3M

Volition™ Network Solutions

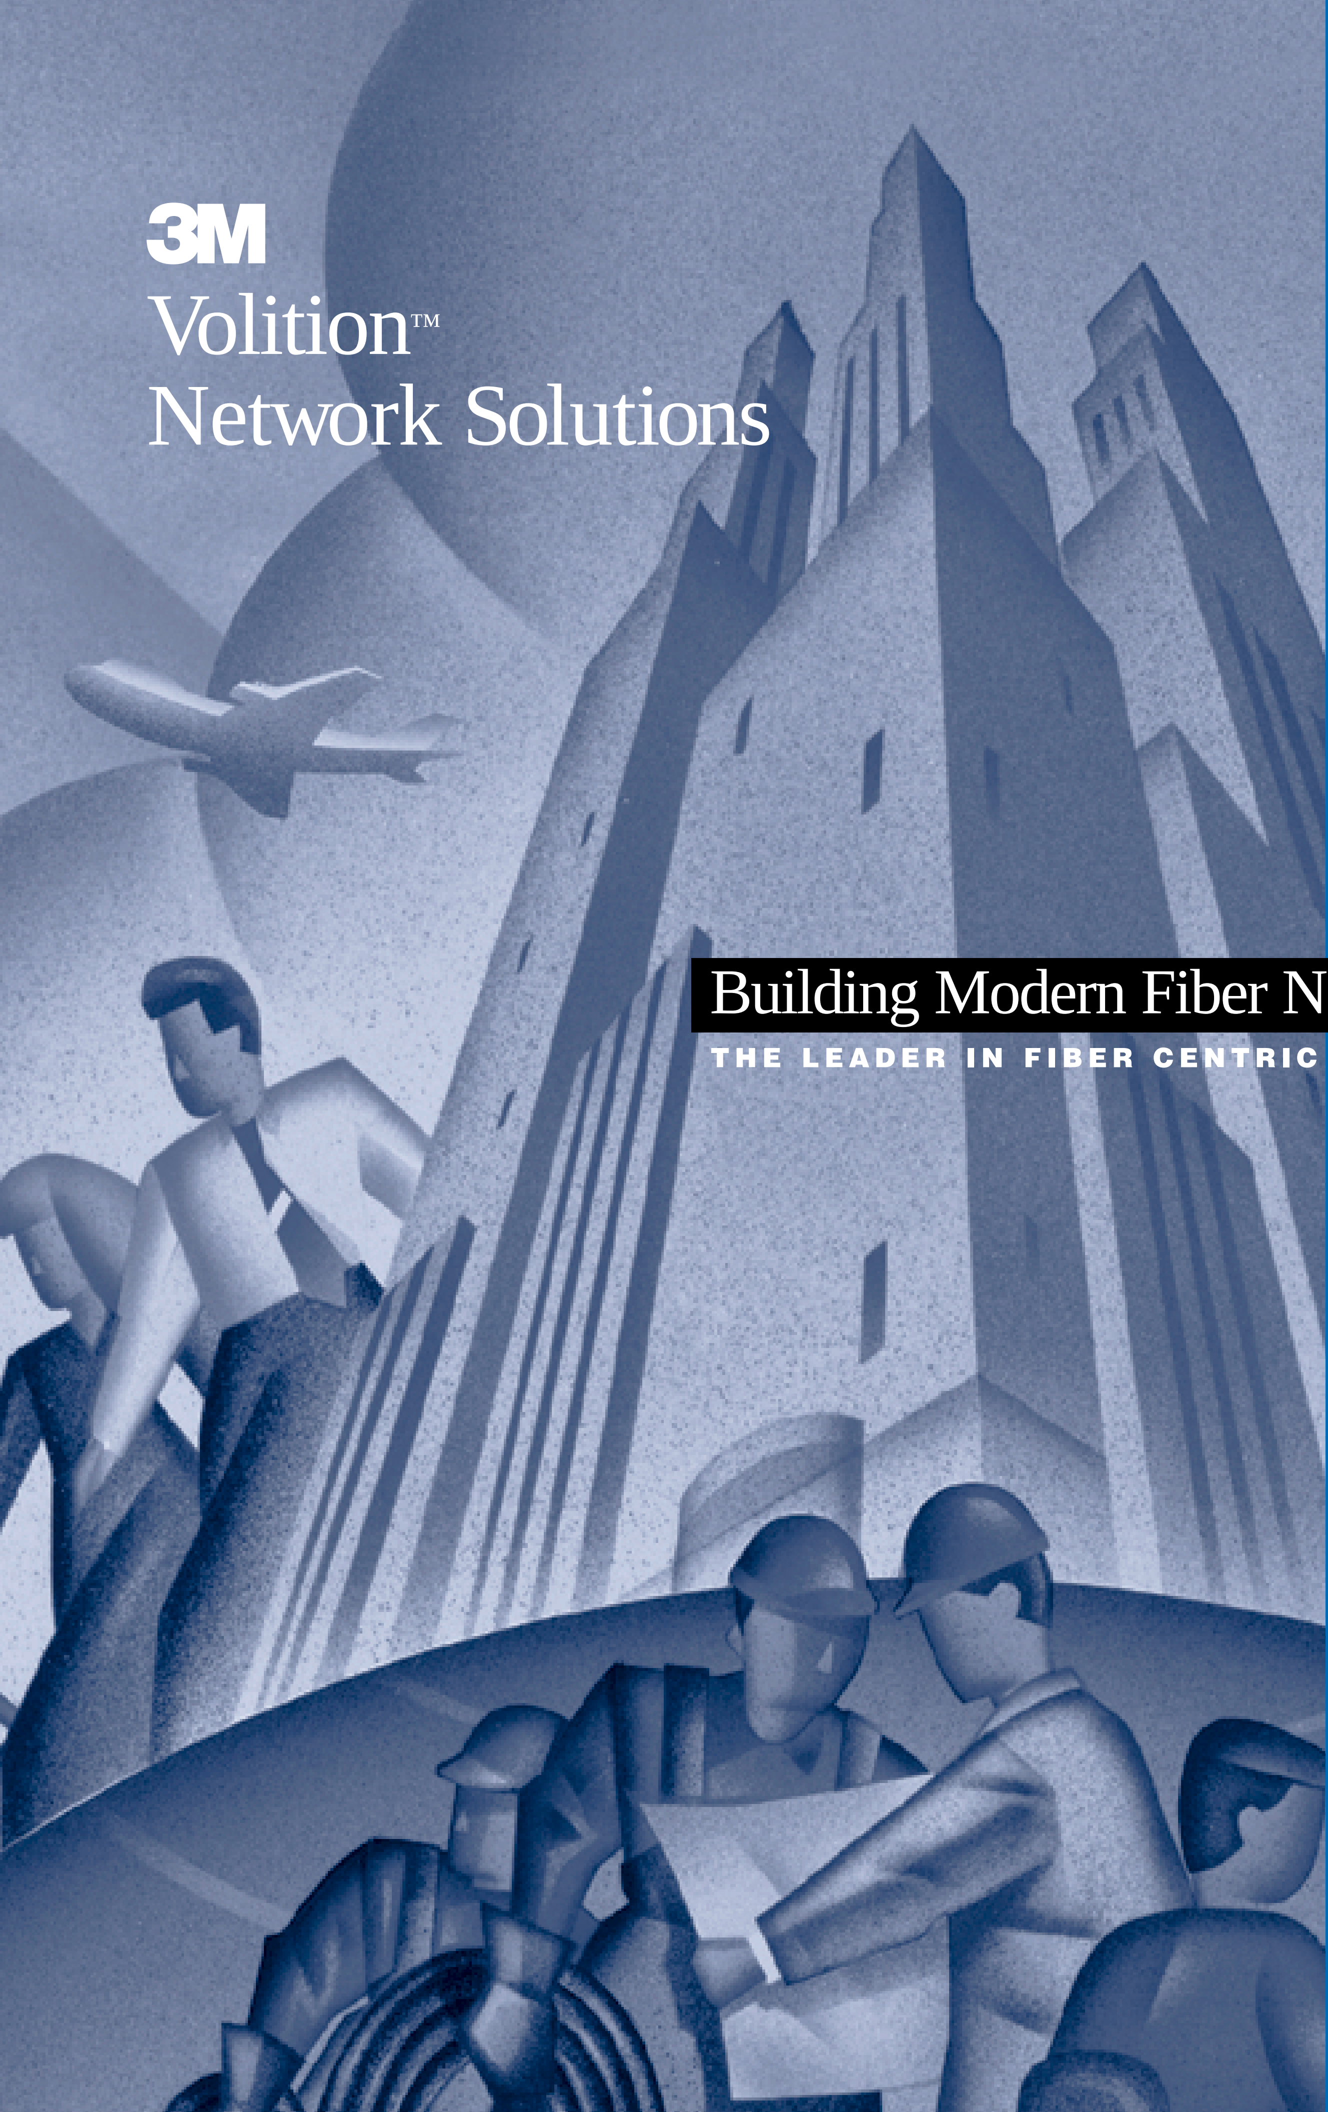

Building Modern Fiber Networks

THE LEADER IN FIBER CENTRIC NETWORK

A Guide To Building
Modern Fiber Networks

Table of Contents

Section 1: Leveraging The Fiber Switch Farm To Reduce Backbone Utilization …………..1

Section 2: Calculating Deliverable Bandwidth: The Cost Per Megabit.………………………9

Section 3: Sample Network Designs ……….……………………………………………18

Foreword

Many would ponder the need for a guide written to specifically address the design of 'Fiber to the Desk" networks. Indeed there is a vast pool of networking know-how both in written form and in practical experience that includes the use of fiber optics. However, the way that fiber is used in a traditional network vs. a "Fiber to the Desk" approach is not entirely the same. In fact, most of our accumulated experience with fiber on the premises is in the backbone and that experience doesn't always translate into '"Fiber to the Desk." The points below illustrate this.

Backbone Fiber

- Accounts for 2% of total port count
- Cost of connector is insignificant
- Install time per connector is insignificant
- Install success rate of 80% is acceptable

Fiber to the Desk

- Accounts for 98% of total port count
- Cost of connector is very significant
- Install time per connector is critical
- Install rate of 80% is totally unacceptable

It is imperative that, as we implement 'Fiber to the Desk' networks, we take advantage of the capability of fiber rather than simply replacing components (switches, NICs and cabling) in a copper network with fiber counterparts. Taking a 'systems' approach allows us to model all of the benefits of the network as a whole rather than by component. Furthermore, it allows us to truly measure the cost / benefit ratio of implementing the system.

This document has been divided into 3 sections:

Section 1 Centralized Fiber Networks

Centralizing takes advantage of the long-distance capability of fiber. This is the most efficient and cost effective way to implement 'Fiber to the Desk.' The 'Switch Farm' is a key element of this design. Centralized networks have more unblocked bandwidth than their distributed counterparts and are better suited to combined voice, video and data traffic that require implementing quality of service. Eliminating intermediate closets simplifies network layout and reduces *overall system cost* even though some of the fiber *switching components* may have higher costs per port than copper alternatives.

Section 2 Calculating Deliverable Bandwidth

Networks should be measured by the amount of bandwidth available to an end user in a worst-case scenario (i.e. when the network blocks). Users attached to a 100Mbps network don't necessarily have that much bandwidth available to them. Once the amount of actual bandwidth available is determined, that amount can be divided into the total cost of the network to come up with a 'Cost per Megabit' measure. Fiber Networks can have substantially lower costs per megabit than copper distributed networks even though they may have higher component costs. This method allows networking professionals to do a true comparison of fiber networks versus copper networks.

Section 3 Sample Network Designs

Although fiber can be installed in any configuration, the design samples presented, focus on centralizing as well as migrating from copper distributed networks to centralized fiber.

Although 3M Volition™ isn't unique in promoting the concepts contained in this document, the refining of these principles came from our own experience in selling 'Fiber to the Desk' components over several years and observing their successful implementation in networks ranging from 10 users to 10,000 users.

Many thanks to the lab engineers, support technicians, product managers and sales reps. for their input and especially the visionaries in those enterprises who understood and implemented these revolutionary fiber architecture designs.

Daniel Harman
December 2001

Leveraging the Fiber Switch Farm to Reduce Backbone Utilization

Centralized Fiber Networks
- Leveraging the Fiber Switch Farm to Reduce Backbone Utilization -

Thanks in part to Sprint and their pin-drop commercials, selling the goodness of fiber optics is like preaching to the choir. Most people have the perception that fiber optics are as fast as the speed of light but that like all exotic technologies are also of limited use for mere mortals, expensive, fragile and reserved for specialized applications.

Fiber Optics Perceptions

• High Speed "Speed of Light"
• Used for Network Backbones
• Specialized Use – (e.g. distance, security, etc.)
• Fragile
• Expensive

Fiber has an exotic range of capabilities

Fiber performs over a very wide range of bit rates: 10 Mbps, 100 Mbps, 1 Gbps and in 2002 will also operate at 10 Gbps per wavelength. In fact all of these bit rates will operate on a single cable type (i.e. 50.0µm cable) as well as various multi-mode and single-mode cables. Of course the fiber must conform to standards-based specifications in order to achieve maximum performance. Fortunately, fiber manufacturers are continuously developing product that can meet these critieria.

Fiber uses a simple duplex transceiver to transmit and receive. To achieve a faster bit rate we can run the transceiver faster. In comparison, transmission schemes such as gigabit over copper require compression codecs, multiplexed transmission over multiple twisted pairs as well as higher frequency cabling like Cat. 5E or Cat. 6.

Of course, fiber has a well-known immunity to electromagnetic and radio interference meaning that it can be run alongside power cables. It is also preferred in secure data environments because it is difficult to tap.

As more and more network cabling is installed, physical size becomes even more important. Six workstations can be connected with a fiber cable that has the same diameter as a single Cat. 5e cable Buildings that have restricted pathways for network cabling will find installation and expansion using fiber cabling easier and less expensive. In addition, the advent of the SFF (small form-factor) fiber connector and transceivers means that fiber switches now have the same port density as copper switches (with RJ-45 connectors).

Long distance capability – the most important attribute

If you had to pick the one characteristic of fiber that can be leveraged to the greatest extent, it would have to be the long-distance capability. Properly deployed, this capability has the potential to make a 100 Mbps centralized fiber network perform better than a 100 Mbps distributed network, or a fiber centralized gigabit network perform better than a distributed gigabit network.

Fiber Capabilities Today

• High Bit Rate - 10, 100, and 1000 Mbps per wavelength
• Simple Duplex Operation – No compression and multiplexing needed for high bit rate
• Immunity to Interference - EMI and RFI
• Security - Very Difficult to Tap
• Small Physical Size – Six-workstation fiber cable is the same diameter as a Cat. 5e cable: Small Form Factor connector enables same density as copper
• Long Distance Transmissions - The most valuable feature today for delivering non-blocked bandwidth over distance

To take full advantage of the long-distance capability of fiber, network managers should implement the Centralized Architecture sometimes referred to as "Collapsed Backbone."

This architecture uses substantially fewer equipment closets, is easier to manage and allows the formation of "Switch Farms" similar to the frequently implemented Server Farm.

Giving workstations direct access to the Central Equipment room (location of the Switch Farm) makes it easier to match users to services across switch backplanes thereby distributing greater bandwidth to the end user.

