A key to deploying enterprise client/server networks is a design approach called *structured networking.*

Structured networking is a precise, highly organized design scheme for building large-scale client/server networks that ensures production-quality performance and reliability, plus sophisticated management.

Introduction

This paper examines many key requirements for client/server networks, explores important design considerations, and describes the specific recommendations of structured networking. What structured wiring did for your LANs in the 1980s, structured networking does for your enterprise networks in the 1990s.

As more enterprises develop mission-critical applications using network-based platforms consisting of Novell, Microsoft, Banyan, UNIX, and others, client/server computing is entering a new phase in its evolution. The new application development efforts leverage high-performance, intelligent desktops and powerful new tools and technologies to deliver a number of the promised benefits of client/server computing: lower cost for hardware and software; faster development cycles; lower maintenance burdens; and, most importantly, better applications that are easier to use and provide access to a wider array of information resources.

The early applications of client/server computing enabled workgroups and departments to take advantage of the file- and resource-sharing benefits of networking. By contrast, the new generation of client/server applications is being implemented using new software development technologies, languages, user interfaces, database engines, and network operating systems to create complex applications for use in operations such as airline reservations, order entry, human resources, materials resource planning, customer service, billing and accounts payable, and inventory tracking systems.

The emergence of these complex, enterprise-wide client/server applications is placing new demands on the networks that must support them. In today's client/server computing model, the network itself must perform many of the functions traditionally associated with mainframe computers in the terminal-host computing model. In other words, in the client/server computing model, the network itself is becoming the distributed computer backplane (Figure 1).

These new client/server applications, and the increasing importance of the network, require new techniques for designing and building networks. The foundation of these networks is an underlying system architecture that Bay Networks calls the *network fabric*.

The network fabric is a managed, high-speed communications system that supports new classes of enterprise and multienterprise applications such as electronic commerce, desktop videoconferencing, and medical imaging. With the intelligent hub as its key building block, the network fabric seamlessly integrates emerging technologies with existing network equipment to create a transparent, self-managing system that easily expands and evolves as an organization grows.

Bay Networks network fabric is based on Bay Networks intelligent hubs, which serve as connectivity, internetworking, and network management platforms that can be configured in various ways to create a powerful system architecture that delivers the specified performance and functionality. The Bay Networks Optivity® network management system manages all the elements of the architecture as a cohesive system. New technologies can be introduced into the existing network fabric simply by configuring hubs to support them or by introducing standalone devices that can be managed by Optivity. The advantage of this approach is that the network fabric can integrate current networking technologies, legacy systems, and emerging technologies seamlessly, offering an organization enormous investment protection benefits.

Figure 1 | **Terminal-Host Versus Client/Server Systems**

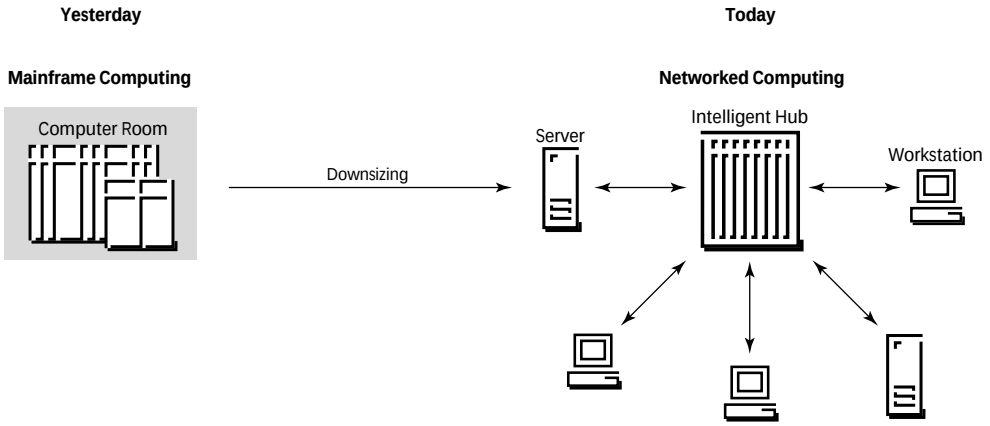

This paper addresses how the emergence of client/server applications is driving the need for structured networking and, specifically, network centers to ensure greater reliability and manageability through centralization. It also discusses the need for a new class of intelligent hub that resides in the network center.

The shift to more sophisticated client/server applications establishes a new set of mission-critical requirements for intelligent hubs that must be embodied in network center hubs to ensure flexibility and reliability.

Localization Gives Way to Distributed Workgroups

In early client/server systems, it was possible to keep most clients on the same local area network (LAN) segment as their server. Today, with a greater number of people in different locations requiring access to these new applications, that type of localization is no longer feasible. Workgroups today normally consist of users spread across buildings in a business campus or across floors in an office tower. Traditionally, these distributed workgroups consist of two or more LAN segments or rings connecting users at different wiring closets. Since desktops will need to reach an increasing number of servers, each one running different applications and potentially located on a different LAN or at different sites, delivering adequate performance for these nonlocalized client/server applications is a challenge.

