White Paper

How Much Bandwidth Is Enough?

Network Design in a Complex World

This paper builds on the previous white paper, "Switching Paradigms," examining the new realities of network design and comparing the capabilities of alternative network technologies.

Bay Networks

Northeast Consulting Resources, Inc.
Boston, Massachusetts, USA
by Christopher Serjak

Network Design in a Complex World

Changes in application and networking technology have complicated network design. Network owners recognize they must re-architect their networks in order to meet the demands of distributed applications. New architectures often mean introducing new technologies. Unfortunately, planners have little data to use in making decisions about which technologies to adopt, how to use them, and what will be the impact on network capacity. This paper examines the new realities of network design and compares the capabilities of alternative network technologies.

New Perspective on Networks

Networks exist to move application data—that's the only reason organizations build and operate them. And the only measure of a network's success is how well it meets that goal. Network planners and managers must remember that performance, scalability, and manageability are not simply goals for the network—they are goals for applications that reside on the network.

The corporate network must scale to support a growing user population, larger numbers of applications, and a heavier transaction load per user. The network must deliver high performance in order to meet the demanding response time requirements. Finally, the network must be manageable so that it can quickly be adapted to meet the ever-changing needs of modern businesses.

As organizations become increasingly dependent on information exchange, applications become more network-centric and reliance on the corporate network grows. This view of the business, application, and network challenges conventional wisdom which views the network as "plumbing". A new IT world view asserts that **today's networks are application platforms**—and forces network owners to rethink strategies for building and operating computing systems. Traditional distinctions between hardware and software, application and network are no longer useful. Network-based distributed computing fundamentally changes the network design equation.

Applications Matter

The client/server architecture is network-based computing. That means the people building the network must work with the people creating applications or both will fail. Optimized application designs use networks up to 10 times more efficiently than "vanilla" out-of-the-box configurations. Entities on a data flow diagram (databases, users and other resources) are distributed across a network, and the arrows interconnecting entities represent traffic flows. Planners must keep these transaction streams in mind when designing the network infrastructure.

When planning a network, application developers and network planners should work together on the following issues:

Break organizational structures which isolate developers from the networking organization. Most IT organizations consist of a network operations group and a development group. Typically, the only manager common to both organizations is the CIO. This often results in poor communication between the two groups, which ultimately leads to failed deployments. Network planners and developers should create multidisciplinary teams that work with users to help them understand requirements and set reasonable expectations.

Develop a budget for the network, just as they include a budget for servers, workstations, coding and training. Budgets for downsizing applications from a minicomputer or mainframe must also include the cost of upsizing the network to support the increase in traffic and shorter response times required by distributed applications.

Discover new applications early in the design phase to provide adequate lead times for re-engineering the network if necessary. Planners must identify new application efforts in the formative stages. A Fortune 100 company recently discovered corporate managers planned to deploy nearly 60 new client/server applications in an 18-month period. The network organization knew of only 12. Proper planning eliminates surprises.

The Performance Phase of Network Evolution

The industry has entered a new phase of network evolution: the Performance Phase. Most organizations have successfully navigated earlier phases focused first on building connectivity and then creating basic interoperability among major groups of systems. Having largely been achieved, connectivity and interoperability are now taken for granted. Businesses have taken on a new technical challenge—deploying mission critical network-based applications in a production environment. Production deployment means that the applications (and network infrastructure) must support large numbers of client/server transactions. But getting transactions done—"moving the data"—is enormously difficult on networks. Networks lack the required capacity because they have not been architected for distributed applications. The result is that many IT organizations are focusing on network performance as the driving design goal.

Network Planning is Getting Harder.

Network designers are realizing that the vagaries of real-world applications and network systems conspire to make planning difficult. Centralized resources, new classes of applications, new usage paradigms, changing traffic flows, more connected users and more powerful computing platforms all must be factored into the network design equation. The large number of technology and architecture options can make the design problem appear almost insurmountable.

Too Many Choices?

Figure 1 on page 3 illustrates the menu of technologies available to network planners. The arrival of frame and cell switching charts new territory in terms of network throughput and delay—how much data networks can carry and how quickly it is moved. The LAN/WAN solution space now encompasses six orders of magnitude in network throughput (from one thousand bits per second to one billion bits per second) and four orders of magnitude in network delay (ranging from one second to one millionth of one second). In relative terms, increasing the network solution space from X.25 and fractional T-1 to ATM LAN/WAN switching is comparable to the size difference between a typical single family home and the North American continent.

Evaluate alternative application topologies and their effect on the network (as well as user response times) before anybody starts writing code. Evaluate centralized and distributed application topologies. Investigate the impact of moving servers closer to large user populations and running background updates of the master database.

Build reasonable test beds early in the development cycle. Develop a network prototype which approximates the application traffic flow(s) when developers build a GUI prototype for users. Network prototypes help reveal potential bottlenecks and demonstrate the performance of alternative networks. Let users help determine appropriate response time and they will become allies in defending network performance or justifying infrastructure spending. And don't demonstrate an application on a T-1 circuit as it will run over switched 56 Kbps.

Evaluate alternative protocol suites to determine tuning requirements. Make an effort to understand how optimizing network operations for one (suite of) application(s) may undermine the performance of others.

Identify users and their locations because that information plays a significant role in determining where servers should be deployed. Will users move periodically? Will they be contained on one subnet? Will they connect over the WAN?

Identify application usage cycles and plan for peak periods. Look at weekly, monthly, quarterly and yearly usage trends. Try to anticipate how a "surge in business" will affect the network and determine how much extra bandwidth is required to keep the business running.

Overlay requirements on the existing production network to determine whether it can carry the load. Look at boundaries and bottlenecks. Try to learn how changes in user demographics will affect the network.