By concentrating all of the unused ports in a network into a single (Switch Farm) location, port utilization can be increased. Typically a distributed copper network has wasted switch ports in every intermediate closet with an overall unused port count as high as 30%. A centralized fiber network, on the other hand, can achieve an unused port count as low as 5-10%.

Long Distance Capability Enable the Following:
- Implementation of Centralized Architecture
  - –Fewer Equipment Closets
  - –Simplified Maintenance
- Formation of "Switch Farms"
- Matching User / Services / Backplane
- Greater Bandwidth and Port Efficiencies

How did LANs evolve?

In the early to mid 80's, Novell and 3COM helped to develop a mass market for the Local Area Network (LAN). Novell built a low-cost network operating system and freely licensed the NE1000 chipset design to manufacturers like 3COM who in turn manufactured low cost PC-based network adapters. Back then, Novell and others used the 80/20 Rule. This rule of thumb states that 80% of a workgroup's traffic remains within the workgroup and that 20% of the traffic moves between workgroups. Every organization has distinct workgroups such as accounting, sales, engineering, etc.

In the past, networks were subdivided into workgroups, each with its own server. Keep in mind that, in a copper distributed network, no end-user can be further than 100m away from the nearest switch. Note that 80% of network traffic remained within the workgroup while 20% moved between them over the backbone.

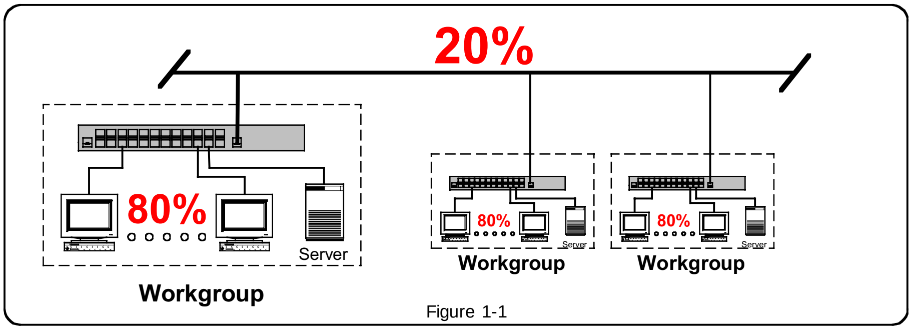

Figure 1-1

In the late 80's, implementation of Enterprise (Global) services like e-mail increased the movement of traffic over the backbone. This occurred as users in other workgroups gained access to these services.

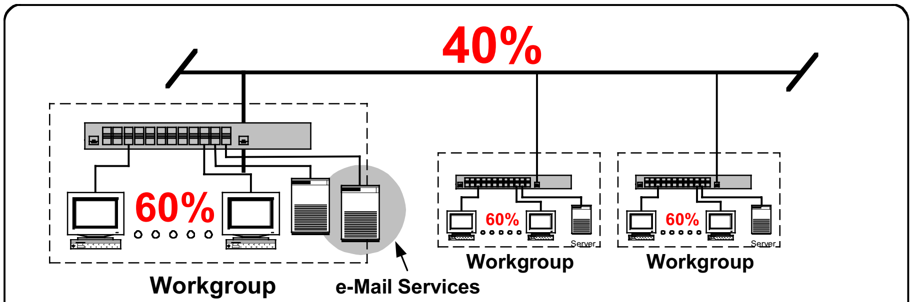

Figure 1-2

Eventually these Enterprise services were moved to a centralized part of the network leading to the development of the "Core." As can be seen, more traffic was now moving over the backbone.

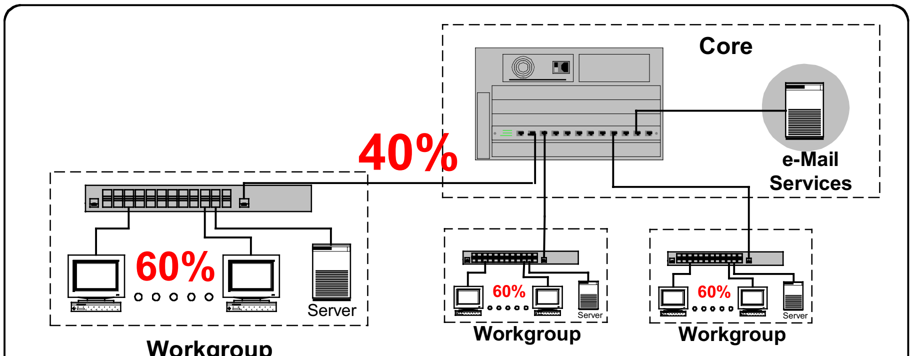

Figure 1-3

Server Farms Drive Backbone Traffic

The Server Farm has flipped the 80/20 rule on its head and created the 20/80 rule. Now 80% or more of the traffic moves over the backbone while only 20% or less remains within a workgroup. Server Farms have driven the demand for more potent core switches and technologies to manage distributed network environments such as Fast Ethernet, Gigabit Ethernet, VLANs, Trunking, etc.

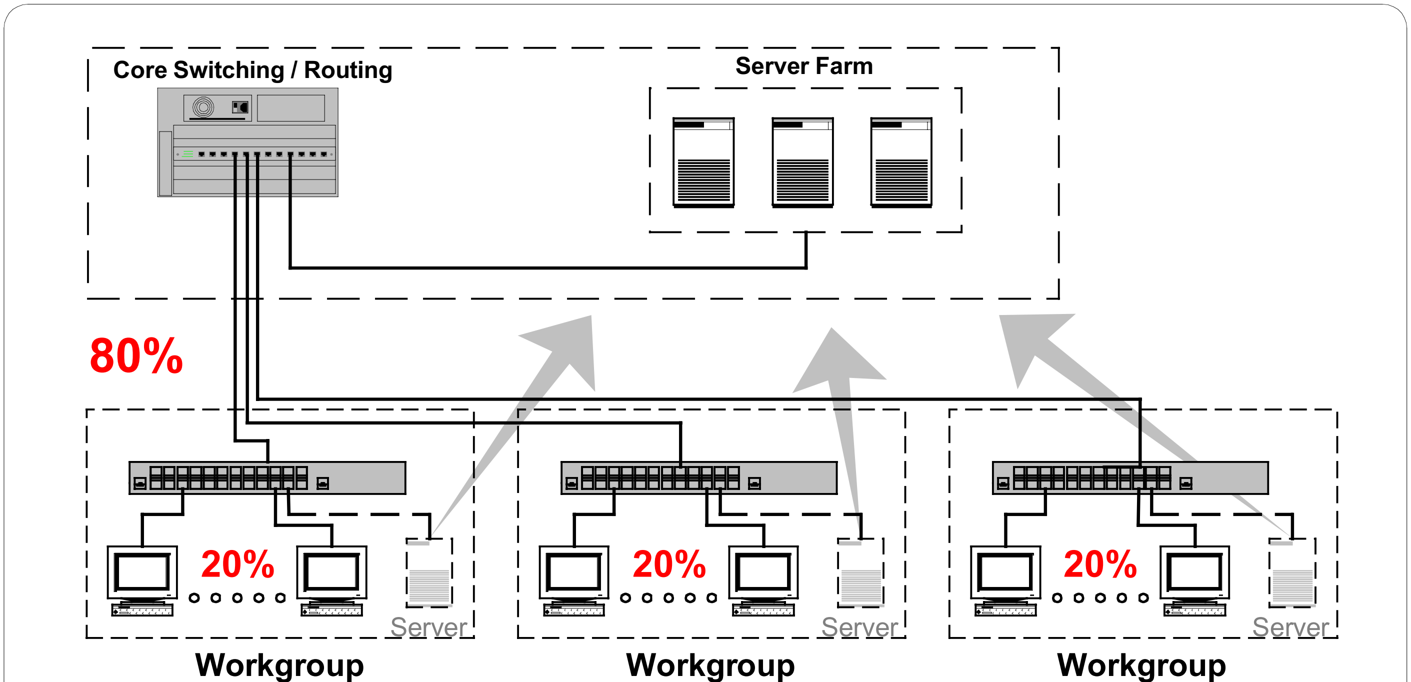

Figure 1-4

Today there is an increasing demand placed on the backbone both because of the proliferation of networked PCs and because of the implementation of new networked appliances. For example:

IP Telephones. Although IP telephony consumes very little bandwidth, it requires Quality of Service, which means high-quality bandwidth with no blocking.

Instrumentation. These devices weren't networked in the past but today are. Heating and environmental control systems, machine tools and even refrigerators and microwaves are all becoming networked Ethernet devices.

Audio-visual systems. These used to run on separate networks in schools and are now running on Ethernet-based IP networks.

Voice and Video Conferencing. Corporations are enabling telecommuting with high bandwidth access.

Security Cameras. These devices are now running over IP networks.

This rapid growth of bandwidth demand means that where there used to be one networked device per 10 people, in the near future there could be 10 networked devices for every person!

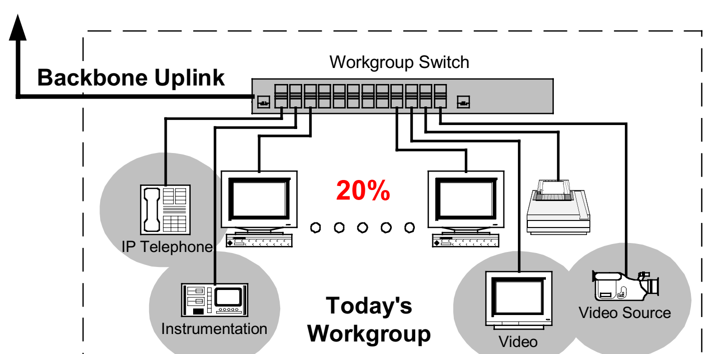

Figure 1-5

Existing Product Lines Reflect Legacy Architecture

It's no mystery that traditional networking equipment manufacturers have built entire product lines based on the distributed network. These product lines usually consist of high-end core switches and low-end workgroup or edge devices.

Core Switches (Figure 1-6) usually have non-blocking architecture, which translates to lots of bandwidth. They are usually chassis based in an attempt to switch all core traffic on a single switching fabric. They have the best feature set, are managed and have the best quality of service. Among switches they have the highest port cost.

Workgroup / Edge Switches have the lowest cost per port of any network switch and thus are the most widely purchased of all switches. Although many manufacturers advertise non-blocking architecture, this is only true on the models with the lowest port counts. The higher density models that include gigabit uplinks and stacking ports can be oversubscribed by a factor of 3 or more (meaning that the switch has 1/3 of the bandwidth its ports could handle; Figure 1-7).

In all fairness, switch manufacturers take several factors into consideration when advertising non-blocking capacity. One very significant factor is that PCs with PCI-based network adapter cards don't actually get full throughput on a 10, 100, or 1000Base link.

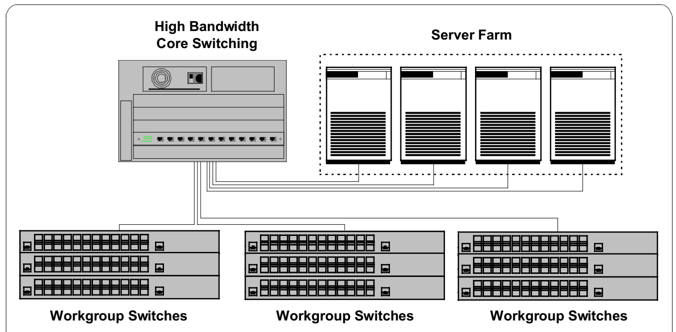

Figure 1-6

In fact a 100Mbps adapter card may only achieve a network throughput of 40 to 65 Mbps while a gigabit adapter may only achieve 400 Mbps. This is a limitation of both the PCI bus and most network operating systems.

Although switch manufacturers plan more capable core switches, and have gradually been increasing the bandwidth of their workgroup switches, this doesn't address the true bottleneck in the network: the backbone link (Figure 1-8). Since it's critical for backbones to be standardized (IEEE sets the standards), individual manufacturers can't rapidly introduce new backbones.

In a distributed network, this is unfortunate since the end users aren't normally connected directly into the powerful core switches, rather they are plugged into the mediocre edge devices and forced to communicate with the core environment via fixed-bandwidth backbones.

Two Types Of Blocking

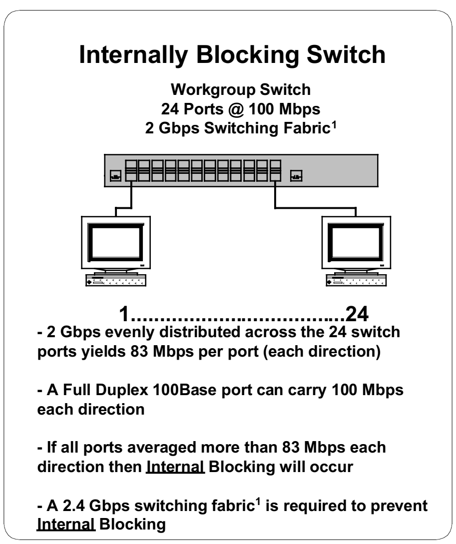

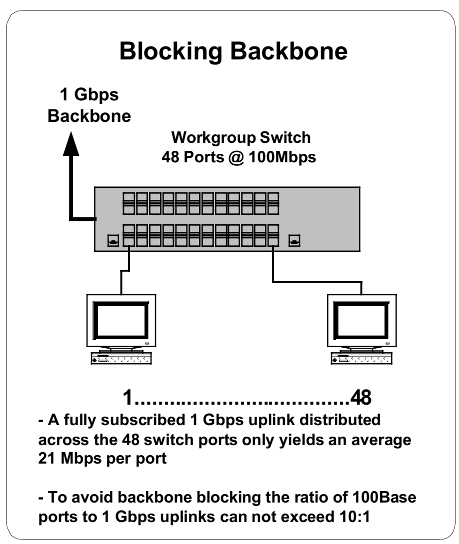

[1] The convention in this document to describe switching fabric and link bandwidth refers to the bit forwarding rate (in one direction) with all interfaces set to full duplex mode.

Figure 1-7

Figure 1-8

Now that we are clear about the limit of a traditional network design, this brings us to the crux of our argument regarding fiber networks:

The long distance capability of fiber allows us to move the workgroup switches from the intermediate closets to the main equipment room and give users direct access to the core environment, thus giving each user, office or classroom a dedicated private backbone to the feature-rich core environment.

Fiber distance capability allows backbones to be collapsed all the way to the Data Center. This creates "Switch Farms" and allows network managers to match users, services and backplanes.

As can be seen in figure 1-9, in a centralized fiber network both core and workgroup switches are combined in the core environment allowing the server load to be distributed among the switches.

This returns us to the old and familiar 80/20 rule of thumb. But now the network backbone resides within the switches and locally between switches in the same room!