As a result of these changes, in this new era of enterprise-wide client/server applications, central IS organizations are playing a much greater role in developing, deploying, and managing these applications and the networks that support them. A number of these new applications actually are being developed by central IS rather than by individual workgroups or departments. Consequently, more traditional thinking about operations, management, change control, and user support is now being applied.

Many users welcome central IS's new role in client/server computing because they realize it will mean greater stability, reliability, and data integrity. However, as central IS solidifies its role as the developer and manager of client/server enterprise applications, there is renewed concern about how to manage and control these new distributed systems. Enterprises across a wide range of industries are now struggling to establish high-performance client/server computing systems. Beyond the procedural changes that must take place, such as new testing and change-control procedures, there are a number of technical changes that must take place in the deployment and usage of computing and networking systems to ensure stability (see Figure 2).

To prepare for new generations of client/server applications, enterprises must address two fundamental challenges: They must achieve more reliable server operation and they must achieve much more reliable network operation.

More Reliable Server Operation

Servers are the main repositories of data and software in client/server systems. Although desktops provide the crucial user interface into the applications, most of the processing and data are found on servers. Placing client/server applications into production means that servers of all varieties now must deliver the same level of reliability, protection, and control as mainframes or minicomputers.

Figure 2 | **The Swinging Pendulum of Computing**

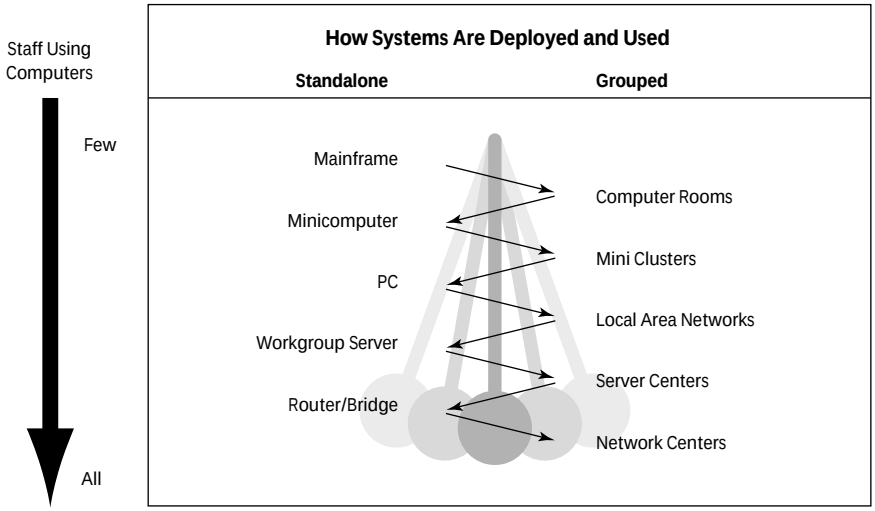

This points to the need for server centers — controlled areas that house potentially large numbers of servers — to provide centralized service to distributed, departmental workgroups (see Figure 3). Server centers are physically secured, protecting corporate data against theft or destruction. Backup battery power supplies ensure that minor power outages or brownouts will not crash servers and corrupt their databases. Server centers also include halon or other rapid-fire extinguishing systems to protect against fire.

A server center functions in much the same manner as a data center, except that it houses numerous separate small machines rather than a few large ones. Another difference is that server centers may not have the same stringent environmental needs as rooms designed to house massive mainframes requiring water cooling or major amounts of air-conditioning. For organizations consolidating their mainframe systems into one or two data centers, there is a large amount of old data center space available for converting into server centers.

Server centralization offers a number of other benefits. Placing dozens of servers in a single room simplifies rapid repair or upgrades, which was much more difficult when the equipment was dispersed across a major site. Today, server centralization also facilitates such functions as software updating and backup, although new network-based system management tools also make it possible to perform these functions remotely.

Choosing the optimal number of server centers for an enterprise, as well as the best way to distribute the servers for different kinds of applications, will take many years to sort out. Design tools are emerging that will help client/server developers and system managers begin to analyze performance engineering issues — such as the right number of servers and users per server — for a particular application.

The trend right now is to centralize servers at the site level, or at least within each building. The cost of wide area network capacity is still too high to think about centralizing large numbers of servers at a single, corporate headquarters site, since client/server interactions have generally been designed for use over high-speed LANs, not slower speed wide area circuits.