Do a trade-off analysis against both the application and the network. Centralizing data often means adding capacity in the center of the network, while distributing data might require augmenting workgroup LANs by installing faster server interfaces to improve response times.

Performance Matters

Although network delay is measured in hundredths or thousandths of seconds, seemingly immeasurable performance differences between different network technologies are important design considerations. A difference of 1/100 second delay can be problematic for users if it occurs one thousand times during the course of an application task. An endless number of these seemingly immeasurable delays are encountered in modern client/server application networks, compounding their effect. In the end, small performance differences between network technologies (bridges, routers, switches, LAN types) result in meaningful differences in application response time and in the number of users a particular network architecture can reliably support.

Figure 1 | **Performance Landscape**

The "Performance Landscape," or range technologies the planner may include in the corporate network. Alternatives offer significantly different levels of response time and total throughput. Technologies such as Ethernet and ATM are available at varying bandwidths, ranging from 10 Mbps to 100 Mbps and beyond (as indicated by the solid bar). Unfortunately, planners have no way of deriving how many application users each technology will support and what response time will be.

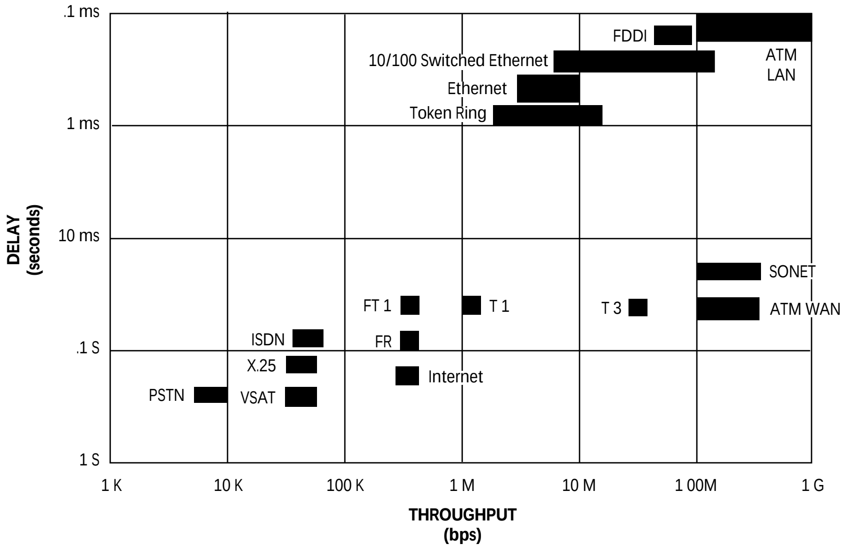

What Kind of Network Will You Build?

The next-generation mission-critical client/server application is already at your door, while others looming just over the horizon promise to consume huge chunks of network capacity. Developers are coding. Users are waiting for better performance and greater flexibility. What kind of network will you build?

Some Guidance for Network Planners.

Hard and fast rules for choosing one network architecture over another are virtually nonexistent. Application requirements are as numerous and diverse as the population of applications in a modern enterprise. If there are any universally true network design rules, they are carefully guarded secrets.

The remainder of this executive briefing is aimed at answering the question "How much bandwidth is enough?" The next three sections provide a side-by-side comparison of four alternative network architectures (router "collapsed backbone"; FDDI; 10/100 Frame switching; and an ATM backbone) running three distinctly different types of applications: classic client/server; imaging groupware; and mathematical modeling.

Each application and network is illustrated on the following pages, with a brief table description to help the reader understand some of the high-level design considerations. A section entitled "How the Networks Stack Up" takes a look at how well alternative architectures meet the needs of each application, which in turn stress different parts of the network—the backbone, server links, and the workgroup/backbone connection.

Three Applications—Which Network?

Figure 2 **The Applications**

A. Transactions include: 1) Download application forms, triggers, rules; 2) Establish remote connections; 3) Query/response to servers; 4) Write statistical analysis.

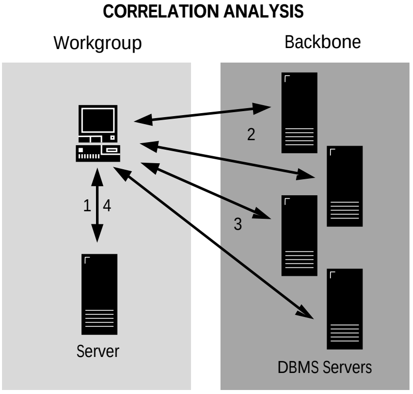

B. Transactions include: 1) View user "in-box" to check for assignments; 2) "Drill-in" to record detail; 3) Update records and files; 4) Check out design documents; 5) Send files to workgroup server for optical conversion; 6) Update design and optical files to server.

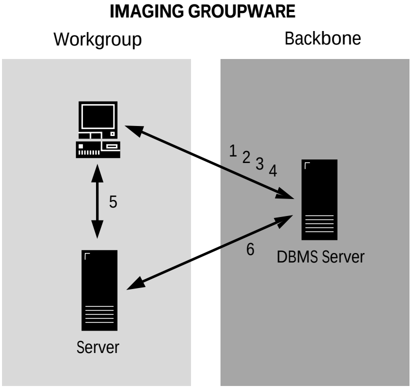

C. Transactions include: 1) Create environmental model, set up simulation parameters; 2) Configure terrain and atmospheric input files; 3) Load modeling input; 4) Load additional files; 5) Exchange working files; 6) Write modeling output files; 7) Write simulation results to workgroup file server; 8) Browse results.