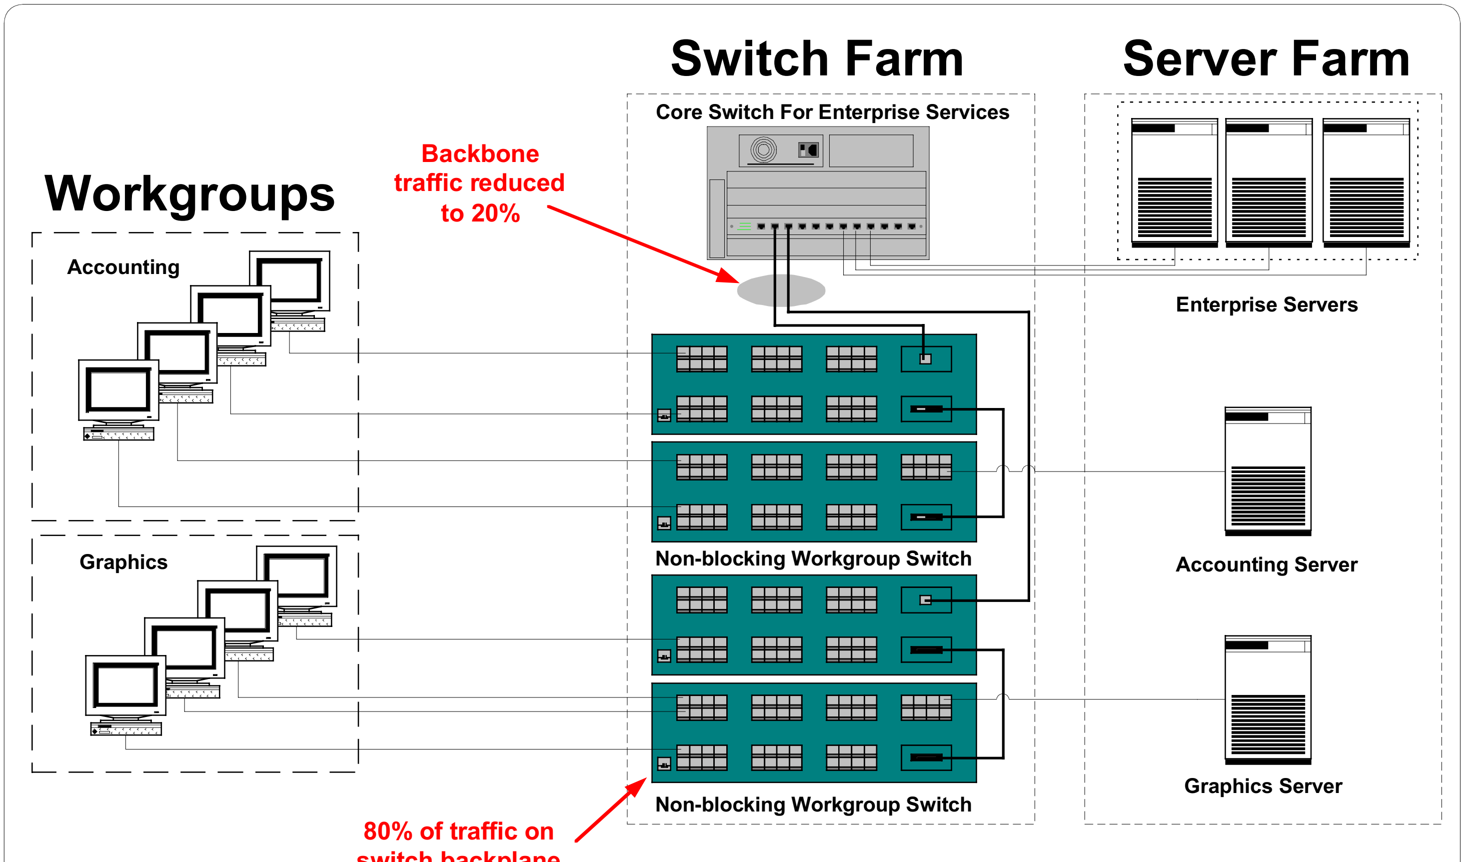

Figure 1-9

Taking a closer look at the Switch Farm we see that both the core switch and the workgroup switches are now in the same room or even in the same rack. Backbone cables connecting the two are very short. Typically, the backbone cable is a patch cord!

A new and slightly different type of switch should be utilized in this environment that is a cross between a core switch and a workgroup switch.

- Non-blocking design
- Configurable as a core switch
- Advanced feature set (i.e. QoS)
- Configurable as a workgroup switch
- Stackable design for scalability
- Fully managed
- Lower costs than a traditional core switch

This describes the Volition™ line of switches developed by 3M. The VOL-5000 is shown in Figure 1-10.

Service Architecture

To reduce the amount of traffic moving over the backbone, only enterprise services should be plugged directly into the core switch (top of Figure 1-10). Enterprise services are those that are routinely accessed by all members of an organization and represent services such as e-mail, Internet access etc.

Group-specific services however, should not be plugged into the core switch, but instead should be plugged into the same switches as the workgroup that uses those services (bottom of Figure 1-10). For example, members of the accounting department and the accounting department server should share the same switch. This, in effect, isolates group-specific traffic within a switch while preserving the backbone link for global traffic. It also does so without breaking up the server farm, provided that the server farm is within cable range of the switch farm.

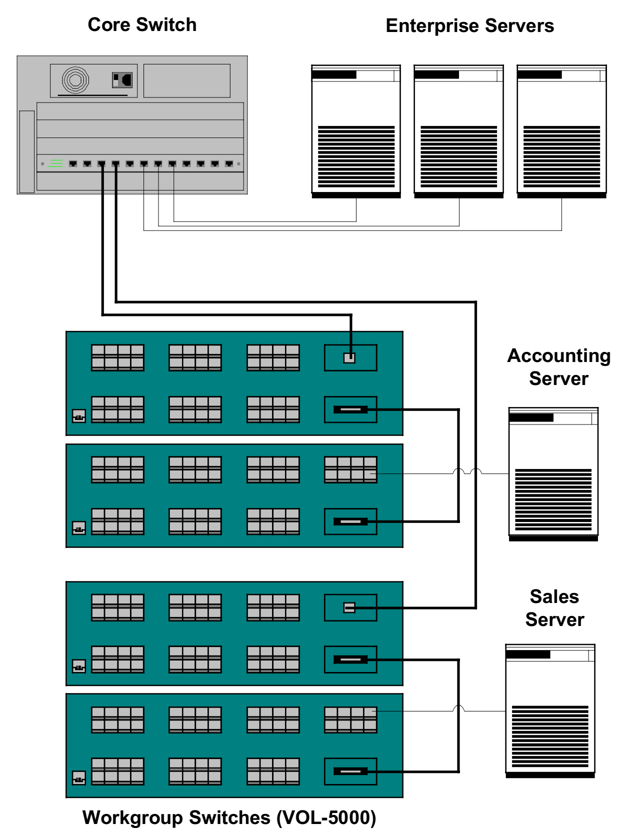

Figure 1-10

Deploying the "Centralized Backbone" architecture in a large organization (with a network radius greater than 100 meters) is only possible using fiber optics.

Benefits of Centralized Networks

Fewer Equipment Rooms. Where new construction is concerned this design can lead to dollars saved, sometimes on a dramatic scale. The first and largest Volition™ Network Solution site, The George Washington University, reduced closets needed from 180 to 11. This offered dramatic savings in terms of reduced real estate and HVAC requirements of the typical 8' x 10' communications closet.

Consolidated Port Counts. Switch ports come in fixed increments such as 8, 12, 16, 24, and 48. When switches are distributed throughout an organization there are often many unused ports. At $50 to $75 or more per port, this may represent a significant waste of money.

Reduced Backbone Requirements. Matching users to services and switches keeps unnecessary traffic off of the backbone and makes it available for the ever increasing demands of internet access and other enterprise services.

Reduced Network Blocking. Using the switch as the backbone reduces network blocking and makes it more suitable for applications requiring high QoS like VoIP and streaming video.

Simplified Network Design. A recent magazine article stated that although 90% of switch buyers wanted to buy switches with VLAN capability only 5% of those people actually implemented VLAN. Network deployment is increasing faster that we can train people to build them. People really don't want more complicated network management or architectures. Collapsed backbone networks can greatly simplify network design and management while giving individual users access to greater bandwidth and quality of service.

Is There a Disadvantage to Centralized Networks?

More cable is needed for longer individual runs. In new construction the extra cost of the cable can be offset by reduced construction costs of the Equipment room and fewer wasted ports; however, in remodeled construction it's often more difficult to account for the extra cost.

One astute customer observed, "I'd rather spend my money on a 20-year component than on a 5-year component."

In fact too many customers put most of their money into switches and ignore the cabling even though they end up re-cabling every 5 years. A properly provisioned centralized fiber network that today delivers 100 Mbps to every end user can continue to deliver increasing bandwidth for years to come.

That is what represents real value.

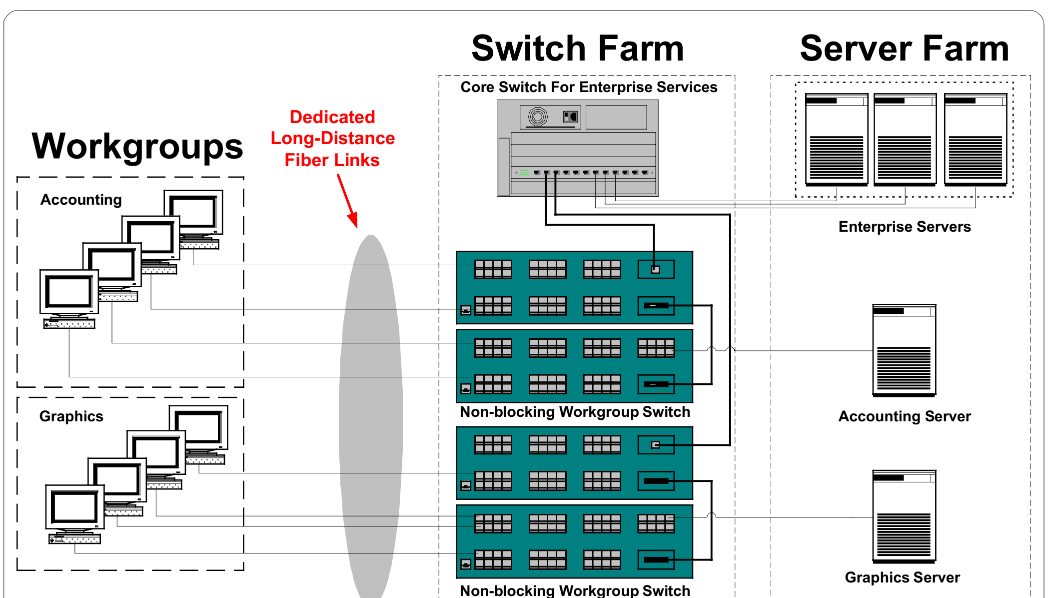

Figure 1-11

Calculating Deliverable Bandwidth

- The Cost Per Megabit -

One of the best ways of accurately comparing a Volition™ solution against any other is by calculating the present and future deliverable bandwidth to an end user, office, or classroom.

Attempting to compare centralized fiber networks to distributed networks switch by switch or NIC by NIC is useless. Since the architecture of fiber networks should be different, the true comparison is a system-wide comparison. Ultimately networks are designed to deliver a specific amount of bandwidth to users under the worst of conditions, so the best comparison is one where we calculate the amount of bandwidth delivered using each design. This provides an "apples to apples" comparison that becomes meaningful.

A traditional distributed network is multi-tiered with core switches at the center, linked by a fiber backbone to edge devices. Users are attached to the edge devices with 10/100 copper links. Users cannot be more than 100 meters from the nearest switch. In Figure 2-1, individuals in a classroom might use a dedicated 10/100 copper link or multiple students could share a single 10/100 copper link using a desktop switch or a wireless hub (common in many school applications, these desktop devices are not considered part of the infrastructure because they don't require permanent installation).

Most users assume that because they are attached to a 100Base network they actually have 100 Mbps bandwidth available.

This is true if only a couple of users are using the network, but as more of the attached users start transmitting, the available bandwidth for each user rapidly drops to less than 100 Mbps!

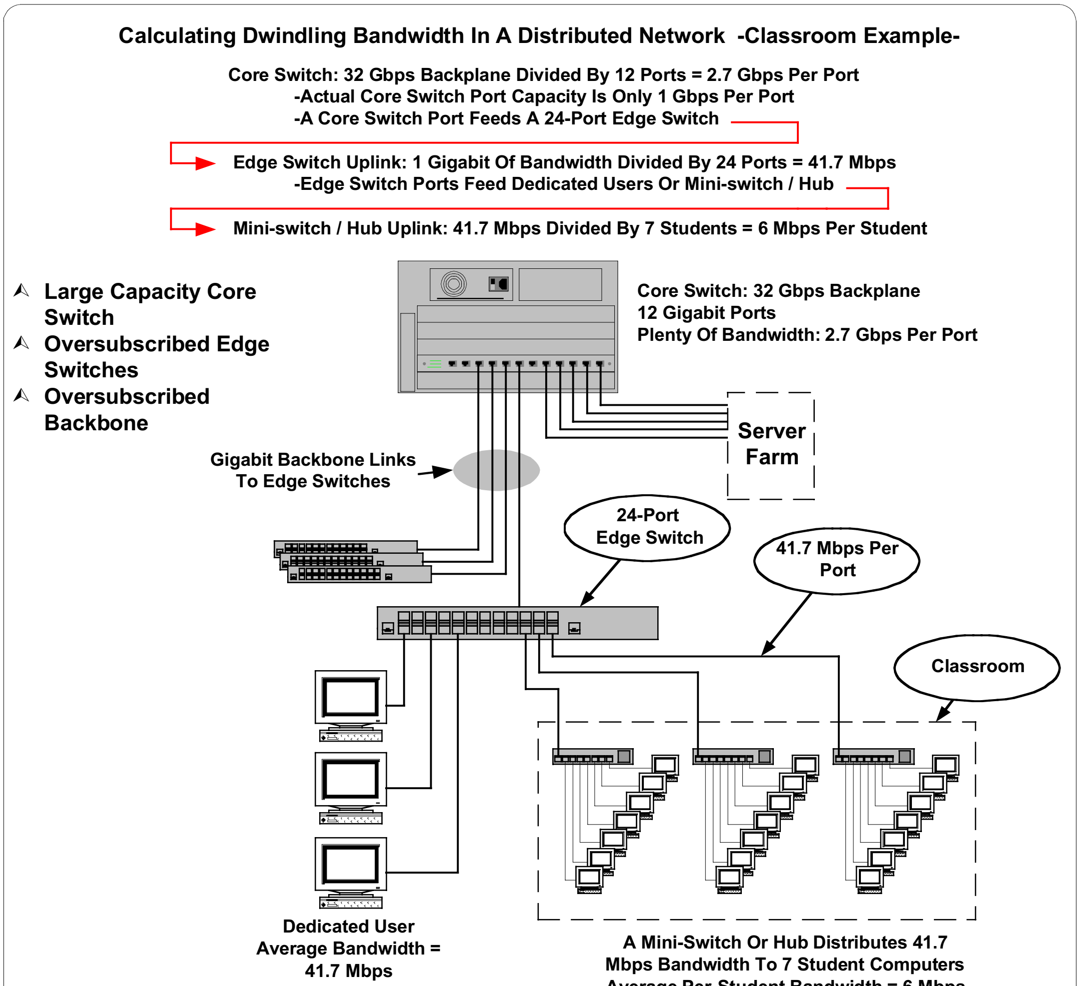

Figure 2-1

Calculating Deliverable Bandwidth
- The Cost Per Megabit -

Traditional switch manufacturers have addressed the issue of delivered bandwidth from the switch perspective. A current chassis-based workgroup switch is branded as an enhanced edge device. Although increasing the bandwidth of the edge device seems like a good thing, it doesn't address the real bottleneck, which is the backbone. While the current trend is to add more ports to the workgroup switch, this only compounds the congestion on the backbone. Unfortunately, changing the backbone is a slow process, determined by committee (IEEE), taking years to implement.