Figure 3 | **Workgroup Server Centralization**

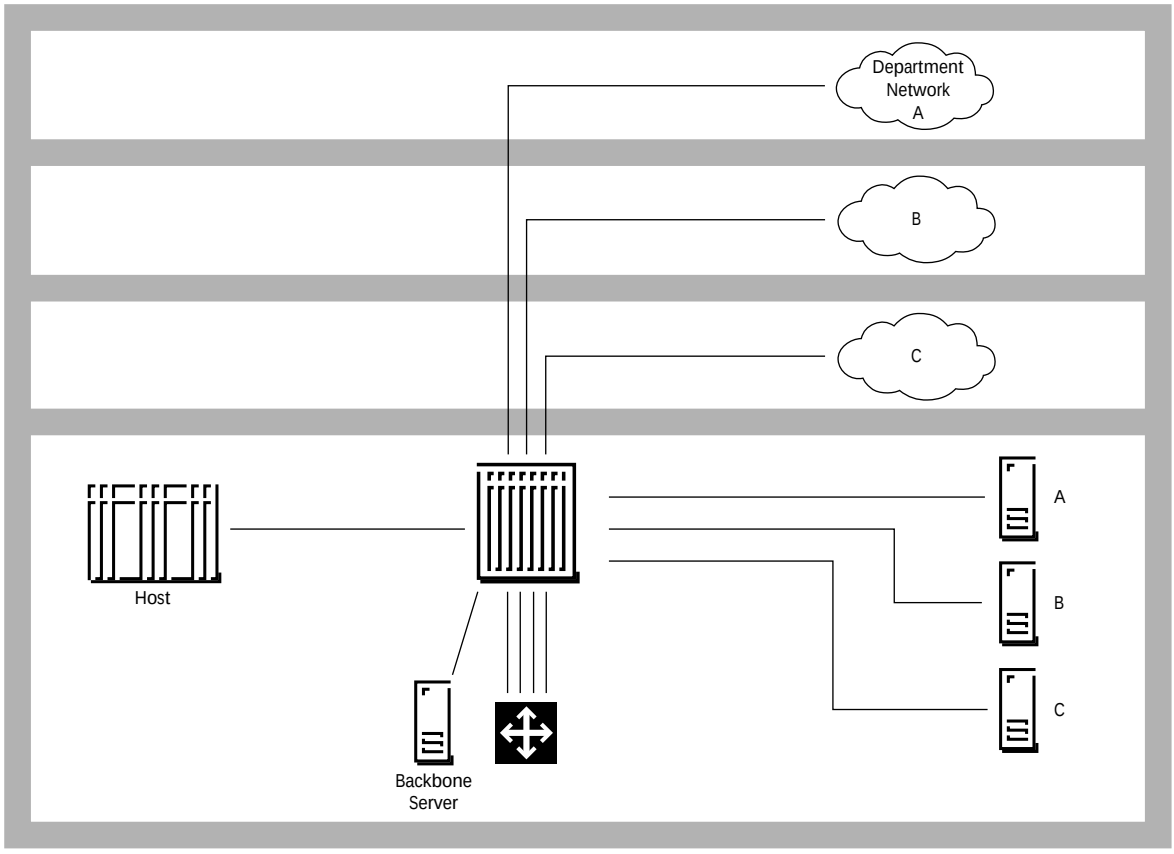

As a result, many enterprises are working today to centralize servers at contiguous sites within a facility where network capacity is less expensive and high-speed, fiber-based networking technologies make it practical to rapidly interconnect large numbers of clients and servers. Since server center space is more expensive than normal office space, site centralization — centralizing servers at contiguous sites — offers the best balance between the cost of network bandwidth and the cost of the server centers themselves.

Increasing server availability also means increasing the reliability of its network links. As more users access servers, and as the amount of data moving between client and server or among servers continues to increase (due to use of imaging, multimedia, and other complex applications), there is growing pressure to increase server network interface speeds. The intense use of servers by large groups of engineers seeking file access already exceeds the capacity of standard IEEE 802.3 Ethernet network interfaces.

This problem can be solved by moving to a higher-speed network interface such as Fiber Distributed Data Interface (FDDI). It can also be accomplished — perhaps less expensively — by employing multiple Ethernet or Token Ring interfaces (see Figure 4). Another major advantage of using multiple network interfaces is the increased availability resulting from having a server homed to multiple networks with separate connections to backbone LANs.

Using multiple LAN interfaces to access a common server also breaks the client population into multiple groups, each on their own LAN segment. This segmentation cuts down on contention for LAN capacity and increases the overall bandwidth available to users. Mission-critical applications frequently require servers with multiple network interfaces and even more reliable, fault-tolerant network configurations.

More Reliable Network Operation

Putting client/server computing into production operation means that the enterprise will need a more reliable network to support ever more intense client-to-server and server-to-server interactions. While older legacy applications relied on relatively simple networking technologies such as leased lines and terminal-to-host protocols, client/server applications depend on more complex LAN internetworking technology and network operating system protocols.

Placing client/server networks into production operation also means increasing the reliability of internetworking — and rationalizing the LAN environment. Enterprises must evolve beyond LAN anarchy in which different workgroups set up and manage their own LANs, including buying their own bridges and routers. Network planners and operations groups must extend their efforts more deeply into the site-level LAN infrastructure, develop standard guidelines for LANs, and develop building-level and campus-level backbone topologies. Enterprises can no longer base their LANs on a foundation of cables snaking through office ceilings, racks of PC-based bridges with little or no network management, and an assortment of black boxes from a variety of suppliers.