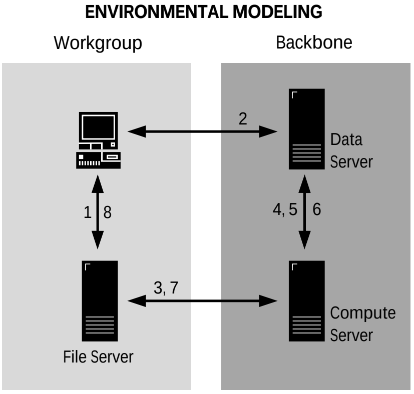

Table 1 **Application Summary**

	Description
Correlation Analysis	**Objective** Replace weekly manual QC data-gathering/analysis with on-line client/server system to provide JIT information for plant operations. Expect >10% reduction in waste and >10% improvement in quality.
	Business Forestry Products Company
	Application Pulp & Paper Manufacturing Quality Control; 20 concurrent users.
	Description Operator/engineer downloads information from servers distributed throughout very large plant. Runs quality/efficiency analysis and creates statistical reports stored on local file server.
	Network Usage Moderate. Majority of client/server traffic flows across workgroup/backbone boundary between LAN client and remote (backbone attached) server. QC information makes network traffic time sensitive. Need flat network architecture to minimize hops to provide acceptable response time. Peak traffic periods during shift changes, switching line to make new products, plant emergencies, long term planning projects.
Imaging Groupware	**Objective** Replace manual document management, storage and retrieval system with automated imaging application. Combined with other work process initiatives to save >20% time to complete projects.
	Business Construction engineering firm
	Application Electronic storage/tracking retrieval of design documents; 300 concurrent users.
	Description Centralized database serves as repository of all design documents. Engineers interact with applications like e-mail/groupware. Engineers query database to determine what design project they should be working on and to update project management information. Download CAD files from image repository; work with files; return updated CAD files to central server.
	Network Usage Moderate to heavy. Typical of groupware applications, this system is in use all day, everyday. Interactive usage and CAD file transfer.
Environmental Modeling	**Objective** Downsize existing application(s) from super computer to client/server system. Expect savings of >$1 million/year in super computer operating costs.
	Business Environmental consulting firm
	Application Atmospheric modeling; 180 concurrent users.
	Description Analyze air emissions from smokestacks, water treatment plants, highways, landfills, etc. Also model toxic waste spills which threaten public.
	Network Usage Heavy. Input file(s) for mathematical modeling as large as 60MB. Heavy server/server transactions in backbone as working files exchanged 1,000 times per model run. Company expects this application/network system to be doing computing 99% of the time, using the network 1% time. Result is high performance expectations. Application in use 24 hours/day; 360+ days/year.

Network Architectures

Figure 3 **Summary of Network Architectures**

The figures at right show two standard network architectures: router collapsed backbone (A), FDDI backbone (B); and two architectural alternatives based on switching: 10/100 Frame Switched Network (C) and ATM & Frame Switching (D).