Most IT Managers try to solve the bandwidth crunch by moving to a larger core switch. Bigger switches allow backbones to migrate from 100 Mbps to 1 Gbps, but still may not reduce congestion on the backbone. Moving to a larger core switch will increase the available bandwidth for backbone links, but this can help only if the core switch was suffering from internal blocking. Otherwise this does nothing to reduce workgroup congestion since the link bandwidth does not increase.

Many companies have built their entire product line and marketing strategy around large-capacity core switches. Much is spent on expensive core switches when the real solution to the bandwidth problem lies with the cabling and backbone infrastructure.

Another trend with many workgroup switch manufacturers is to increase the bandwidth of the edge switch and the port density (from 24 ports to 48). Increasing the bandwidth of the edge switch does nothing to solve congestion on the backbone. On the contrary, adding more ports to the edge switch **increases** congestion over the backbone and cuts the available bandwidth per user even more (Figure 2-2). In a distributed network where the backbone is the limiting factor, one thing that will increase end user bandwidth is reducing the number of ports per uplink by increasing the number of backbone links (i.e. re-cabling!).

Another option is to increase backbone bandwidth (i.e. from 1 gigabit to 10 gigabit). Unfortunately, backbone standards are decided in committee, not by any single switch manufacturer, and deploying bigger backbones can take years!

Assuming non-blocking switching has been implemented, the only way to increase bandwidth in a network is to **reduce** the number of end-users per backbone link. This means putting in edge devices with **fewer** ports rather than more ports. This also means that the number of backbone links must be increased. Unfortunately, adding more backbone links means that the network must be re-cabled.

Therefore, the ultimate implementation of backbone expansion is to provide a dedicated backbone link to every end user.

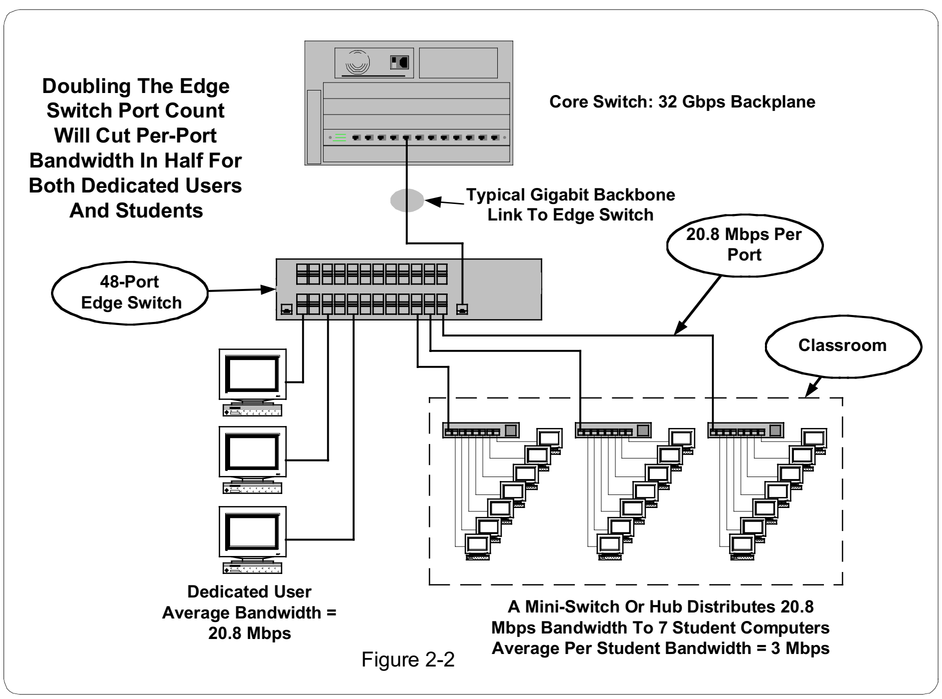

Figure 2-2

Calculating Future Bandwidth In A Distributed Network

The next part of the problem is to figure out how much bandwidth could be delivered over the proposed infrastructure with equipment upgrades (without changing the cabling) using known and planned technologies.

In comparing centralized fiber networks to distributed networks, it's also important to calculate the potential of the infrastructure. This will give us an idea of the expandability of the network.

Assume that at some future date switches are upgraded to take advantage of forthcoming IEEE standards. Fiber links are upgraded from Gigabit to 10 Gigabit[1], 10/100 copper links are upgraded to Gigabit copper links[1]. These are technologies that are foreseeable today.

Calculations using the new standards demonstrate the potential increase in bandwidth. As can be seen, dedicated end users increase from 41.7 Mbps to 417 Mbps and Student shared bandwidth increases from 6 Mbps to 60 Mbps.

[1] Upgrading to higher bandwidth requires that the original cabling (copper or fiber) meets the appropriate specifications. For example, Cat. 3 copper cabling will not carry 1000Base-T per the IEEE standard.

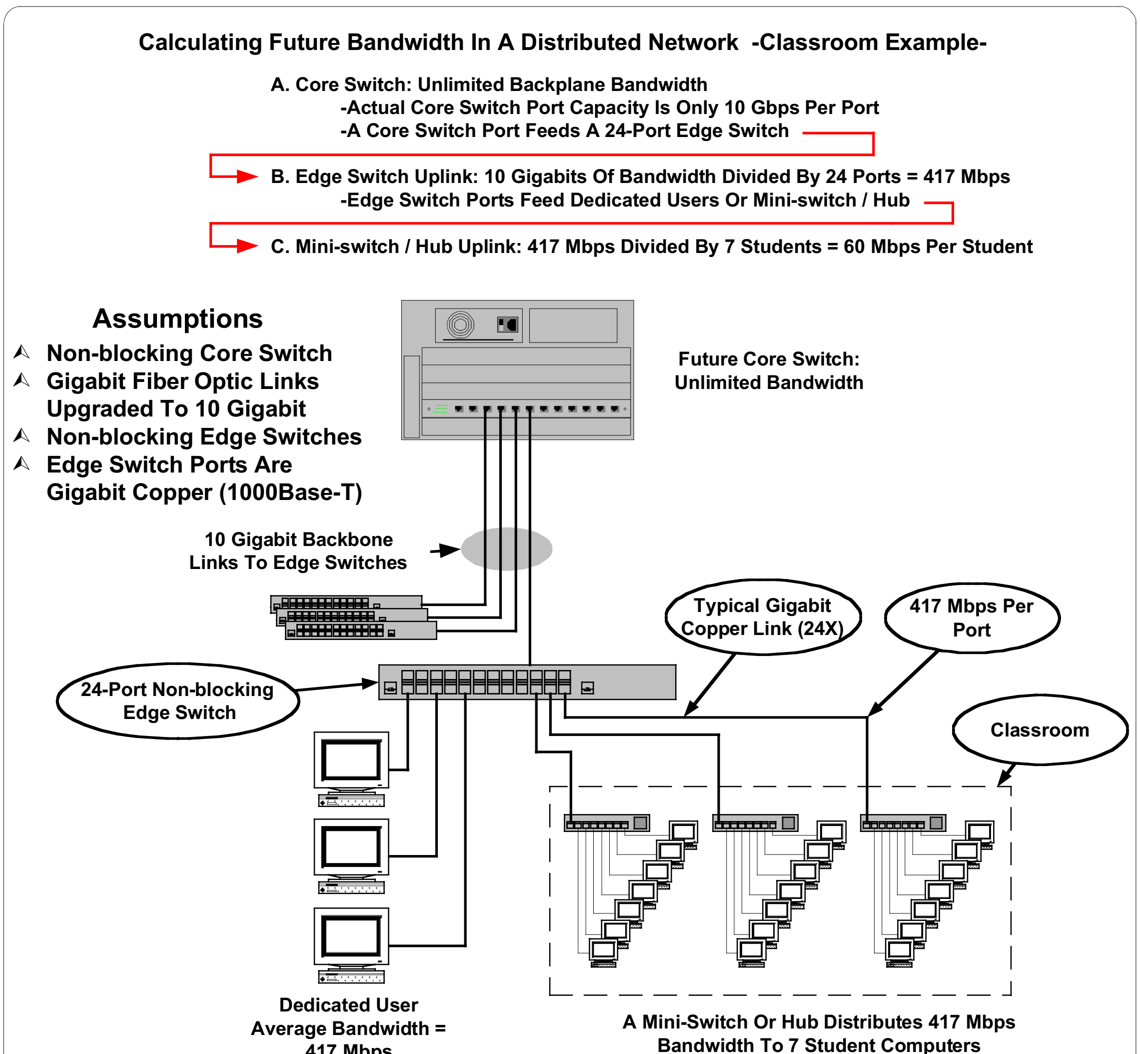

Figure 2-3

Is a 100 Mbps centralized fiber network actually better than a 100 Mbps distributed network?

A centralized network reduces the number of tiers commonly found in networks today. In smaller networks (with a radius smaller than 100 meters), the collapsed architecture can be implemented with either copper or fiber cabling. Larger networks require dedicated fiber to operate with a single tier. Campus networks that deploy fiber to the desk may eliminate 1, 2, 3, or more tiers and still deploy the switch farming technique with network radii up to 2 kilometers.

There are fewer points to analyze in a centralized fiber network. The calculations in Figure 2-4 show the bandwidth available at the core, on the backbone, for a dedicated user and for students sharing a desktop mini-switch.

As can be seen from the calculations, the limiting factor in a centralized fiber network is the core switch. This is great news for switch manufacturers, because it allows them to build switches that meet the demand for bandwidth without waiting for IEEE committees to develop new backbone standards. It also allows IT managers to increase available bandwidth by buying new hardware but without re-cabling. In addition, having all of the switches in a single location forms an easier to manage switch farm. Switch farms distribute the load of the servers across many switches instead of forcing all traffic into a single gargantuan core switch.

As illustrated in Figure 2-4, dedicated end-users attached to a 100 Mbps link <u>actually get 100 Mbps!</u> Students who are sharing a 100 Mbps link get 14.3 Mbps which is much better than the 6 Mbps delivered in the distributed network classroom example (Figure 2-1).

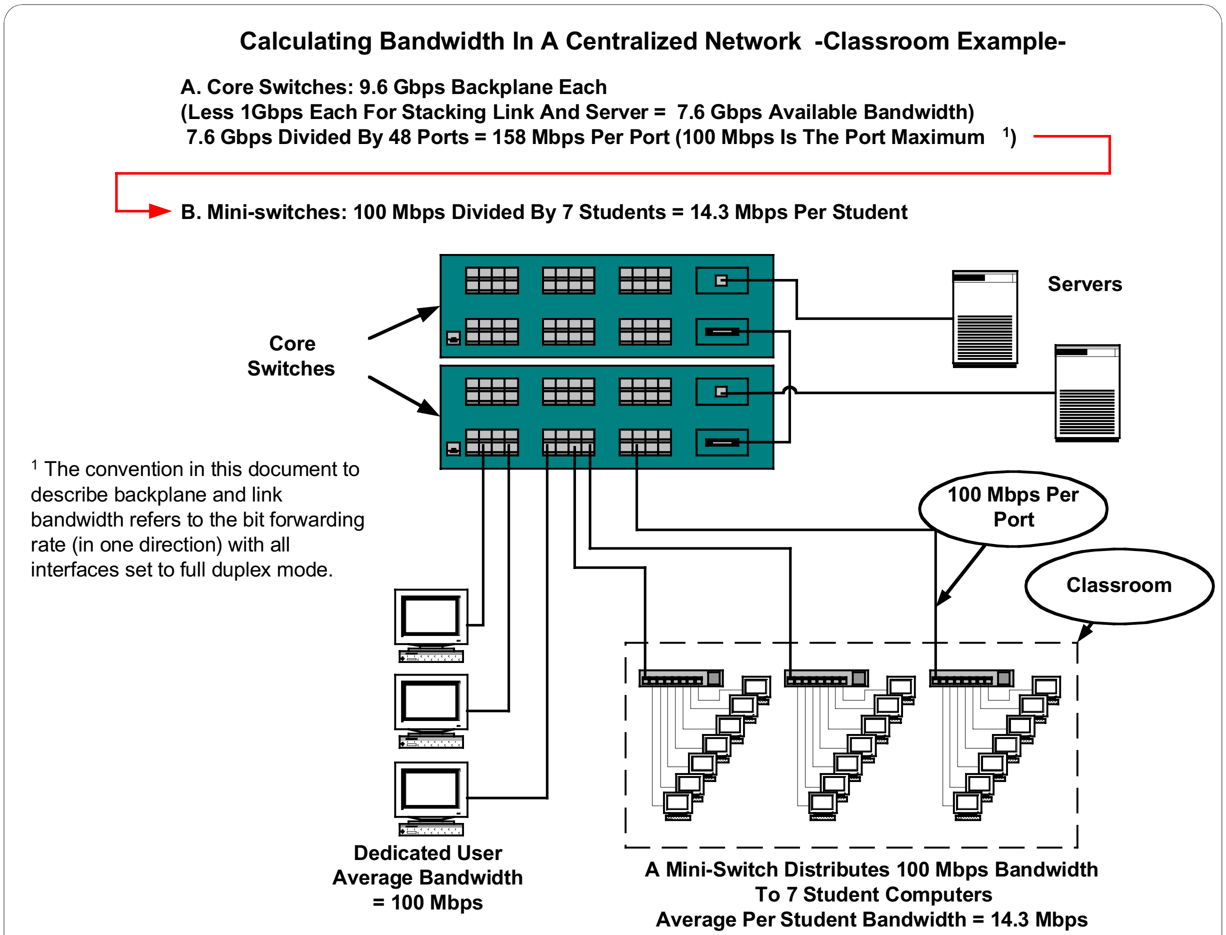

Figure 2-4

How will fiber networks fare in the future?

The next part of the problem is to figure out how much bandwidth can be delivered over the proposed centralized fiber infrastructure with equipment (not cabling) upgrades using known and planned technologies.

Assuming future switches will be non-blocking, 100 Mbps fiber links will be upgraded to 1 Gbps and then to 10 Gbps.

Upgrading to higher bandwidth requires that the original fibers meet the appropriate specifications. For example, 62.5µm fiber with a modal bandwidth of 160 MHz_km will not carry 10GBase-S more than 26 meters per the IEEE draft standard (Feb. 2002). However, 50µm fiber with a modal bandwidth of 2000 MHz_km is expected to reach 300 meters with 10GBase-S. Other 10 Gigabit technologies can reach even farther.

For example, CWDM transceivers can enable 2-kilometer distances over multi-mode cabling. There is also work in compression encoding. (CWDM = Coarse Wavelength Division Multiplexing)

As can be seen in Figure 2-5, dedicated users increase from 100 Mbps to 10 Gbps, while student shared links go from 14.3 Mbps to 1.4 Gbps. This represents a truly remarkable value and one that can keep pace with Moore's Law (i.e. processing speeds will double every 18-24 months).