Structured, hierarchical, physical star wiring architectures and intelligent hub-based LAN implementations are two essential steps to rationalizing the building-level network environment and making it a solid base for enterprise-scale, client/server computing. Implementing LANs using intelligent hubs and structured wiring creates LAN infrastructures that are more easily managed. Troubleshooting and fault isolation are achieved more rapidly because each end station is attached to the LAN on its own individual port — one that can be monitored individually and, if needed, easily turned off. Moves, additions, and changes are also simplified when the wiring and hubs are standardized throughout.

Figure 4 | **Server Attachment Techniques**

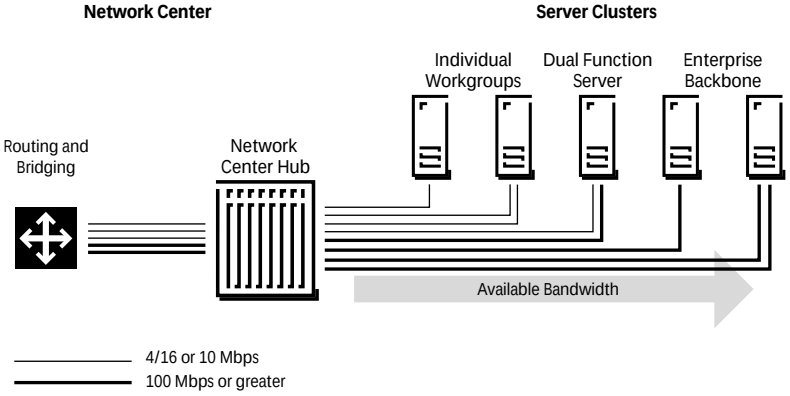

Intelligent hubs provide end-user connectivity while delivering backbone resources to individual workgroups. They combine this structured approach to wiring distribution with an open, standards-based management architecture. Modular hubs make it simple to add ports, routers, or bridges at minimal cost. Intelligent hubs implement a physical star-wired topology that provides each device with its own port, delivering a new level of management visibility, control, and statistical collection. Designed for a high density of connections and ports, hubs make maximum use of small wiring closet space.

Network security also is improved when hubs are located in locked equipment closets where individual network interfaces can be activated and disabled remotely from a network management system. Intelligent hubs provide an increasingly automated substructure that supports higher and higher levels of self-healing capabilities and adaptation to changing network conditions.

Advanced intelligent hubs go even further by monitoring all traffic on each port and providing the ability to monitor what each end-user device is doing. This monitoring of device activity makes it possible to link logical network information, such as MAC or IP addresses, to physical network locations such as a port number on a specific module in a specific hub in a specific wiring closet. Information of this type is extremely valuable for network troubleshooting, repair activities, and security and access control functions.

With greater emphasis on enterprise-wide client/server applications, there is growing pressure to internetwork all enterprise LANs into a private internet (see Figure 5). The resulting internet provides a single utility for all data communications.

Most enterprises have adopted routing as the basis for their enterprise internets. Routers use the information contained in network protocols to intelligently forward packets through the internet. They support arbitrary mesh topologies of any diameter, plus they support highly sophisticated dynamic adaptive routing algorithms that allow the network to recover from link and node failures if backup paths are available.

As might be expected, routers are complex devices. As a result, there is a strong tendency today to place routers under the control of the corporate networking group and standardize them. Minimizing the differences between routers and standardizing on a single, robust routing update protocol are important steps in creating an enterprise-wide client/server network. Another important strategy to combat router complexity is to centralize routers within a building or site, minimizing the number of routers that must be configured and managed.