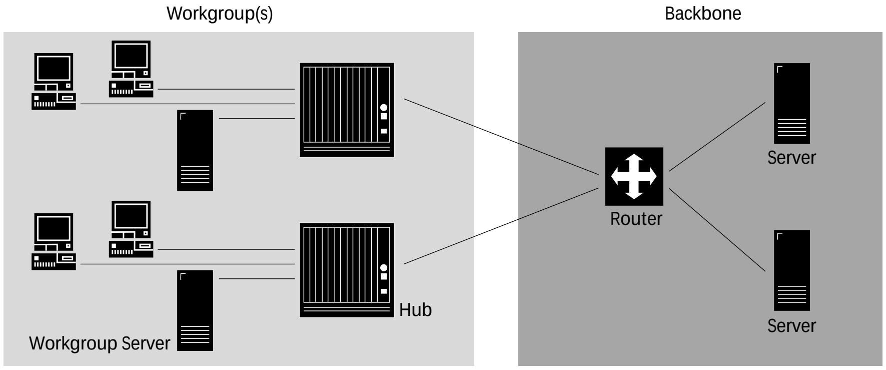

A: Router Collapsed Backbone

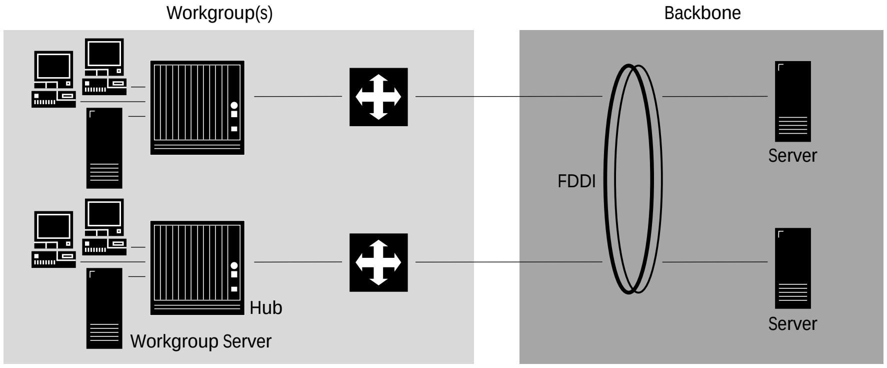

B: FDDI Backbone

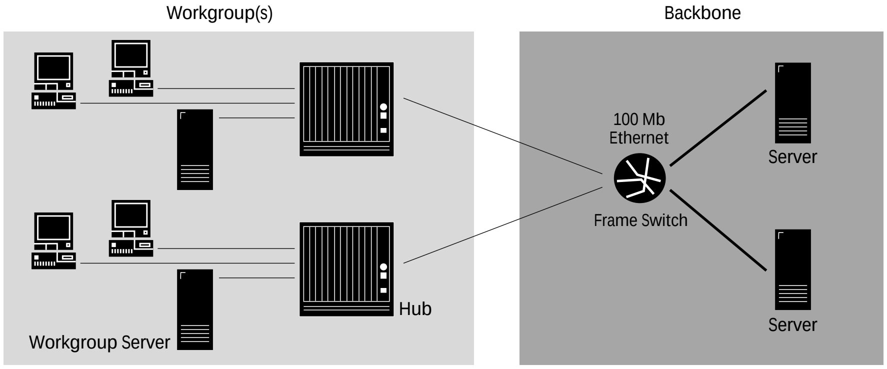

C: 10/100 Frame Switched Network

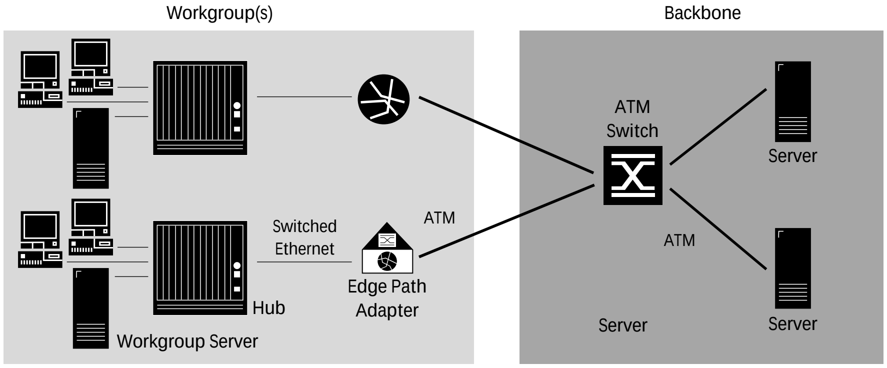

D: ATM & Frame Switching

Table 2 **Summary of Network Architectures**

	Description	Strengths	Weaknesses
Router Collapsed Backbone	Mid-range multiprotocol router interconnecting Ethernet workgroups and segment of centralized server(s). Suitable for building backbone and workgroup interconnection.	Excellent protocol control and separation of workgroups through routing at network center. Little complexity; one internetworking device to manage.	Limited scalability: Must upgrade to higher-performance router in order to get better throughput, or upgrade to FDDI ring. No "backbone" or bandwidth hierarchy: Servers connected via 10 Mbps Ethernet. These connections will become bottlenecks as many users access shared databases.
FDDI Backbone	Single FDDI network interconnects FDDI-attached servers and Ethernet workgroups via mid-range router(s). Suitable for building or small campus backbone.	Excellent protocol control and firewalling through routing at the workgroup/backbone boundary. High-speed backbone handles traffic aggregation and high-speed server/server transactions. Available; mature standard.Good scalability: more workgroups/routers can be connected to FDDI until it is saturated.	Complexity: Scaling network will result in large numbers of routers/router ports, each (potentially) with its own subnet address. Large number of devices and address administration may become burden to manager. Hops between routers and across shared-access Ethernet and FDDI LANs may hurt application performance. Problems are exacerbated if backbone needs to be segmented to carry more traffic.
10/100 Frame Switched Network	A single frame switch provides 12 dedicated 10 Mbps connections to workgroup concentrators (or end stations) and two 100 Mbps interfaces for servers. Suitable for high-performance workgroups or building backbone.	Simplicity eases management burden. Ethernet-only solutions are generally plug-and-play. Virtual LAN capabilities create logical workgroups and support firewalls. High performance and low latency provides good response time to client/server applications with transactions to/from centralized resources.	Primarily an Ethernet-only solution; Token Ring products from most vendors only by late 1995. Most valuable Virtual LAN capabilities will be vendor-specific, requiring users to choose strategic vendor for frame switched products to obtain maximum leverage from VLANs.
ATM & Frame Switching	ATM switch interconnecting ATM-attached servers and Edge Path Adapters which provide dedicated 10 Mbps Ethernet links to workgroups (or end stations) and 155 Mbps ATM connections to a backbone cell switch. Suitable for high-performance workgroups or building/campus backbone(s).	High-speed backbone to handle traffic aggregation and high speed server/server transactions. Excellent scalability: Mesh of frame or cell switches may be easily extended to provide more links to end stations or servers; individual interface rates may be increased as needed using Fast Ethernet or multiple higher-speed ATM links. Virtual LAN capabilities create logical workgroups to ease management. Servers may be centralized for ease of administration, yet remain logically close to users.	ATM is a relatively new technology; standards are still maturing; industry is gaining experience with ATM in production environments. ATM solutions will require close relationships with equipment vendors. A very high-end solution for networks with the top 10 percent of performance requirements.

How The Networks Stack Up

Correlation Analysis: Network Architectures

Figure 4 shows the network response time required to complete a large client/server query/response transaction in the correlation analysis application. Although response time is important, since this application is used in a manufacturing environment, the network only needs to support 20 concurrent users. Response times (on the Y-axis) indicate only network time. To determine the true response time for this query/response, server processing time must be added. For example, given 30 seconds of server processing overhead, this transaction would take roughly 50 seconds in the FDDI environment (19 seconds network time + 30 seconds server processing time) vs. 40 seconds on the ATM network.