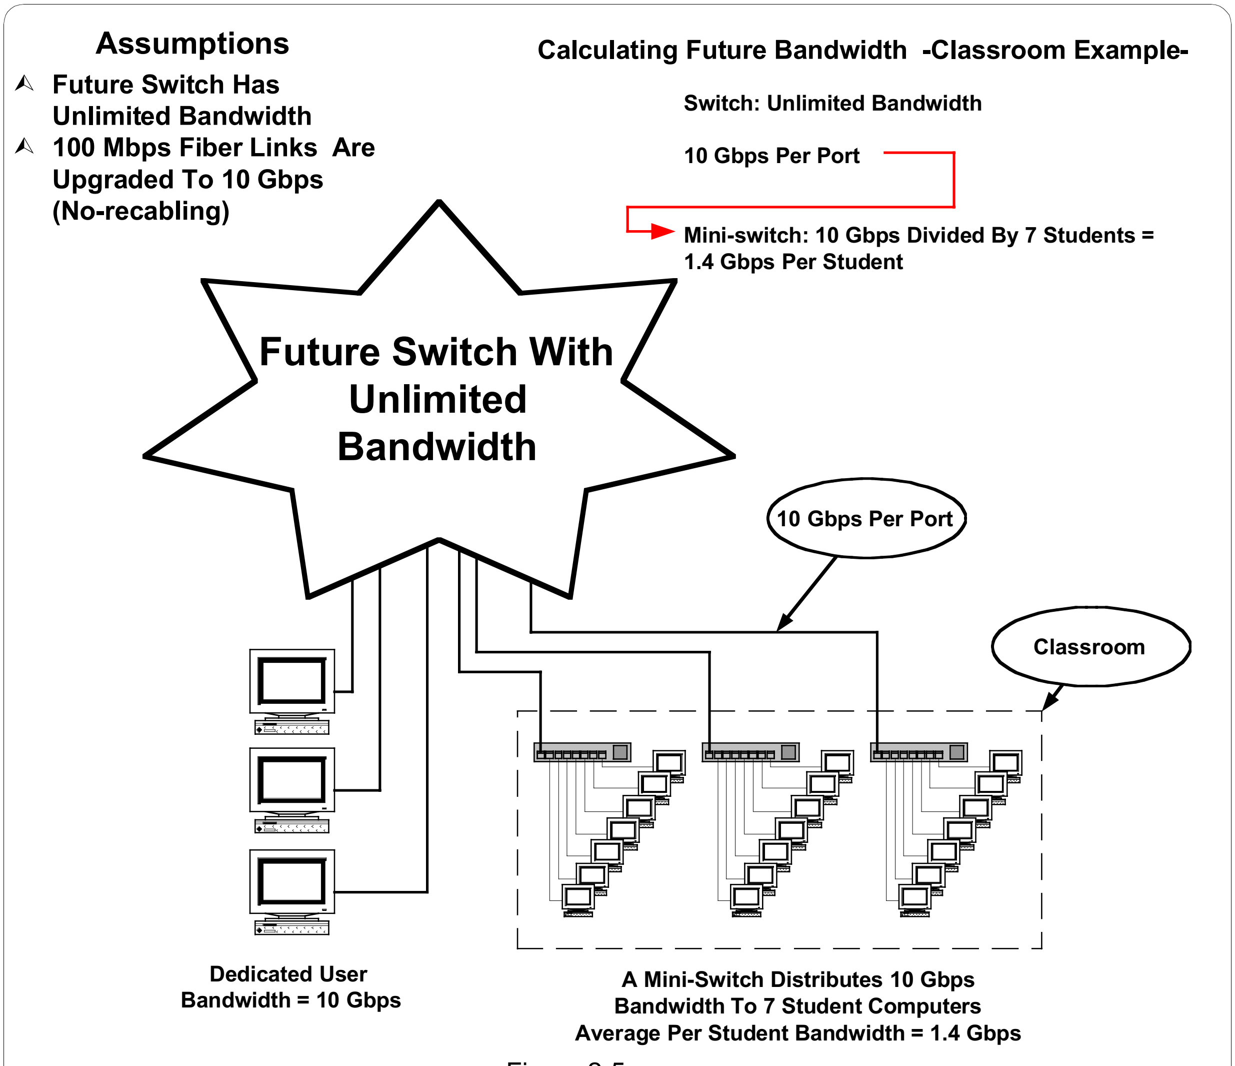

Figure 2-5

Present Day Distributed Networks Compared to Centralized Fiber Networks

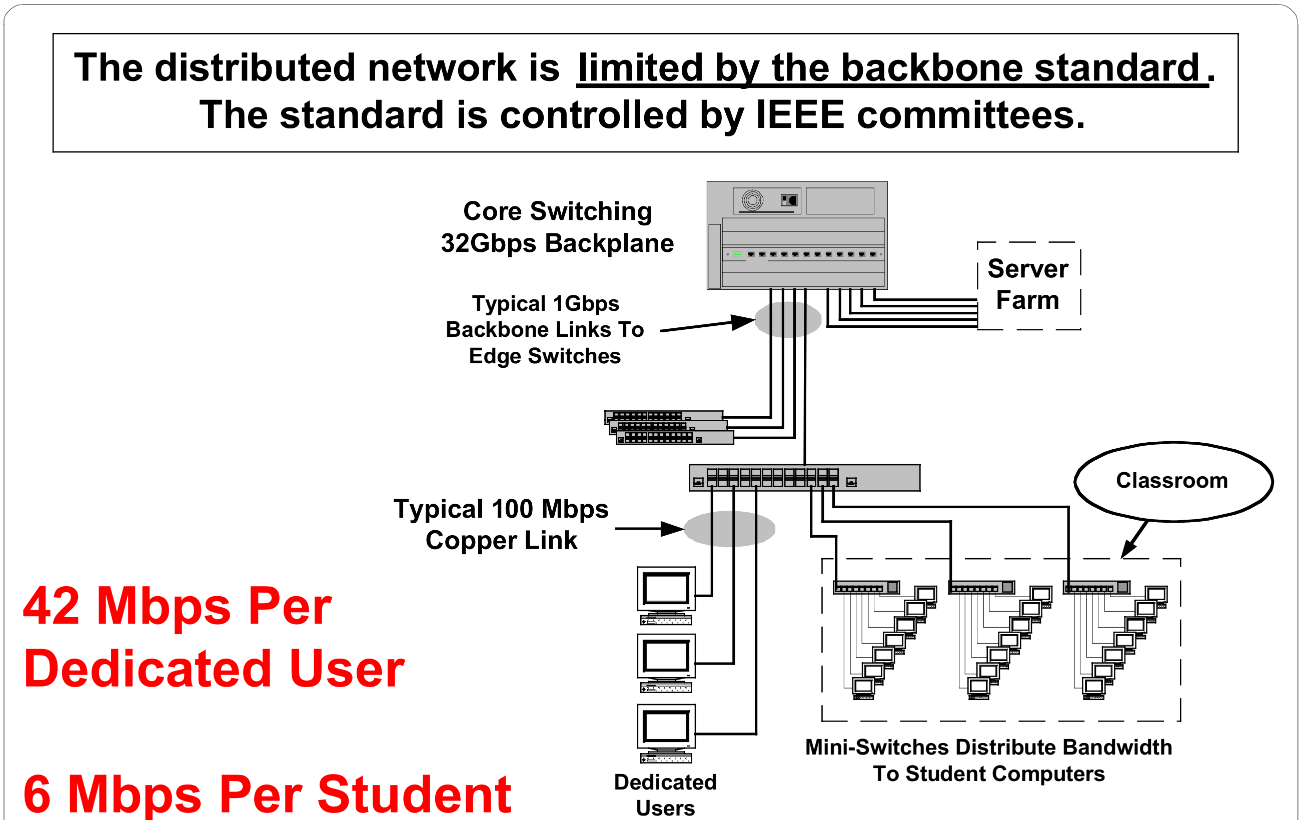

Figure 2-6

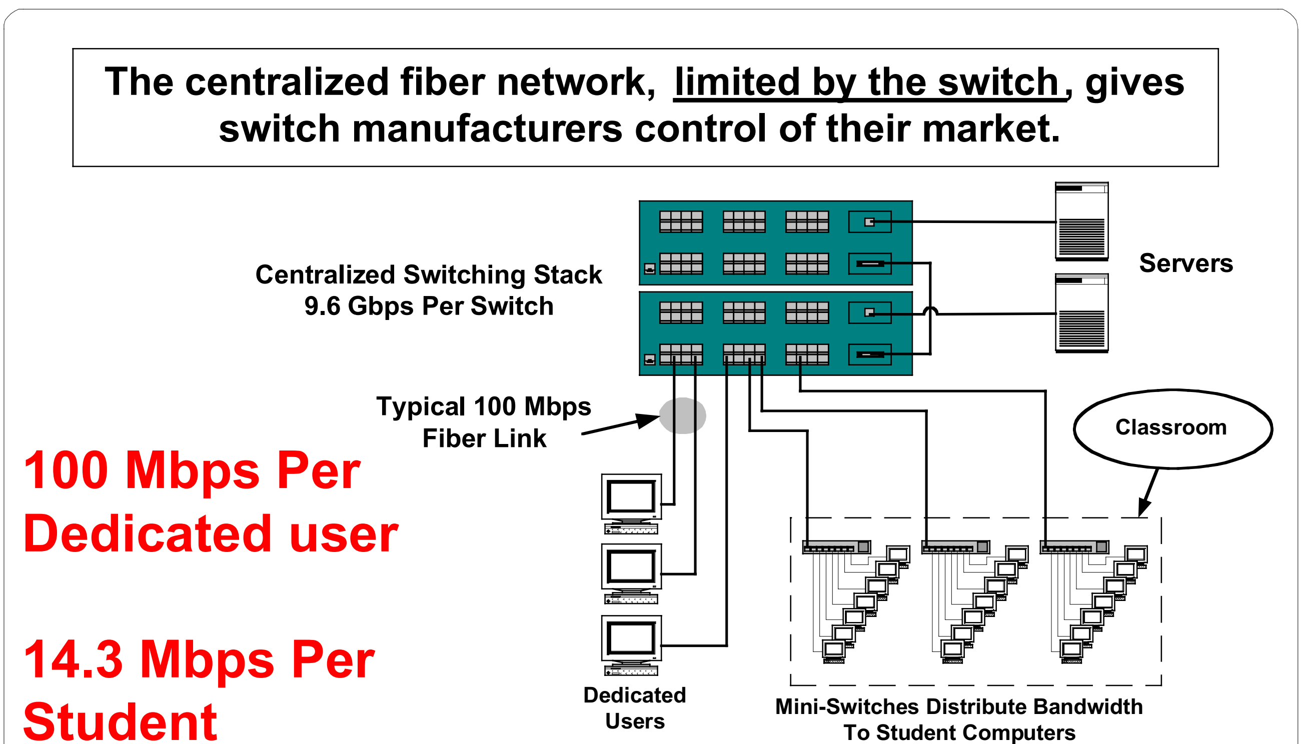

Figure 2-7

Future Distributed Networks Compared to Centralized Fiber Networks

The distributed network edge-delivered bandwidth can be increased <u>10X</u> without re-cabling.

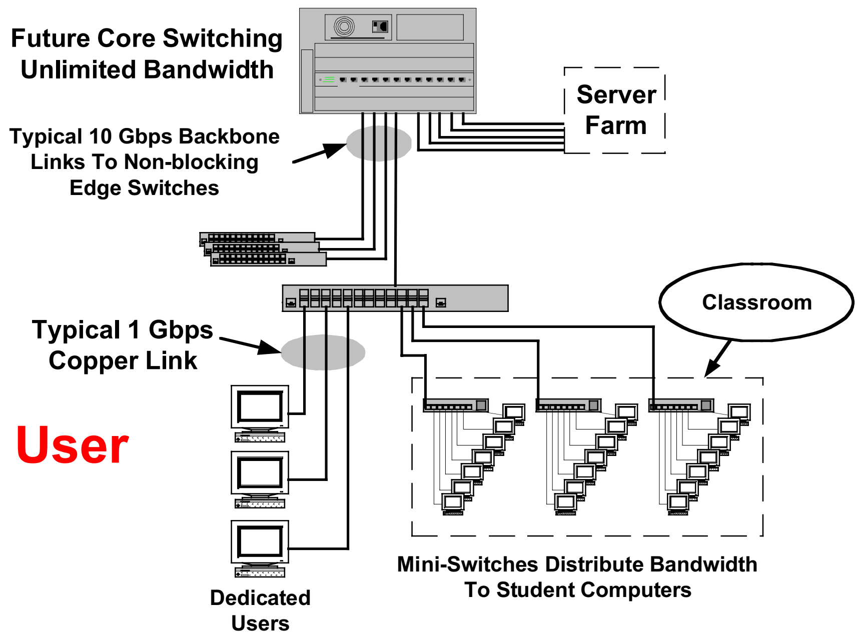

417 Mbps
Per Dedicated User

60 Mbps
Per Student

Figure 2-8

The centralized fiber network edge-delivered bandwidth can be increased <u>60X</u> before applying more exotic technologies like WDM, compression, etc. to increase it further.

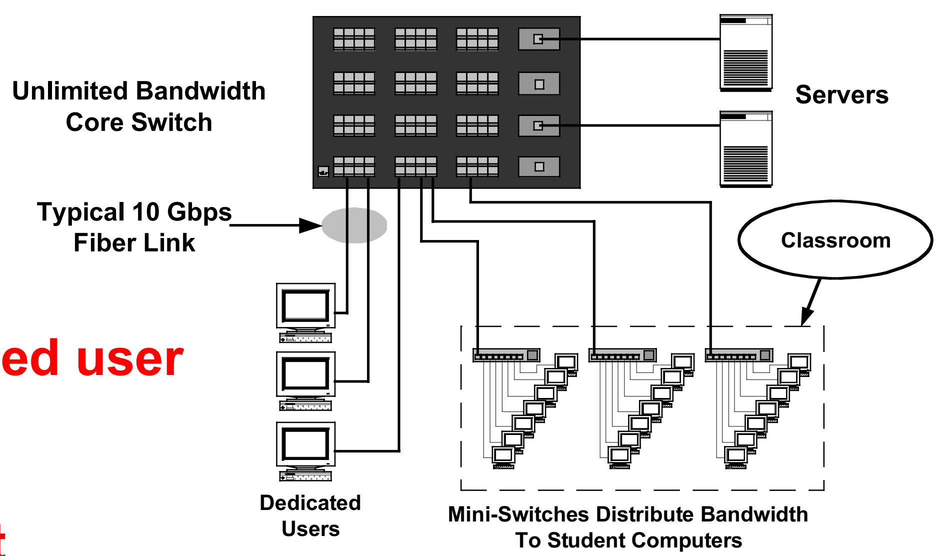

10 Gbps
Per Dedicated user

1.4 Gbps
Per Student

Figure 2-9

Applying a dollar value to the results

Assigning a dollar value to the delivered bandwidth puts this presentation into the proper focus. 3M commissioned a white paper[1] by the Tolly Group to compare traditional distributed network costs to those of a centralized fiber network. Their cost comparison was on a per-user basis. Using the same network designs and costs found in the Tolly Group white paper the cost per megabit delivered can be seen below.