Figure 5 | **Today's Internets**

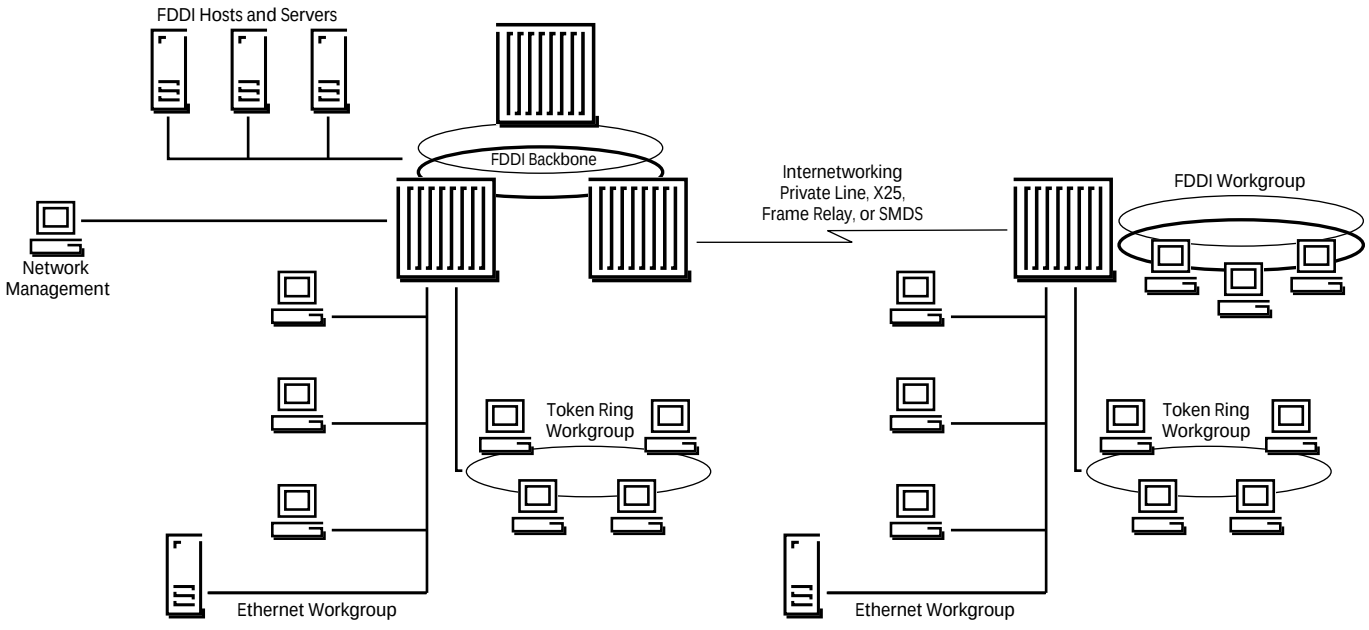

Fewer routers also means fewer "hops" for client/server traffic to traverse. The performance of client/server applications degrades with each router hop due to the delays that occur each time a packet is read into a router, a routing decision is made, and the packet is read back out of the router onto a new interface. To maintain high-performance levels, network designers are seeking to minimize the number of routers in the path between clients and their servers. The ability to do so depends on both the backbone topology and the location of the servers relative to the backbone.

There are two primary backbone architectures: distributed backbones and collapsed backbones. Both approaches represent segmented network architectures, with the primary difference being the use of distributed or centralized routers or bridges. Table 1 describes how backbone architecture relates to the placement of servers and the number of internetworking hops that result from each design.

Continuing to provide a high level of services to the growing number of client/server applications and their users requires a flexible internet infrastructure that can be easily changed to meet shifting requirements. Today, no one is capable of accurately estimating the network traffic demands created by these new client/server applications — especially those using the latest multimedia and imaging technologies. Application designers and network engineers must still do a tremendous amount of fine tuning in terms of server placement, LAN segment sizes, LAN internetworking capacities, and application behavior.

Enterprise network managers will need to quickly and easily resegment their LANs. Breaking a LAN in two and connecting the resulting segments with a bridge or router increases available bandwidth due to the filtering properties of the internetworking device. LAN segmentation strives to localize the majority of traffic within a single LAN segment so that bandwidth on other LANs — particularly backbone LANs — is not consumed unnecessarily.

In addition, network managers will need the ability to easily add new servers or additional interfaces to existing servers to meet growing requirements for more processing or network access capacity. The enterprise internet must be based on a modular, flexible network hardware implementation that can support large numbers of LAN segments using a variety of network interface speeds.

And perhaps most importantly, the enterprise internet must be monitored at a very detailed level to provide comprehensive

Table 1 **Impact of Backbone Topology and Server Placement on Internetworking Hops**

	Collapsed Backbone	Distributed Backbone
Centralized Workgroup Servers Ethernet- or Token Ring-attached. Located in data, server, or network center. Used for word processing, printer sharing, and other applications unique to workgroup.	No hops required	2 hops minimum
Distributed Workgroup Servers Ethernet- or Token Ring-attached. Located in wiring closets or near clients. Used for word processing, printer sharing and other applications unique to workgroup.	No hops required	No hops required
Enterprise Backbone Servers FDDI-attached. Located in data, server, or network center. Used for e-mail, shared applications, gateway services, wide area network access, and other applications common across workgroups.	1 hop minimum	1 hop minimum

measurements of network traffic, loading, and performance. Network management tools must be able to use the network measurement data to analyze specific applications to help application developers and network engineers make intelligent decisions about tuning and reconfiguring both network and server assets.

The network must become a fully capable application measurement system for the new generation of network-based applications. These network measurement and analysis capabilities must be built into the very fabric of the network, structured to provide meaningful, actionable information to the network manager.

Structured Networking for the Enterprise

Demands for more reliable server and enterprise internet operation are creating a new set of requirements for site-level network topologies. The site-wide network design problem is the focus of intense scrutiny by network designers and vendors because it is at this level that enterprise-level network groups still have total control over topology, bandwidth, and cost. Using today's internet and LAN technologies, enterprise networks can be optimized to support the growing use of client/server computing for enterprise-scale applications and prepare for a new generation of high-speed networking technologies that are just now appearing on the market.

The trends already discussed in this white paper present a strong set of arguments for designing office tower and campus network infrastructures using a two-tiered, hierarchical star approach. It is clear that individual floor-level LANs should be implemented using structured wiring and intelligent hubs. Individual connections radiate from hubs in a physical star topology from wiring closets on each floor, or possibly at every third floor (to accommodate floors above and below in addition to those where the hub is located). Hubs must be placed at floor level due to distance limitations imposed by existing Ethernet and Token Ring twisted pair standards — plus those of future higher-speed technologies — that limit individual wire runs to less than 100 meters (330 feet).