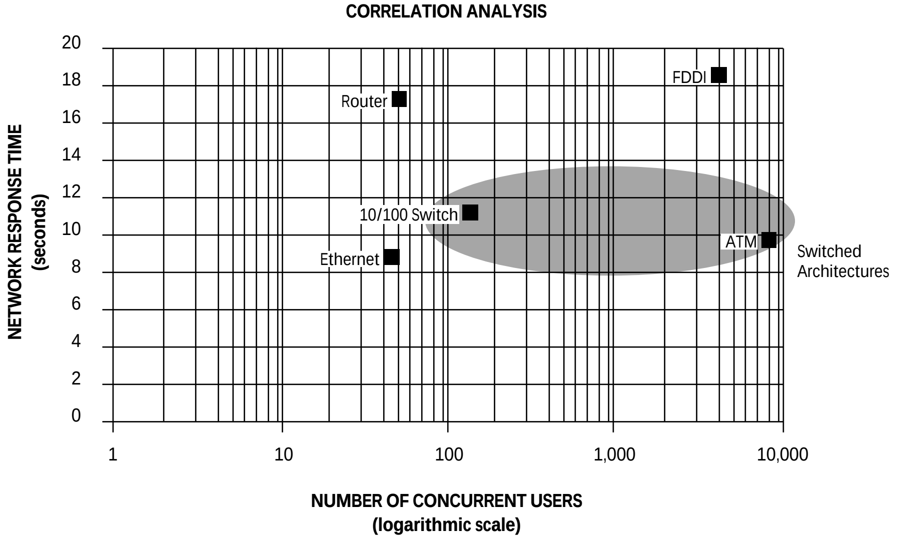

Figure 4 **Correlation Analysis**

The ability of five networks to support the Correlation Analysis application. The Y-axis shows the network time required to complete a large client/server query. The total number of users supported by each network is shown along the X-axis.

The Details: How The Networks Were Modeled

The results graphs in the next section show: 1) the number of users and network response times, and; 2) alternative networks provided in each of the aforementioned applications. Data is based on modeling the performance of alternative applications using NCRI's Traffic Mapping® planning process and is meant to provide a guideline for determining the relative performance between architectures. The performance of actual networks may vary depending on the product implementation(s), application design and tuning, as well as application usage. Several important assumptions underlie the results:

- Response time shows the **network time** required to complete a (group of) select transaction(s); client and server processing time has been subtracted from the equation since the figure shows a **network performance benchmark**.
- Ethernet is used for all workgroups in order to provide a basis of comparison across all architectures.
- Each application is assumed to be running independently on each network. For example, correlation analysis results show that 40+ users can run on a single Ethernet. This is true if **correlation analysis is the only application running on the network**.
- TCP/IP is assumed in all examples.
- The performance characteristics of networking devices, such as delay, latency, aggregate bandwidth, etc., are typical of router and switch products available today.

Ethernet Provides Enough Throughput.

Performance for a single Ethernet workgroup is shown as a baseline. The single Ethernet LAN provides the best response time of all network alternatives, since there is nothing to get in the way of client/server communications except Ethernet's CSMA/CD link layer, which arbitrates usage of the shared Ethernet wire. The single Ethernet supports more than 40 users (with Correlation Analysis the only application running on the network). It is therefore reasonable to expect that the network can support the target population of 20 concurrent users and still have enough bandwidth for other users/applications. Unfortunately, distance limitations alone make it unlikely that all devices in a pulp and paper mill can be connected to a single Ethernet cable. Therefore, a single Ethernet segment is probably not an acceptable solution for this scenario.

Switching Provides Better Response Time.

Response times for switched architectures (both the 10/100 frame switched and ATM) is better than traditional routed or shared-access FDDI architectures. In fact, switching delivers nearly 50 percent better network response times, due largely to the specialized switching engines vendors are developing for these products. Virtually all vendors agree that frame switches will provide better latency/delay characteristics than the vast majority of bridges and routers available on the market and used in corporate networks today. Faster networks translate to better client/server application response times.

The Details: Calculating the Number of Supported Users

The following four-step process is used to calculate the number of users supported on each network:

1) Calculate the number of users (for each application) which may run concurrently in a single workgroup.
2) Determine how much backbone traffic a single workgroup generates.
3) Determine how many workgroups may use backbone resources concurrently.
4) Multiply the number of workgroups on the backbone by the number of users in a workgroup to determine the total user population.

This formula provides a reliable estimate of supported users on each network. This "maximum number of users" is reached when one part of the network is saturated and cannot carry any more traffic. Normally, a network manager would "fix" the bottleneck by adding segments, routers, switches, etc. However, experience shows that alleviating one bottleneck quickly creates another. Therefore, the number of users shown in the results graphs provides a guideline for a "typical" network architecture, not a network that has been "tweaked" to support every last possible user.

Clever network planners might also be able to squeeze a few more users onto the network by tuning it carefully to support the requirements of a specific application. However, this is usually done at the expense of other application users. For example, while tuning a network to support large file transfers may increase the number of users for Imaging Groupware, it will hurt the performance of an OLTP application. In each case, the model is a general purpose network suitable for running several applications.

Bottlenecks Between Router and Server.

The router collapsed backbone supports just over 50 users, not many more than the single Ethernet. The limiting factor is the Ethernet connection between the router and the centrally-located server. Because the Ethernet server connection runs at the same speed as the rest of the network—10 Mbps—that link becomes a bottleneck as users from several workgroups attempt to access the server simultaneously. There is no bandwidth hierarchy in this router collapsed backbone architecture, although proponents of this approach contend that the router's internal backplane is essentially the network backbone. Unfortunately, the fastest router in the world can't alleviate this bottleneck, since traffic can't move any faster than 10 Mbps on a standard Ethernet link. Planners can eliminate the bottleneck by adding a higher-speed connection to the server such as FDDI or 100 Mbps Ethernet, allowing many more users to concurrently run Correlation Analysis.

Switching Eliminates Server Bottlenecks.

Deploying Correlation Analysis on a 10/100 frame switch solves the server bottleneck problem. Adding a private switched 100 Mbps interface to the server allows more users from several workgroups to access the server simultaneously, enabling the network to support 150 users versus the routed network's 50. Furthermore, users gain the response time advantages of switching technology.

Even with twelve Ethernet workgroups accessing the backbone server simultaneously, the dedicated 100 Mbps Ethernet server interface is not saturated. The only factor limiting the number of users this network architecture can support is the number of switched 10 Mbps ports connecting workgroups. Each of the twelve 10 Mbps Ethernet ports is saturated with client requests and responses headed to the server. Since there are no more ports left to connect workgroups, the switch cannot support any additional users. (Of course, more users could be supported if several switches are connected to provide more Ethernet ports.)