Cat 5e Distributed Network
(16-port edge switches w/gigabit uplink)

267 users @ $962.76 each
Total Estimated Cost: $257,056.92

62.5 Mbps x 267 ports = 16688 Mbps
$257,056 / 16688 Mbps = **$15.40 per Mbps**

Fiber Optic Centralized Network

267 users @ $806.80 each
Total Estimated Cost: $215,415.60

100 Mbps x 267 ports = 26700 Mbps
$215,415 / 26700 Mbps = **$8.07 per Mbps**

The noted centralized fiber network actually costs **<u>48% less</u>** per delivered Mbps than the equivalent distributed network while delivering **<u>160% of the bandwidth</u>**. It should be noted that as the network bandwidths increase, the centralized fiber architecture increases its performance lead over the distributed system quite handily. In fact, comparing this distributed network's bandwidth with future centralized fiber network bandwidth results in a **<u>staggering 160X difference</u>** (62.5 Mbps vs. 10,000 Mbps)!

In conclusion, there is no doubt that properly designed centralized fiber network architectures will deliver less costly and easier to manage bandwidth. Additionally, this type network will provide maximum performance both today and well into the future.

[1] The Tolly Group White Paper No. 200505, Migrating to Fiber: The Case for Centralized Cabling

Sample Network Designs

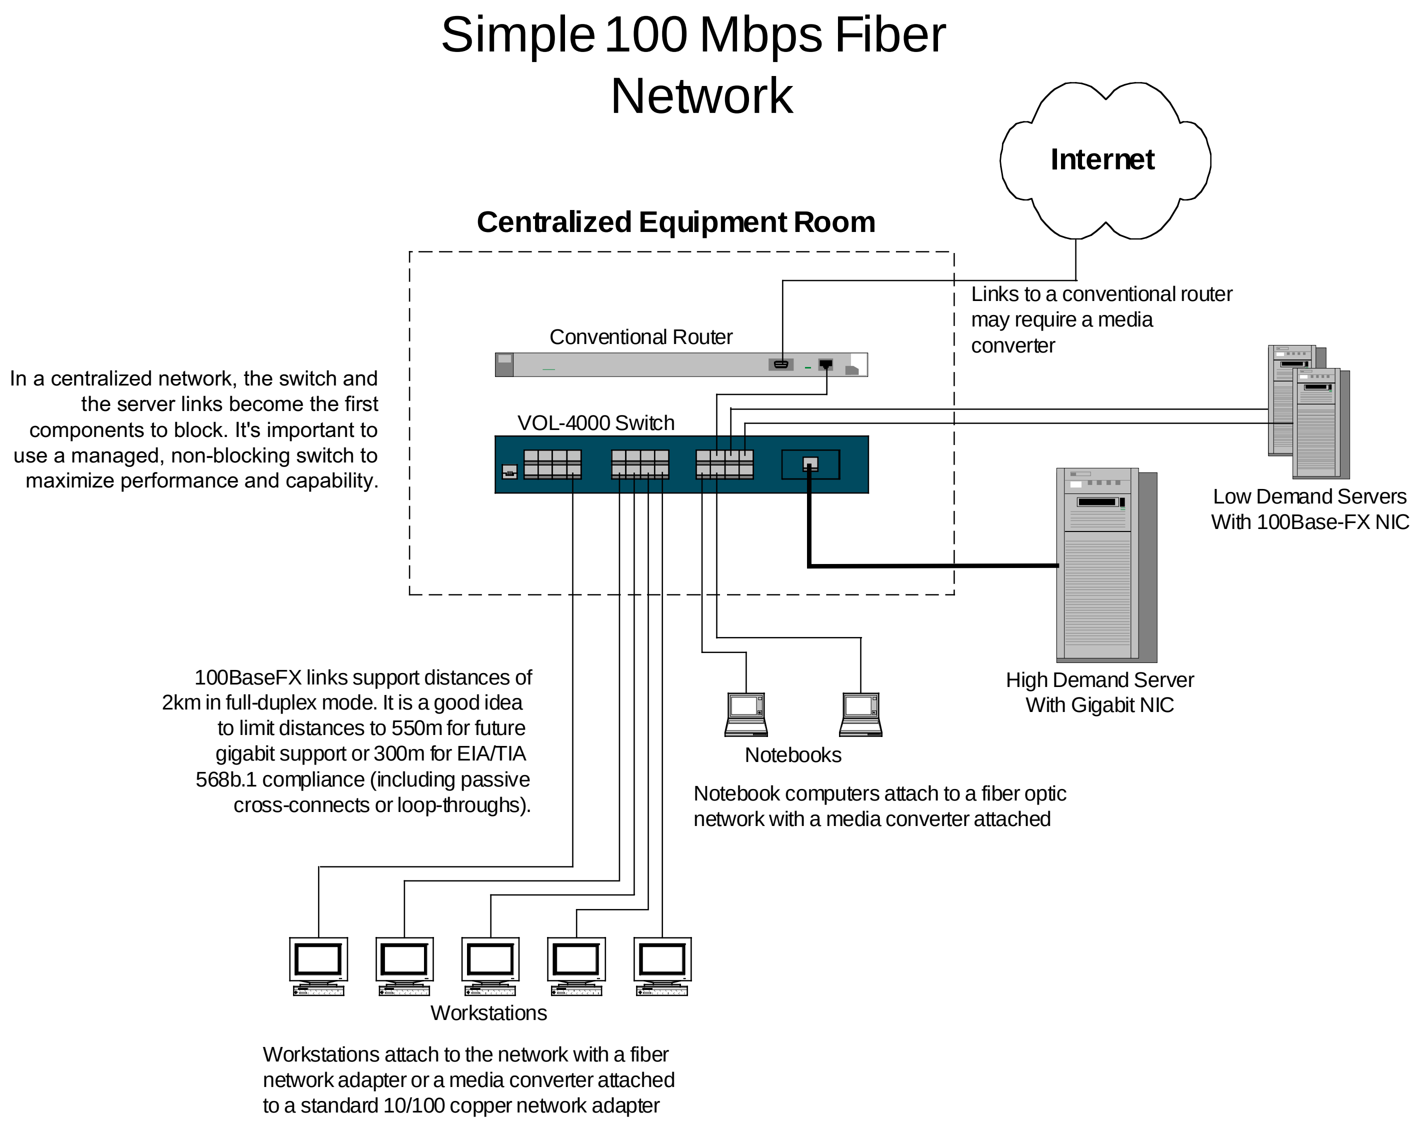
Simple 100 Mbps Fiber Network
Internet
Centralized Equipment Room
Links to a conventional router may require a media converter
Conventional Router
In a centralized network, the switch and the server links become the first components to block. It's important to use a managed, non-blocking switch to maximize performance and capability.
VOL-4000 Switch
Low Demand Servers With 100Base-FX NIC
100BaseFX links support distances of 2km in full-duplex mode. It is a good idea to limit distances to 550m for future gigabit support or 300m for EIA/TIA 568b.1 compliance (including passive cross-connects or loop-throughs).
High Demand Server With Gigabit NIC
Notebooks
Notebook computers attach to a fiber optic network with a media converter attached
Workstations
Workstations attach to the network with a fiber network adapter or a media converter attached to a standard 10/100 copper network adapter

Figure 3-1

96 Port 100Mbps Fiber Network

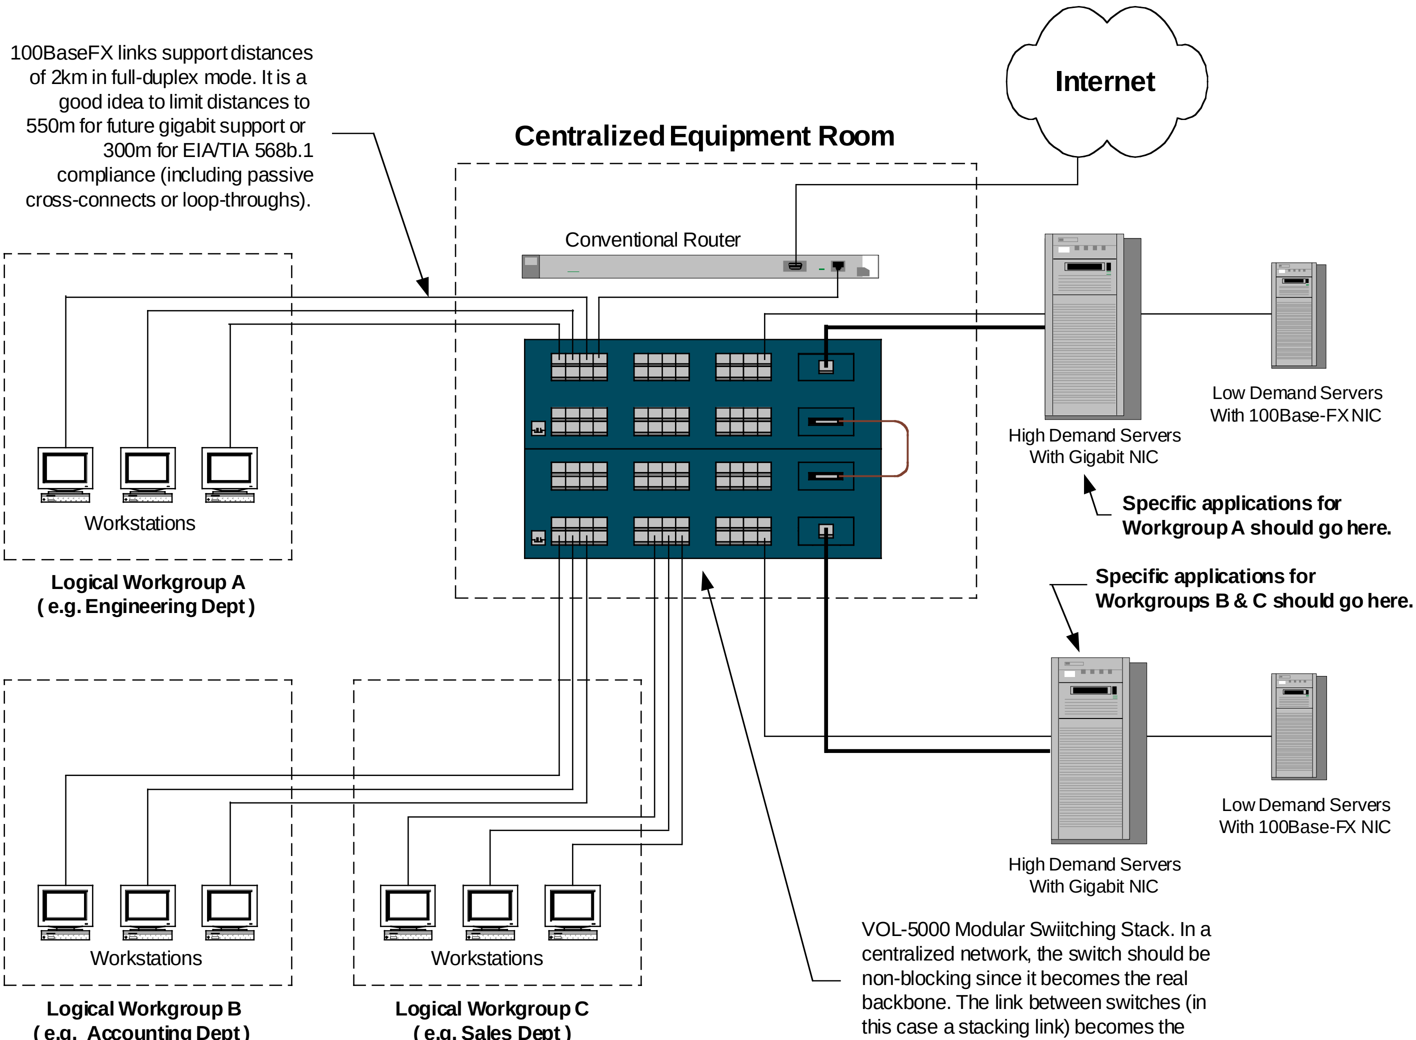

Figure 3-2

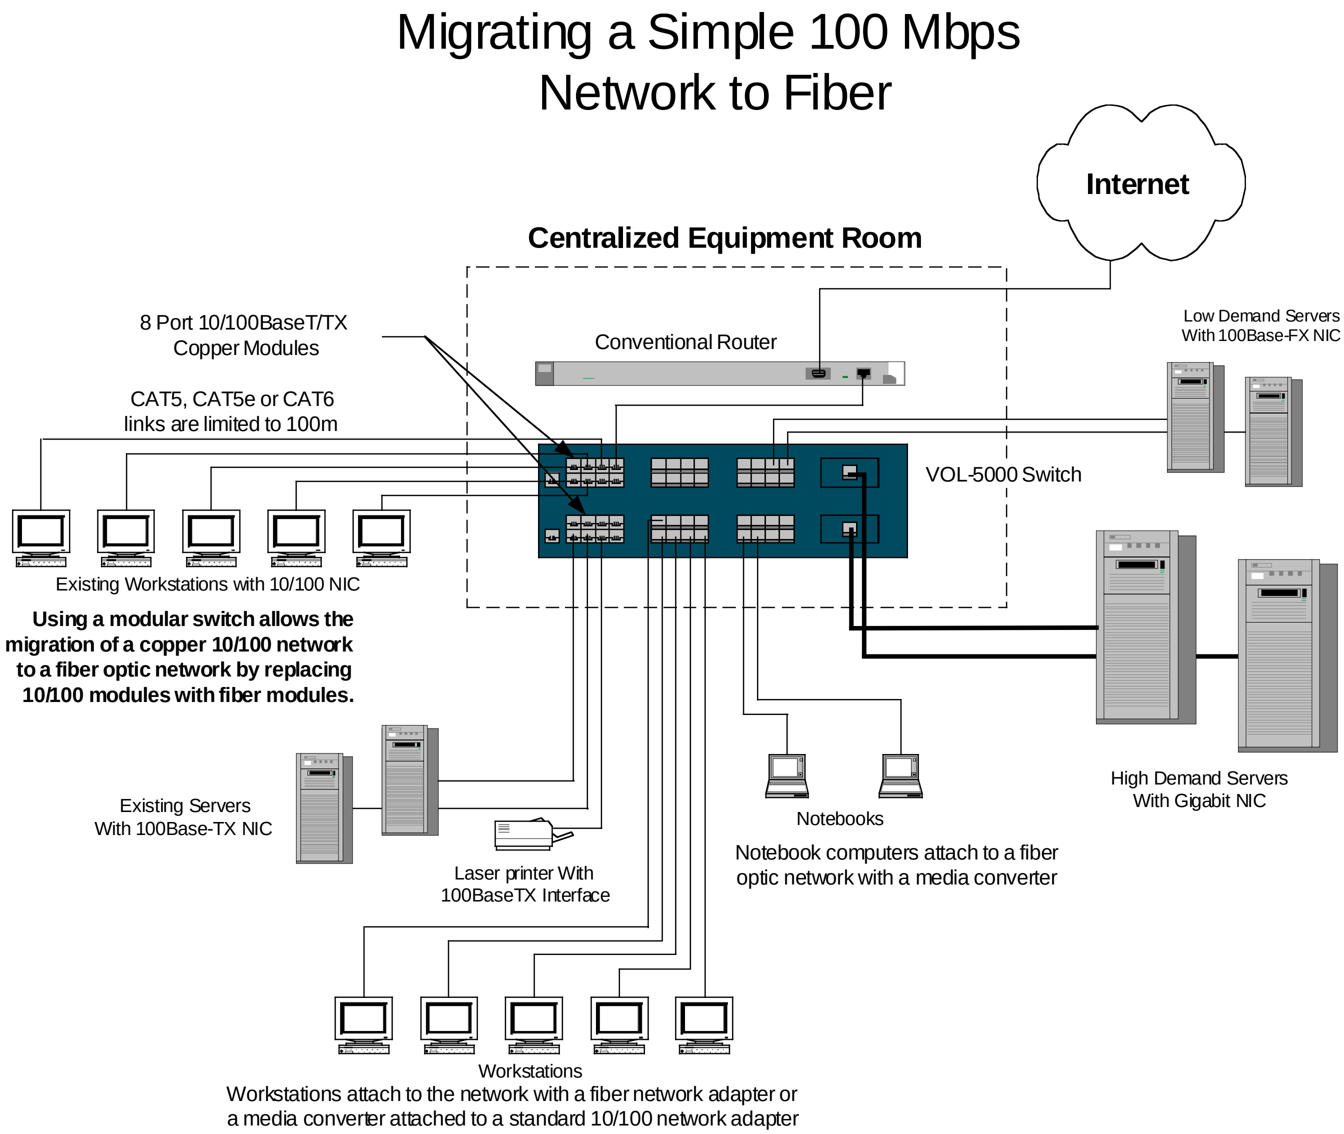
Migrating a Simple 100 Mbps
Network to Fiber
Internet
Centralized Equipment Room
8 Port 10/100BaseT/TX
Copper Modules
Conventional Router
CAT5, CAT5e or CAT6
links are limited to 100m
Low Demand Servers
With 100Base-FX NIC
VOL-5000 Switch
Existing Workstations with 10/100 NIC
Using a modular switch allows the
migration of a copper 10/100 network
to a fiber optic network by replacing
10/100 modules with fiber modules.
High Demand Servers
With Gigabit NIC
Existing Servers
With 100Base-TX NIC
Notebooks
Notebook computers attach to a fiber
optic network with a media converter
Laser printer With
100BaseTX Interface
Workstations
Workstations attach to the network with a fiber network adapter or
a media converter attached to a standard 10/100 network adapter