Today's Network Centers

Based on the advantages of structured wiring and intelligent hubs, the site-level network design problem becomes a question of how to network wiring closet hubs. The same logic that drove the move toward physical star-wired topologies at the floor level is now driving a move to star-wired topologies at this next higher level. It appears that enterprises will be drawn increasingly toward the creation of site-level or building-level network centers that implement highly reliable internetworking for wiring closet hubs.

At these network centers (see Figure 6), downlinks from wiring closet hubs should be based on fiber optic cable rather than copper wire. This practice handles the potentially large distances between the network center and upper floors in an office tower, or between buildings in a business campus. In addition, the hub system that carries out the network center function has a somewhat different set of requirements than the wiring closet hub, although it is still based on a similar modular architecture.

Figure 6 | **Building-Level or Site-Level Network Center**

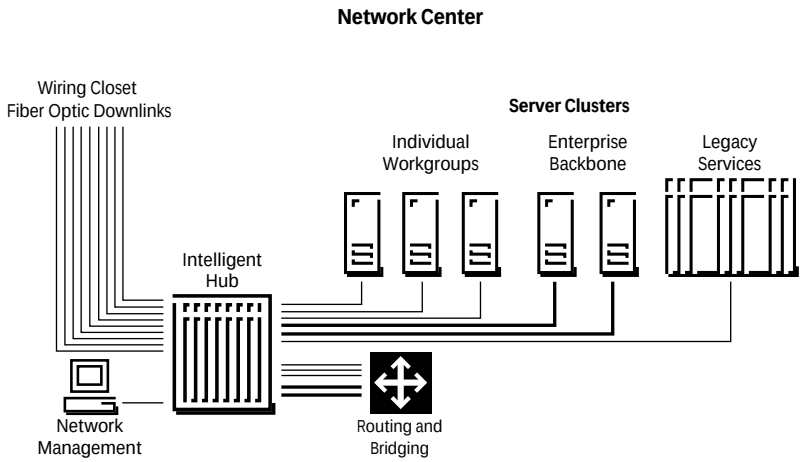

Physical star topologies provide increased manageability, control, and configuration flexibility. By connecting all wiring closet hubs to central points at the building or site levels, it is possible to implement a network in which traffic flows shifting from LAN to LAN are easily handled and in which the number of hops between any two LANs is minimized. Implementing ad hoc floor-to-floor connections creates the same tangled confusion of wiring that originally led enterprise networks to adopt structured wiring and intelligent hubs at the floor level. Figure 7 illustrates how this approach accommodates users who are part of the same workgroup but are distributed throughout a business campus or office tower. Network centers offer structured networking for the enterprise — and bring the benefits of hub architectures to the business campus or office tower network design problem.

This structured network design philosophy aligns perfectly with the trends of server and router centralization mentioned previously. The network center approach actually has the ability to accelerate the move to server centers, as well as improve the enterprise's control and manageability over its server assets.

Network centers will most likely be placed within or near primary site server centers, with intelligent hubs providing server connectivity at various network interface speeds. LAN segments may be established specifically for the needs of enterprise backbone servers. It is typical to have this type of server internetworked on one or more high-speed LANs within the server center. Doing so facilitates access to backup devices, tape libraries, and management systems.

In this structured networking approach, wiring closet hubs use bridges to segment LANs for high-performance workgroups and to provide direct paths between devices on LANs on the same floor. In general, however, LAN internetworking is performed at central network centers using large, multiport routers that also provide access to wide area network connections. This physical star topology allows multiple parallel connections to be run between a wiring closet hub and network center hubs to increase capacity and/or add additional redundancy (see Figure 7).

