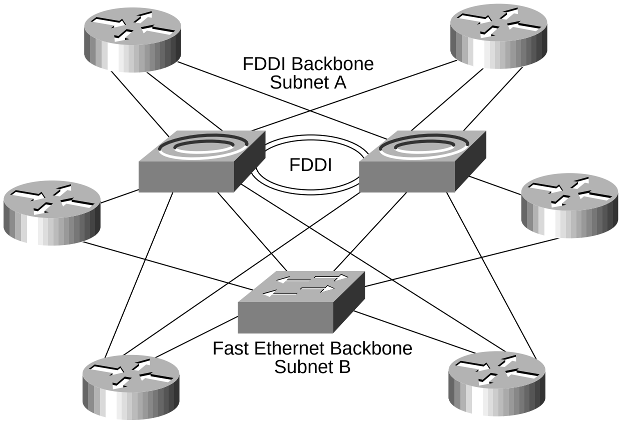

As Figure 50 illustrates, the first step in the transition/migration involves deploying a Fast Ethernet subnet in the backbone. Naturally, this process assumes that all ingress routers have an available slot for the Fast Ethernet connection. Additionally, this process assumes that extra cabling is available for the Fast Ethernet segment.

At this point in the phased migration, two paths will exist across the backbone—one over FDDI and one over Fast Ethernet. From a routing perspective, traffic will be dynamically load-balanced across both paths. During this period, use the **show** and **debug** commands to monitor traffic patterns across the backbone and the two paths. The commands in Table 2 for the Cisco Internetwork Operating System (Cisco IOS™) will be particularly helpful.

Table 2 Cisco IOS "show" and "debug" Commands

Command	Result
show interface	Display interface-specific statistics
show ip route	Display route table information
show ip eigrp interfaces	Display interface-specific Enhanced IGRP information
show ip eigrp neighbors	Display Enhanced IGRP neighbors
show ip ospf interfaces	Display interface-specific OSPF information
show ip ospf neighbors	Display OSPF neighbors
debug fddi events	FDDI-specific debugging information
debug fastethernet events	Fast Ethernet-specific debugging information
debug ip packet	Per-packet IP debugging
debug ip eigrp	IP-Enhanced IGRP packet information
debug ip ospf packet	OSPF packet information

Assuming that backbone operation stabilizes with the two paths, routing metrics can be manipulated to establish the Fast Ethernet path as primary and the FDDI as secondary. With Enhanced IGRP this can be accomplished using the **bandwidth** command in interface configuration. And with OSPF, this can be accomplished using the **ip ospf cost** command, also found in interface configuration. The goal in either case is to weight the Fast Ethernet path as a more favorable route than the FDDI path.

The next step in the phased migration is to add a second Fast Ethernet path in the backbone. This adds fault tolerance to the Fast Ethernet part of the backbone.

Figure 51 Phased Migration to Fast Ethernet—Step 2

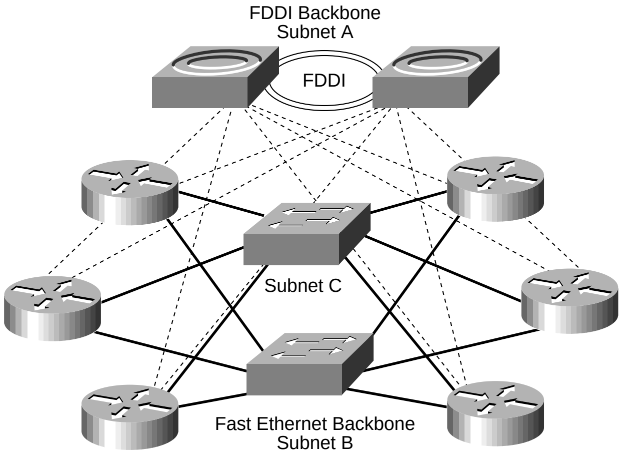

As illustrated in Figure 51, at this point in the migration process there are two Fast Ethernet paths in addition to the original FDDI backbone. To ensure successful deployment, use the **show** and **debug** commands again to monitor backbone operation. Provided that everything is stable, there are two options at this point: 1) further customize routing operations to employ the

Fast Ethernet paths, or 2) shut down the FDDI interfaces. The former option enables the FDDI path to be used as a last resort backup path, while the latter option eliminates FDDI all together from a backbone connectivity standpoint, leaving the design shown in Figure 52.

Figure 52 Phased Migration to Fast Ethernet—Step 3—FDDI Shutdown

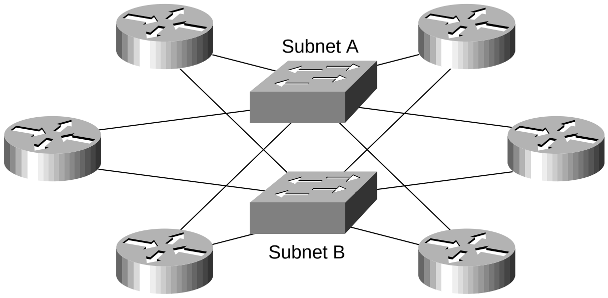

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
World Wide Web URL:
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic-Batiment L1/L2
16, Avenue du Quebec
BP 706-Villebon
91961 Courtaboeuf Cedex
France
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
Tel: 408 526-7660
Fax: 408 526-4646

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
Tel: 81 3 5219 6000
Fax: 81 3 5219 6010

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at http://www.cisco.com.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark Finland • France • Germany • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Malaysia • Mexico The Netherlands • New Zealand • Norway • Philippines • Poland • Portugal • Russia • Singapore • South Africa • Spain • Sweden Switzerland • Taiwan, ROC • Thailand • United Arab Emirates • United Kingdom • Venezuela

1/97 LW