Figure 3-3

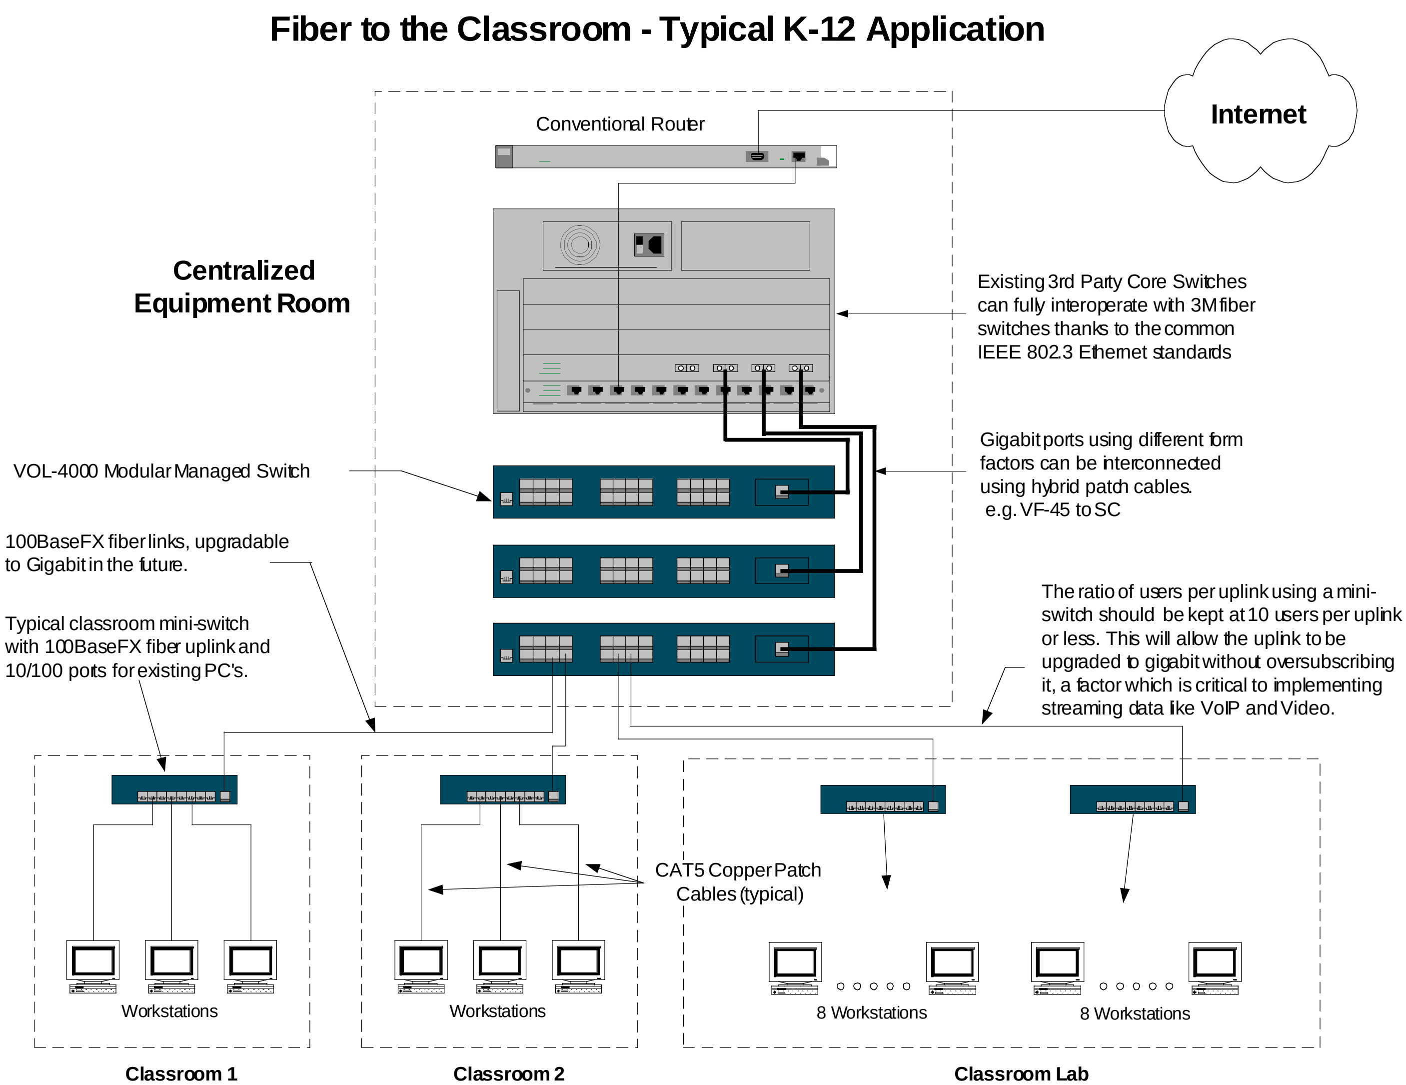

Fiber to the Classroom - Typical K-12 Application
Internet
Conventional Router
Centralized Equipment Room
Existing 3rd Party Core Switches can fully interoperate with 3M fiber switches thanks to the common IEEE 802.3 Ethernet standards
VOL-4000 Modular Managed Switch
Gigabit ports using different form factors can be interconnected using hybrid patch cables. e.g. VF-45 to SC
100BaseFX fiber links, upgradable to Gigabit in the future.
Typical classroom mini-switch with 100BaseFX fiber uplink and 10/100 ports for existing PC's.
The ratio of users per uplink using a mini-switch should be kept at 10 users per uplink or less. This will allow the uplink to be upgraded to gigabit without oversubscribing it, a factor which is critical to implementing streaming data like VoIP and Video.
CAT5 Copper Patch Cables (typical)
Workstations
Workstations
8 Workstations
8 Workstations
Classroom 1
Classroom 2
Classroom Lab

Figure 3-4

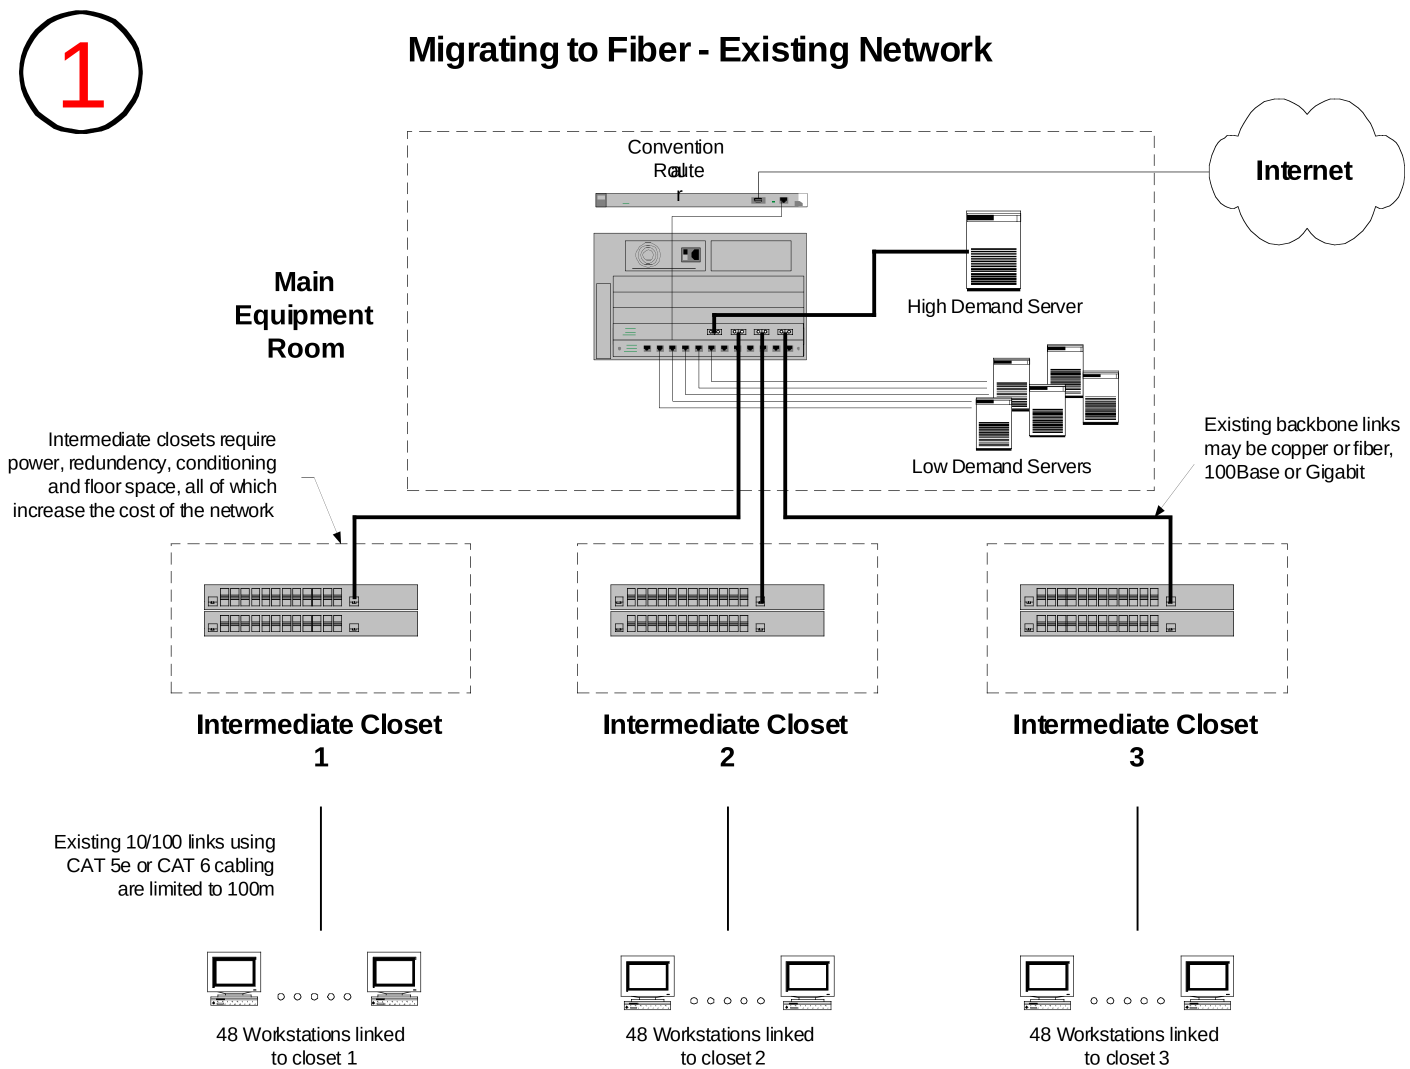
1
Migrating to Fiber - Existing Network
Convention Router
Internet
Main Equipment Room
High Demand Server
Low Demand Servers
Intermediate closets require power, redundency, conditioning and floor space, all of which increase the cost of the network
Existing backbone links may be copper or fiber, 100Base or Gigabit
Intermediate Closet 1
Intermediate Closet 2
Intermediate Closet 3
Existing 10/100 links using CAT 5e or CAT 6 cabling are limited to 100m
48 Workstations linked to closet 1
48 Workstations linked to closet 2
48 Workstations linked to closet 3