Figure 7 | **Distributed Workgroups**

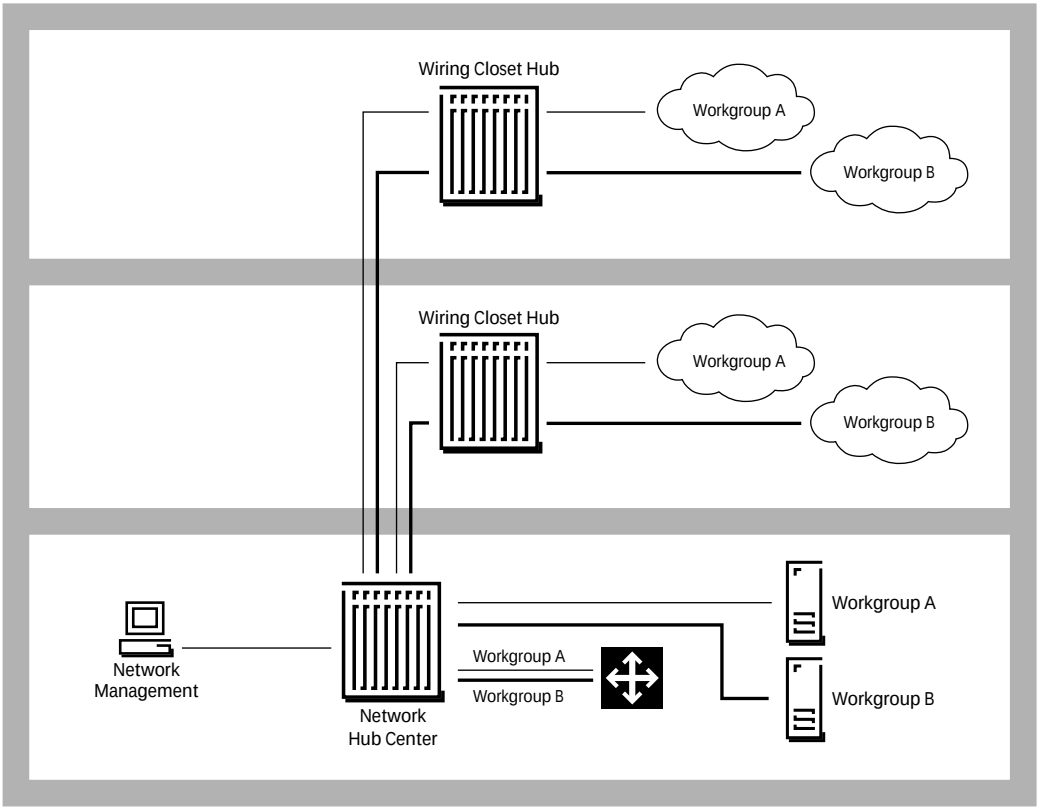

Ideally, each LAN segment at the floor level has its own separate link to the network center. This minimizes backbone contention and allows workgroup servers to be placed in the network center while still residing on the same LAN as the workgroup's floor-level clients. These benefits can be realized today by using multiple individual fiber optic cables. This approach scales easily by multiplexing multiple links onto a single, higher-speed fiber optic cable using available technologies such as ATM cell switching.

Requirements for Network Center Intelligent Hubs

As enterprises reengineer to realize the benefits of client/server networking, new requirements have emerged for networking systems. These include techniques to accommodate centralized workgroup servers, optimize the use of router ports, and provide global management with end-station visibility.

Workgroup Servers

With central IS becoming more involved in deploying these systems, frequently one of the first responses is to relocate workgroup servers into the controlled environment of a server center or network center. These servers provide file sharing for a discrete group of users requiring word processing, scheduling, departmental budgets, and the like.

An important design objective is to reduce the number of internetworking hops between users and their resources. In this case, moving servers to a central location may impose an additional router hop since network designers frequently terminate wiring closet downlinks directly on router ports. The distributed workgroup model is not well-served by this configuration for several reasons. First, if workgroup servers are located at the floor level, users on different floors must traverse a router to access them. Second, installing workgroup servers in the network center traditionally imposes a router hop from the downlink to the LAN segment where the server resides. To rectify this, intelligent hubs for network center use must be able to terminate a large number of fiber-based wiring closet hub downlinks, combine these segments as logical workgroups, and present twisted pair ports to connect the segments to workgroup servers as well as routers and gateways. Doing so eliminates intermediate router hops to workgroup servers and also enables logical localization at central locations.

Enterprise Servers

In addition to accommodating workgroup servers at network centers, enterprise servers — used for cross-workgroup applications such as electronic mail or database access — also require end-user connectivity with a minimal number of hops. Combining, or splicing, segments within an intelligent hub, as previously described, means users on each segment are just one router hop away from FDDI-attached enterprise servers.

This flexibility requires that network center hubs be capable of terminating large numbers of fiber optic downlinks from wiring closets plus have the ability to support many network segments. By contrast, wiring closet hubs are optimized for relatively few segments but many end stations. Thus, it is clear that an intelligent hub's feature requirements are dictated by its specific position within the network.

Network Monitoring and Control

Not only do functional requirements for intelligent hubs differ throughout a network, but management issues also are influenced by position in the network. Integrated network management systems for network center use must provide administrators with the management control required to rearrange workgroup segment relationships with resources, plus splice new servers, routers, or gateways into segments on-the-fly — without rearranging existing cabling. This software-controlled segmentation and assignment of ports to workgroup segments supports the needs of client/server computing with its ever-increasing need for bandwidth.

This new management control paradigm represents a superset of the network monitoring functions found in most intelligent hubs today. Network monitoring is focused on conveying a view of network health to one or more management sites by reporting network traffic levels and building a historic model of network usage. The major purpose is to flag problems as early as possible, to aid in their resolution, and to provide a histogram of usage required for successful network planning. While these tools are critical to network operations, it is the array of management control features that will enable users to keep pace with workgroup and enterprise dynamics in the years ahead. In either case, a critical network management design objective must be to "reel in" the distributed network view — whether a local wiring closet hub, router, server, or even an end-user workstation — to the network manager's desktop (see Figure 8).

Mission-Critical Reliability

Consistent with their position, network center hubs also must incorporate various mission-critical capabilities that make the overall network more fault-resistant, thereby maintaining extremely high availability. This starts with a very high-quality design aimed at alleviating single points of failure within the hub. Such a design includes fully redundant power supply configurations that support the entire hub, not just a part of it, plus the ability to retain configuration data as modules are hot-swapped. Network center hubs also must be able to extend fault resistance throughout the network — for instance, redundant fiber optic links for wiring closet downlinks to protect against cable outages. From a management control standpoint, this means that network administrators must have the ability to set up redundant configurations on an as-needed basis at critical points throughout the network.