FDDI Free of "Port" Limitations.

The FDDI solution, like 10/100 frame switching, provides a high-speed connection for servers—in this case, a shared-access FDDI ring running at 100 Mbps. Unlike the 10/100 frame switch, there are no physical port constraints limiting the number of Ethernet workgroups that can attach to the backbone. The network manager can keep connecting workgroups to the 100 Mbps ring, using routers to translate between Ethernet and FDDI, until the backbone is completely saturated. This solution allows the architecture to scale to support more than 1,300 users.

Unfortunately, FDDI response times are not as good as the switched architecture, and these responses will degrade even further if the network planner segments the backbone once it's finally saturated. That's because segmenting the shared-access FDDI backbone may force traffic to cross two routers during each client/server query and response, adding additional latency.

ATM—More Horsepower Than Required.

The ATM solution provides both excellent response times and supports an extremely large number of users. In fact, the ATM solution supports more than 400 times the target user population, making it a less-than-optimal solution unless other high-performance applications will run on the same network.

In retrospect, the best design options are either 10/100 frame switching or FDDI. Both provide high-speed interfaces and can use fiber optic cabling to connect users over considerable distances. While FDDI is the more "traditional" choice in most manufacturing environments, the 10/100 frame switch solution delivers better performance and is probably a more economical choice, since it doesn't require any additional investments in FDDI interfaces for routers and servers.

Imaging Groupware: Network Architectures

Figure 5 shows the network time required to write an optical image to a client/server groupware database using the Imaging Groupware application. In this scenario, response time will drive the network technology decision for this application.

Typical of groupware, this application is used throughout the day, with engineers constantly downloading and updating CAD documents and interacting with the associated project management database. A slow network or server will degrade response time, discouraging users from maintaining current information. Slow file transfers also cost the company money; a difference of two minutes per transfer, with eight transfers per day for 300 engineers, wastes a total of 80 work hours per day. If an engineer bills at $100 per hour, the company can lose $2 million a year to a slow network.

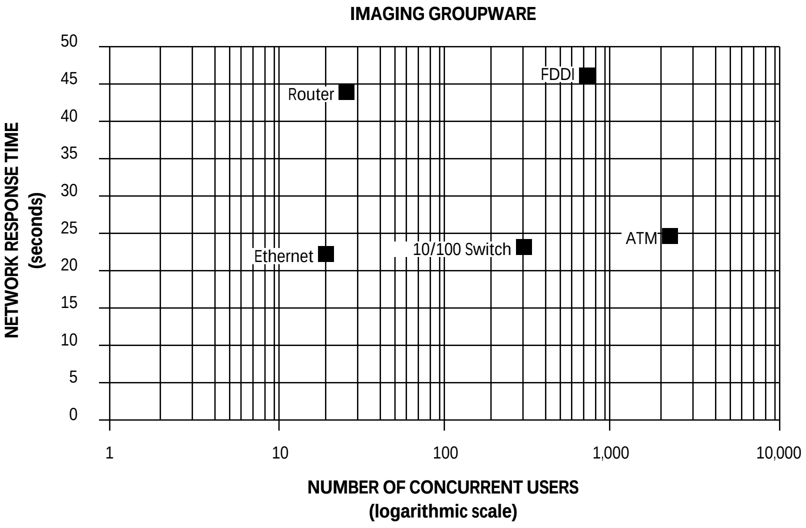

Figure 5 **Imaging Groupware**

The ability of various networks to support the Imaging Groupware application. The Y-axis shows the network time required to write/retrieve an optical image to the database server. The total number of users supported by each network is shown along the X-axis.

Ethernet and Routers Fall Short.

The baseline Ethernet LAN shows excellent response times but, as in the Correlation Analysis example, it cannot connect 300 engineering workstations, servers and associated peripherals. The router-centric architecture can support multiple segments to overcome the physical limitations of a single Ethernet, but it lacks the horsepower required to support the high-volume, time-sensitive transaction stream between workgroups and the "backbone" Ethernet segment. The router's packet forwarding engine runs out of steam before the interfaces are saturated; the mid-range routers deployed over the last several years are not designed to support production client/server traffic streams under bandwidth-intensive applications. Neither a single Ethernet segment nor the router-centric network can support enough users to be considered an optimum design for this application.

Capitalizing On Switching's Response Time.

The 10/100 frame-switched network provides a response time that rivals the single Ethernet architecture while supporting the target 300 concurrent users. Unfortunately, at that size, there is no headroom left for other applications. Just as with the Correlation Analysis application, the number of ports is the culprit that limits how many users may be connected. Unlike the router, however, the 10/100 frame switch solution has been designed to sustain wirespeed frame forwarding across all interfaces, ensuring the switching engine itself will not become a bottleneck.

A New Topology: "Clustering"

The network planner who wishes to capitalize on the superior response time of 10/100 frame switching technology (over both routing and FDDI) might choose to divide the user population into networked clusters of users, each with its own "centralized server." (See Figure 6). While clustering is a popular technique with frame-switched architectures, the process demands cooperation between developers and network planners. To create clusters, developers distribute an application's database across three or four servers while network planners connect users and servers to a high-speed switch. Clustering forces developers and the networking staff to work out any traffic flows and database replication schemes to keep application data accurate and traffic flows manageable.