Figure 3-5

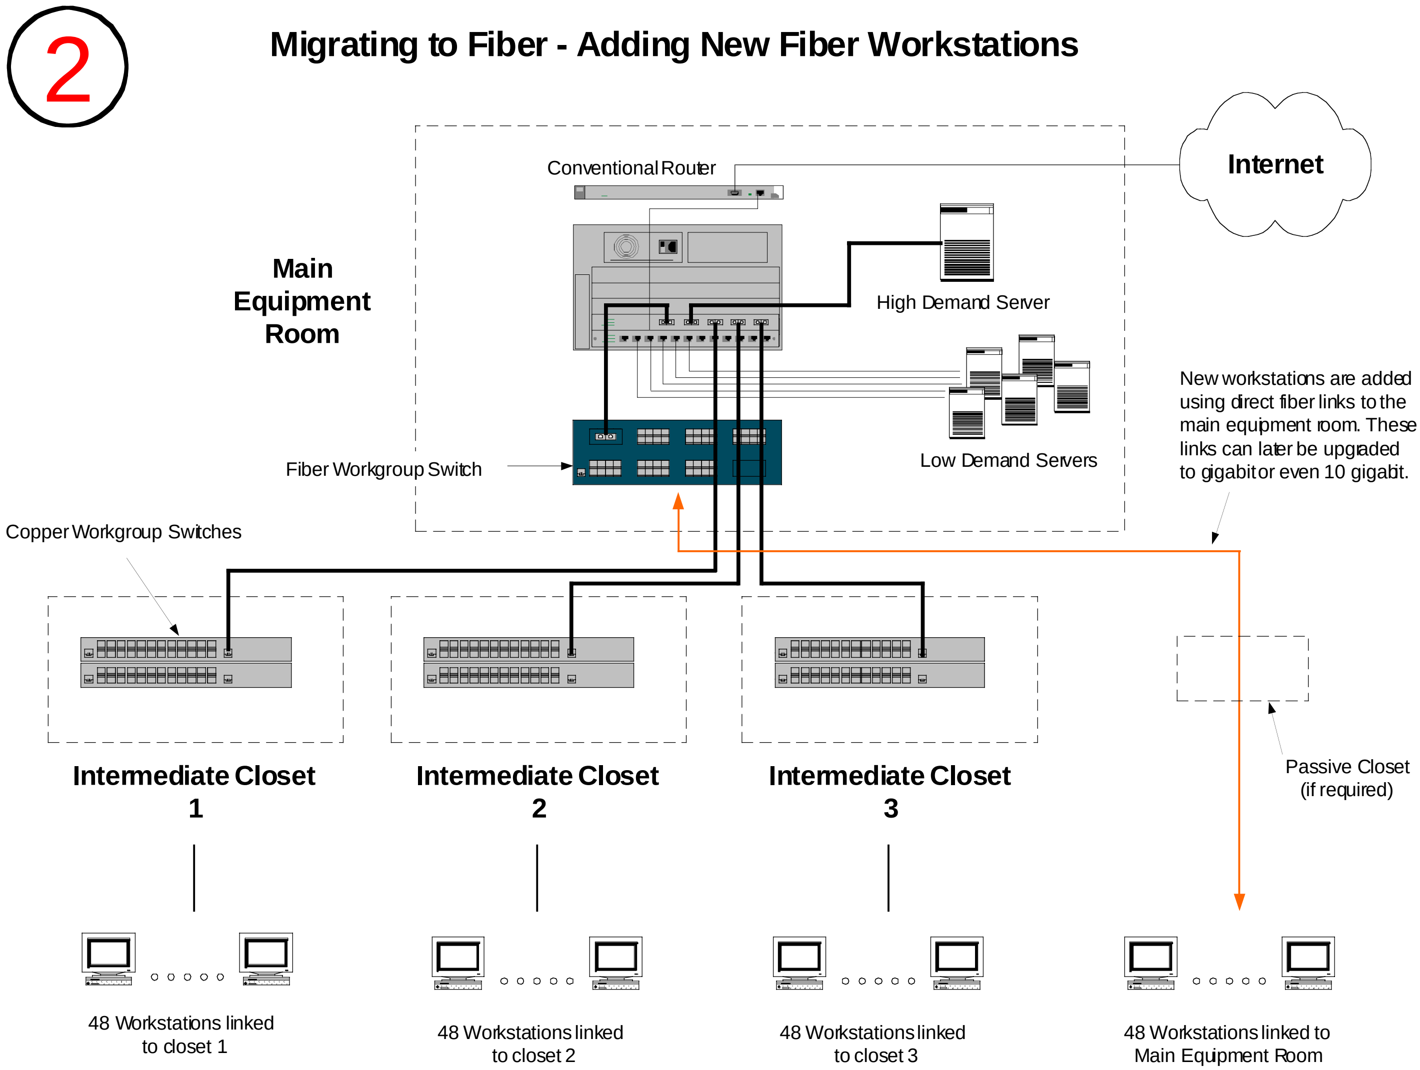
2
Migrating to Fiber - Adding New Fiber Workstations
Conventional Router
Internet
Main
Equipment
Room
High Demand Server
New workstations are added using direct fiber links to the main equipment room. These links can later be upgraded to gigabit or even 10 gigabit.
Low Demand Servers
Fiber Workgroup Switch
Copper Workgroup Switches
Passive Closet
(if required)
Intermediate Closet
1
Intermediate Closet
2
Intermediate Closet
3
48 Workstations linked
to closet 1
48 Workstations linked
to closet 2
48 Workstations linked
to closet 3
48 Workstations linked to
Main Equipment Room

Figure 3-6

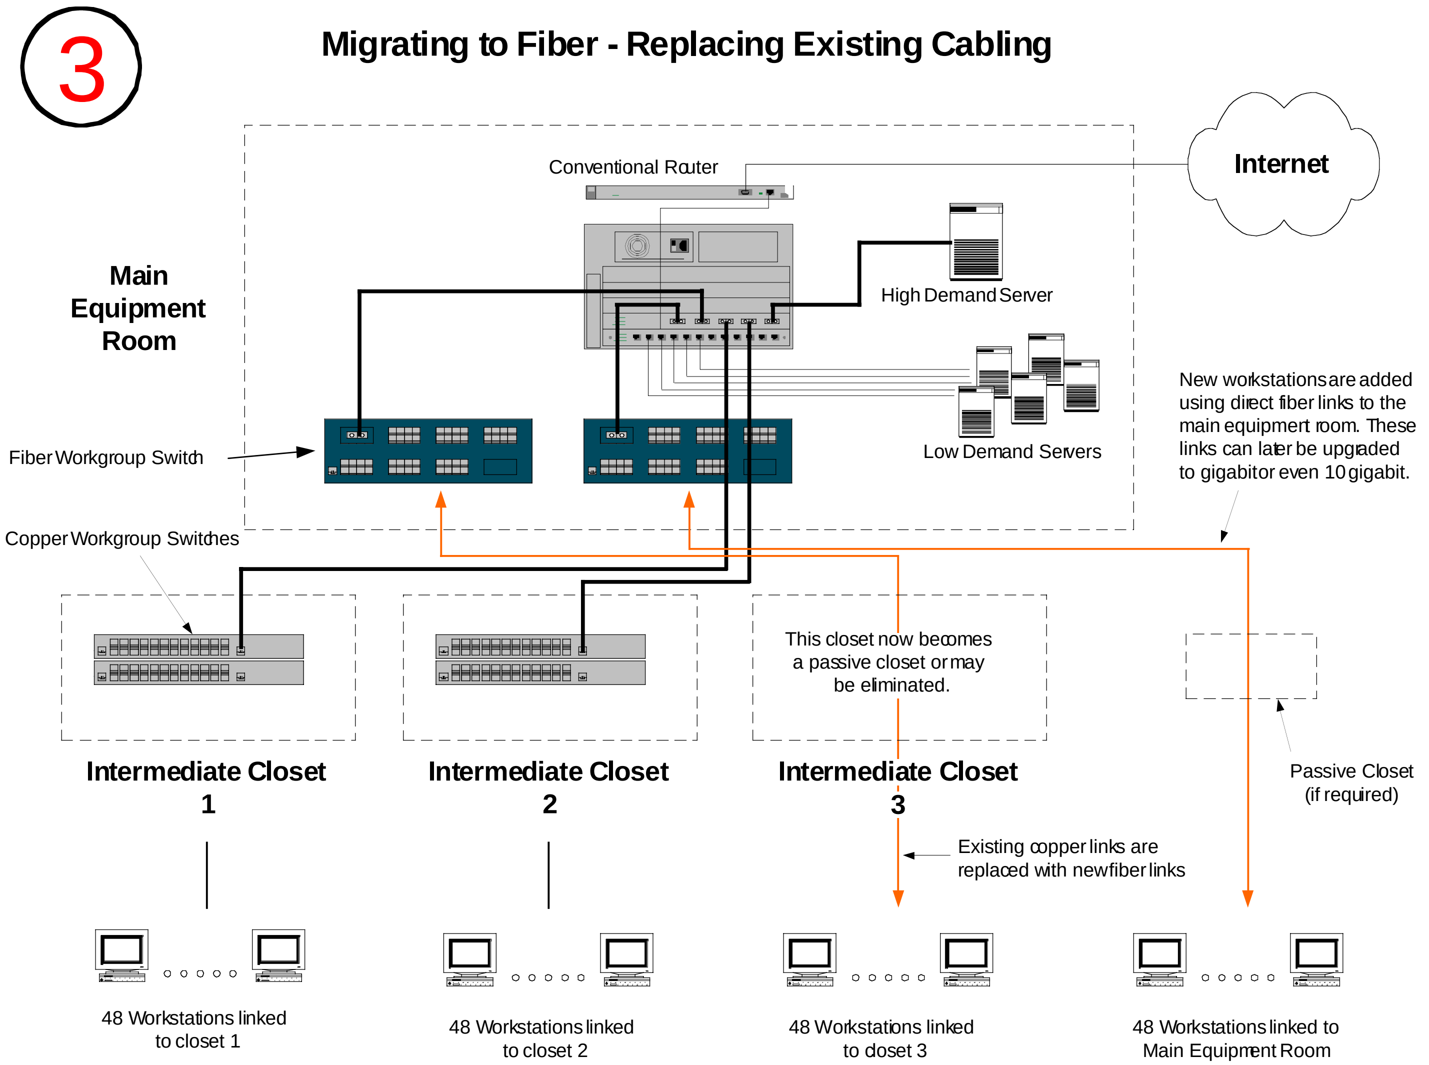
3
Migrating to Fiber - Replacing Existing Cabling
Internet
Conventional Router
High Demand Server
Low Demand Servers
Main Equipment Room
Fiber Workgroup Switch
Copper Workgroup Switches
New workstations are added using direct fiber links to the main equipment room. These links can later be upgraded to gigabit or even 10 gigabit.
This closet now becomes a passive closet or may be eliminated.
Intermediate Closet 1
Intermediate Closet 2
Intermediate Closet 3
Passive Closet (if required)
Existing copper links are replaced with new fiber links
48 Workstations linked to closet 1
48 Workstations linked to closet 2
48 Workstations linked to closet 3
48 Workstations linked to Main Equipment Room

Figure 3-7

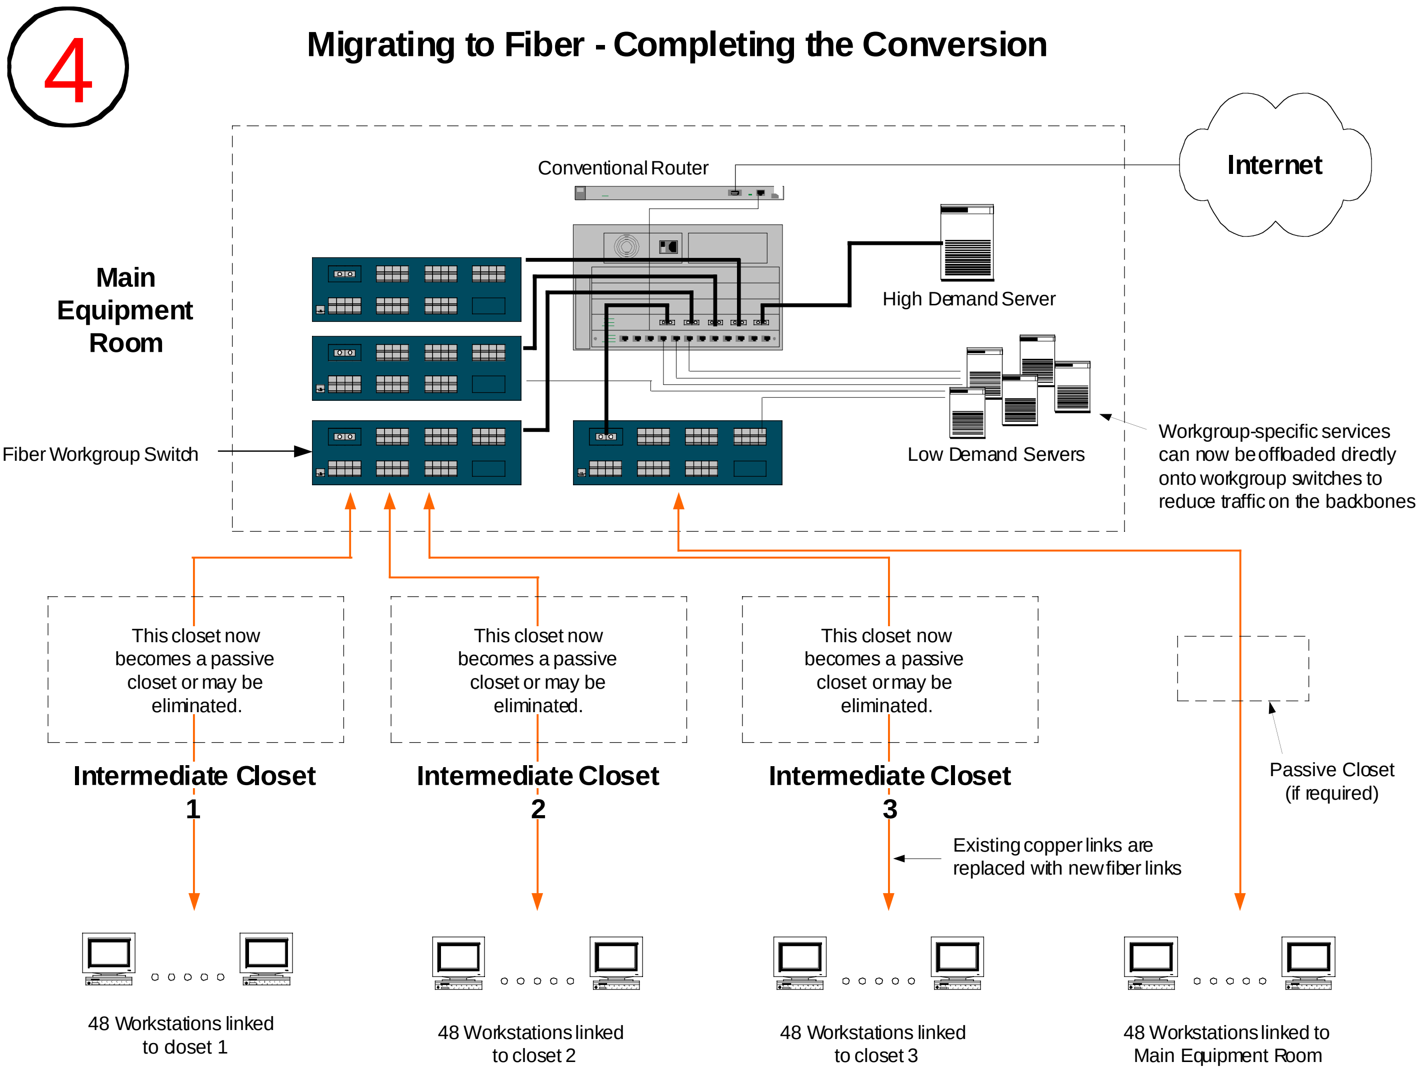
4
Migrating to Fiber - Completing the Conversion
Main Equipment Room
Conventional Router
Internet
High Demand Server
Fiber Workgroup Switch
Low Demand Servers
Workgroup-specific services can now be offloaded directly onto workgroup switches to reduce traffic on the backbones
This closet now becomes a passive closet or may be eliminated.
This closet now becomes a passive closet or may be eliminated.
This closet now becomes a passive closet or may be eliminated.
Passive Closet (if required)
Intermediate Closet 1
Intermediate Closet 2
Intermediate Closet 3
Existing copper links are replaced with new fiber links
48 Workstations linked to closet 1
48 Workstations linked to closet 2
48 Workstations linked to closet 3
48 Workstations linked to Main Equipment Room

Figure 3-8