Integration of Switched Network Technologies

The growing need for more network capacity and lower latency connections between clients and servers is driving the development of next-generation networking technologies. These new technologies are based on switching, rather than traditional shared-media LANs, because switching architectures become more economical as network speeds increase and they provide a simple way to dedicate capacity to individual interfaces.

There are two distinct types of switching: cell switching and frame switching. Cell-switching techniques, notably Asynchronous Transfer Mode (ATM), are receiving the greatest attention because their capacity is particularly well-suited for multimedia communications at the desktop and, more importantly, because of their applicability in enterprise backbone networks. Using short, fixed-length (53 byte) information cells, ATM also provides the low latency needed for demanding backbone applications where clients and servers are separated by a few hops. Short cells, switched in fast hardware, traverse a switch matrix in a fraction of the time it takes today's large LAN-generated packets to traverse a router.

Another switching technology that will be widely applied in the near term is frame switching. If LAN segmentation is carried to its logical conclusion, the LAN is eventually broken down to one end station per segment — in other words, a dedicated LAN.

Figure 8 **Strategic Points of Network Monitoring and Control**

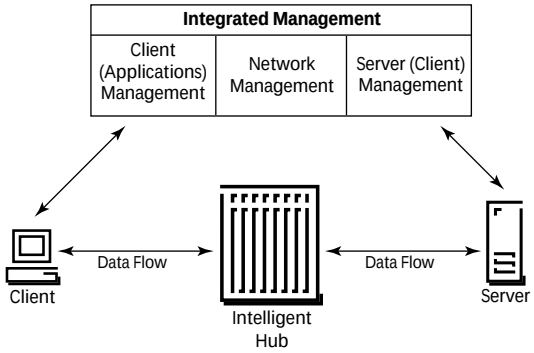

In this form, Ethernet, for example, becomes a point-to-point protocol providing 10 Mbps transmission between the end station and the wiring closet hub, or between the wiring closet hub and a network center hub. Within the hub, a switching engine provides sufficient capacity to handle multiple dedicated Ethernet interfaces simultaneously. Frame switching is a particularly important technology for the desktop interface and for numerous server interfaces (where it enables the provisioning of multiple interfaces). It offers a significant increase in network capacity and investment protection because it does not require any changes to the end stations or to the wiring scheme.

In the near future, LANs containing both cell switching and frame switching will be available to provide high-performance backbone connections to legacy networks. The network center approach to site-level networking perfectly positions the enterprise to take advantage of these new switching solutions. The physical star wiring approach fits well with these emerging switching technologies, and the network center becomes the ideal location for one or more large cell switches that will eventually offer interfaces to public, wide area, and metro area ATM services.

Future-Proofed Network Designs

Future-proofing may be thought of as forward-looking investment protection — making sure today's expense will prove to be a wise investment tomorrow. In designing dynamic, mission-capable networks, this takes the form of ensuring that the intelligent hub infrastructure is extensible enough to incorporate higher-performance technologies as the enterprise's needs evolve. The greatest return on bandwidth investment is found in switch-based technologies: frame switching at the fringes of the network for power users and high-performance servers; cell switching at backbone points-of-presence and for high-capacity workgroup applications. The structured networking approach, discussed earlier, is a key to realizing the benefits of switch-based solutions with a minimal amount of reengineering.

Summary

Server centralization and collapsed backbone implementations are enabling the industry shift from legacy to enterprise, production-level, client/server networks. This change — both in terms of topology and operating environments — establishes a new, mission-critical set of requirements for intelligent hubs designed for use in network centers. While wiring closet hubs are optimized to support many end stations but relatively few segments or rings, network center hubs must be designed to monitor and control large numbers of wiring closet connections as well as many LAN segments or rings. As the centerpoints in enterprise networks, network center hubs must offer a newer, greater level of flexibility and reliability to help solve today's problems while preparing networking systems for the easiest, smoothest integration of emerging technologies.

Bay Networks

For more sales and product information, please call **1-800-8-BAYNET**.

United States

Bay Networks, Inc.
4401 Great America Parkway
Santa Clara, CA 95054
Phone: 1-800-8-BAYNET

Bay Networks, Inc.
8 Federal Street
Billerica, MA 01821-5501
Phone: 1-800-8-BAYNET

Europe, Middle East, and Africa

Bay Networks EMEA, S.A.
Les Cyclades – Immeuble Naxos
25 Allée Pierre Ziller
06560 Valbonne, France
Fax: +33-92-966-996
Phone: +33-92-966-966

Intercontinental

Bay Networks, Inc.
8 Federal Street
Billerica, MA 01821-5501
Fax: 508-670-9323
Phone: 1-800-8-BAYNET

World Wide Web: **http://www.baynetworks.com**

3/96 WP510-913WC-B

Printed on recycled paper.