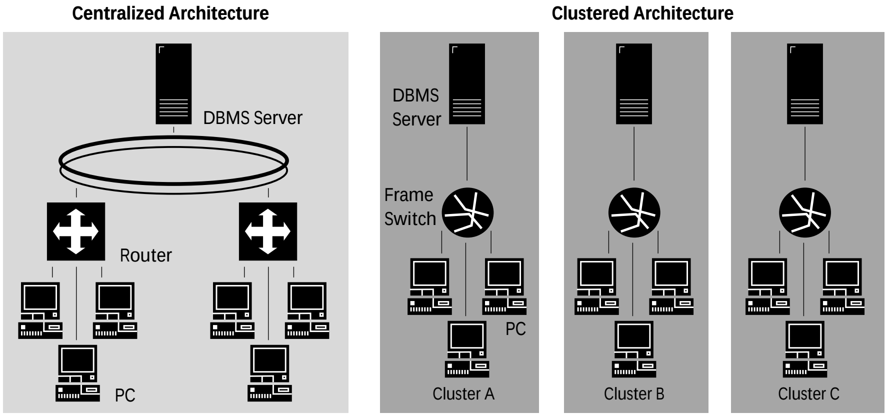

Figure 6 **Clustering**

The "clustering" concept in which a centralized database is distributed across several smaller servers located near concentrations of users. Putting servers close to users improves response time and often reduces traffic across a large enterprise backbone. This architecture forces planners to design networks and applications together.

FDDI and ATM Support the Load.

FDDI and ATM both fully support the user population. However, while FDDI supports roughly 700 users, response time is comparatively poor; the FDDI network requires 20 more seconds than a 10/100 frame switch or an ATM switch to complete this same file transfer operation. To this company, those 20 seconds can translate into losses of $327,000 annually—lost engineering time related directly to the network's slow performance for this application task. By comparison, ATM's superior response time might help the technology pay for itself within a year. The question is whether to invest in an ATM network that supports more than 2,000 Imaging Groupware users.

Environmental Modeling: Network Architectures

Figure 7 shows the relative performance of the five network architectures and their ability to support the Environmental Modeling application. Coordinates along the y-axis indicate the time required to exchange a set of working files 1,000 times between two centralized (backbone) servers. This file exchange occurs during every mathematical modeling. A slow network causes servers to suspend processing while the file exchange takes place. The network owner requires that this client/server application system perform computing (modeling) 99 percent of the time and move data less than one percent of the time. Very high networking and computing performance is required for this business to succeed in downsizing this application from a super computer to the client/server environment.

Ethernet LAN and Router Are "No Shows."

The server/server transaction load on an Ethernet backbone—even if it is firewalled from workgroups with a router—swamps the network with only three Environmental Modeling users. The result coordinates (network time vs. number of users) for the Ethernet LAN and the router-centric network are identical in the graph above. Both architectures use Ethernet backbones and, as a result, support the same number of users. Both architectures also have the same network response time. The server/server transaction above does not require a router hop, even in the router collapsed backbone architecture, since both servers are on the same shared Ethernet segment.

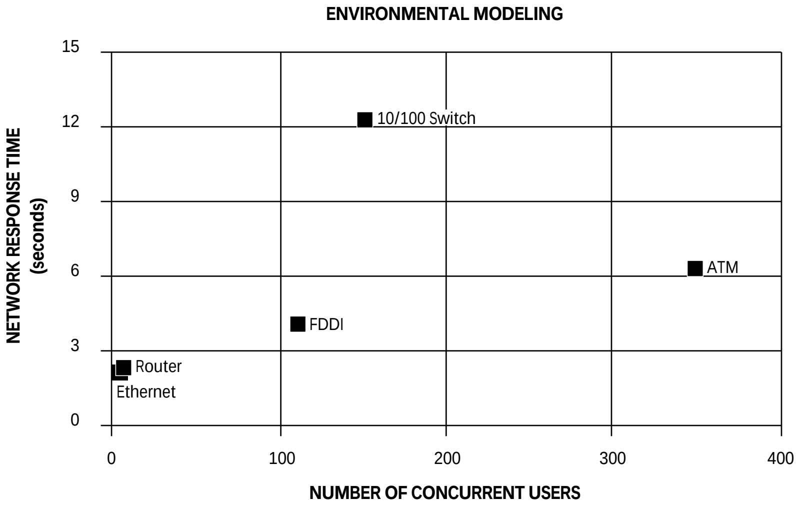

Figure 7 **Environmental Modeling**

The ability of various networks to support the Environmental Modeling application. The Y-axis shows the network time required to support 1,000 server/server backbone transactions. The total number of users supported by each network is shown along the X-axis.

FDDI Falls Short.

The FDDI network supports roughly 110 users, considerably short of the 180 the business requires. A clustering approach, similar to the one discussed for the Imaging Groupware application and 10/100 frame switching, might work if the network owner wishes to deploy an FDDI solution. FDDI clusters, however, will not be as efficient as 10/100 frame switching clusters, since each cluster would require one or more standalone routers interconnecting Ethernet workgroups to FDDI rings.

The FDDI network's response time for this server/server transaction is excellent. The transaction takes place over a single FDDI ring, so no router hops are required. As a result, the FDDI ring provides a better response time than both the 10/100 frame switch, which processes traffic across its 100 Mbps interfaces, and the ATM solution, which must also switch traffic between two ATM interfaces.

Although FDDI shines for this backbone-only transaction, the technology provides much worse workgroup-to-backbone transaction response times than 10/100 frame switching or ATM. In fact, for some transactions, FDDI is up to four times slower than 10/100 frame switching or ATM.

10/100 Frame Switching Better Than FDDI.

The frame switching solution supports more than 150 users compared to FDDI's 110. The reason: while FDDI offers only a single shared-access ring, the 10/100 frame switch offers a private switched 100 Mbps interface to each of the Environmental Modeling backbone servers. This allows the switch to move more data from server to server and from workgroup to backbone server. Not many network planners would suggest a user buy a 10/100 frame switch over FDDI, yet switching is a better solution for this application because of the way it uses the network. Still, frame switching does not support the required 180 users. That leaves only two options: either create clusters of 10/100 frame switches, or move to ATM.

ATM Does the Job.

Only the ATM architecture delivers the horsepower required to meet the demands of Environmental Modeling. As a switched network, it provides excellent response times for both transactions between the workgroup and the backbone and transactions wholly contained within the backbone.

New Technology, New Architecture = More Planning

The design cases described in this paper highlight a number of important issues and shed some light on several critical application/network design decisions. Where a network planner may have been successful drawing upon intuition and experience in the past, they may find it difficult in the future. Network design decisions cannot be made in a vacuum—without careful consideration of the applications the network serves. The "art" of network design must evolve into a science if distributed computing is to be successful. The danger is a potentially mismatched pairing of application and network that results in failed implementation.

The fact that 10/100 frame switched Ethernet supports more Environmental Modeling users than an FDDI backbone network comes as a surprise to most planners. For this application, "intuition" would have probably led to an FDDI deployment—a more costly decision that would ultimately support fewer users. The same is true for the Imaging Groupware application, where a frame-switched or ATM network might save the company $300,000 annually over an FDDI-based network by providing better response times and making the engineering team more productive.

Today's network design equation contains too many variables for planners to trust their intuition. There are few "obvious" answers to complex application/network design problems, and the stakes are too high. Instead, cooperation, understanding and careful planning by developers and network planners alike are the keys to success.

Journey To High Performance: Old vs. New Roads

The road to higher network performance didn't have many turns or intersections until the introduction of frame and cell switching technology. Ethernet and Token Ring segments were connected to routers. Token Ring LANs that needed more bandwidth could move up to 16 Mbps rings. After that, network planners deployed FDDI backbones to aggregate traffic between workgroups.

Today, Ethernet switches with 10 Mbps ports are complementing or displacing routers in certain environments, while 10/100 frame switches are challenging FDDI closer to the center of the network. As the examples in this paper show, frame switching, on average, provides significantly better response times and supports a larger user population when compared with both router-centric networks and FDDI backbones. And switches are being deployed in new "cluster" configurations which provide dedicated connections to users and high-speed links to servers or the enterprise backbone. These switching characteristics make for more scalable, more manageable networks.

Networks are now being re-architected so that routers no longer lay in the path between client and server, since most have not been designed to support low-delay, high-performance client/server transaction streams. Routers are resuming their original role of interconnecting dissimilar networks (Ethernet to Token Ring, for example) and providing firewalls between large, independent networks.

Although FDDI is still the mainstay for large backbone networks, ATM cell switching has begun to displace FDDI as the backbone technology of choice in organizations pushing the performance envelope with leading-edge application technology. This technology promises to become mainstream by the end of the decade, and this trend will be accelerated by the technology's emergence in the public network.

Beyond mere performance advantages, switching promises to improve network manageability through virtualizing network resources, allowing managers to create logical collections of users and network resources. Logical groupings are easier to maintain and adapt much better to changes than the physical subnets defined by router placement.

Frame switching also leverages existing network technology. In cases where 10/100 frame switching is comparable to an FDDI backbone, the technology offers the additional advantage of connecting directly to any device with an Ethernet or Token Ring interface. By contrast, building an FDDI backbone requires an investment in routers and FDDI interfaces. In the long run, a switching implementation will cost less.

Finally, frame and cell switching define a new price/performance frontier that will reduce the cost of networking over the long term. Network owners often fall into the trap of evaluating the cost of networking in terms of price per port, a popular measure of network cost. Although an accurate yardstick back when connectivity and interoperability were networking's primary goals, price-per-port is no longer a valid equation for determining a network's true cost.

Today, the goal of networking has shifted from building connectivity to moving large amounts of data for operational distributed applications. With this goal in mind, a true measure of the cost of networking must reflect the ability of a network to move data. Here, cost-per-port is irrelevant, since it is almost impossible to determine how much bandwidth a shared access LAN can provide. Instead, a better metric is the "cost per usable megabit," which measures how much bandwidth a client or server is guaranteed. With switching technology, each end station receives dedicated bandwidth—network capacity which can be used to complete data transactions. On the basis of cost per megabit, switching becomes far more economical than traditional shared access LANs.

The Bottom Line

Networking vendors and owners have learned a lot about data communications over the last 10 years. This experience has helped them take a dramatic step toward solving their most troubling problems by creating a new generation of products based on high-speed switching technologies. Switching delivers high performance, provides excellent response time, and allows planners to design networks for manageability—three features that are critical for networks of the future.

Agenda for Change

The challenge for network planners is to continue solving problems more effectively. This challenge manifests itself when the network becomes congested or users demand better response time. Forward-thinking planners see a congested Ethernet interface on a router; recognize the problem as an application traffic flow; then consider several solutions to fix the problem. Planners who practice the art of intuitive network design will feel compelled to apply traditional techniques such as segmentation, firewalling, or the adoption of faster shared access LAN technologies such as FDDI. These conventional techniques, however, have already begun to yield diminishing returns; they cannot possibly serve as the foundation for a network that is expected to carry the business into the next decade. Instead, planners must begin considering new technologies that deliver significant advances in both performance and manageability. The rise of switching technology signals a major intersection on the road to better networks. The industry, as a whole, has endorsed the technology with an investment and development cycle unrivaled in the history of networking. Any sound network architecture must include provisions for using this important technology—whether on an experimental, pilot or permanent basis—in order to derive maximum benefit from new application initiatives.

Looking Ahead

The final installment of this executive briefing series—"Getting To Switched Networks"—will examine the issues surrounding the adoption of frame and cell switching throughout an enterprise network. Strategies, standards, topologies, technologies and industry "religious wars" will be addressed and, in some cases, debunked. The briefing will also include a test to decide whether you should consider ATM, plus a questionnaire which determines if your installed technology base, plus a comprehensive three year plan, add up to a next-generation architecture.

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**

Prepared for Bay Networks Inc., by Christopher Serjak of Northeast Consulting Resources, Inc.

8/95 WP510-1619WC-B

Printed on recycled paper.