Enterprise Application Management Made Easy

Technology and Trends For the Next Decade

CITRIX

Executive Summary

MIS managers face the daunting task of deploying client/server Windows applications across enterprise networks, which can easily grow to regional, national, or global proportions. Unfortunately, traditional client/server technologies rarely rise to the enterprise-wide challenges faced by MIS. In fact, the established approaches usually hinder strategic application deployments by inflating costs, complicating management, and performing poorly.

Traditional client/server application architectures and the accompanying deployment models established by distributed PC-based LANs, remote control, and remote node technologies all fail to deliver fast, cheap, efficient application deployments. The problem is inherent to traditional client/server architecture, which emphasizes client-side computational power. In today's widely distributed enterprises, the client/server model breaks down as the client moves farther away from the server yet is required to perform the same tasks as a local machine.

The solution to the enterprise-wide application deployment challenge lies beyond the traditional client/server model. New *thin-client/server* technologies are creating a new, network-centric computing paradigm. This evolutionary network architecture, which employs multi-user application servers, centralizes application deployment and puts a thin-client at the end-user workstation. The resulting model reduces costs, increases manageability, and improves application performance.

Challenges of Enterprise-wide Deployment

At most large companies, MIS staff involved in Windows application deployments must take into account the enterprise network's diversity. On the client workstation side alone, MIS managers will likely find Pentium PCs running Windows NT or Windows 95 sharing network resources with 286 PCs running DOS as well as Macintosh, OS/2, and Unix clients.

Remote network connections will also run a gamut. Wide-area links to remote LANs may be established via T1 lines or they may be made with little more than ISDN or 56K frame relay links. Mobile users and telecommuters are likely using asynchronous dial-up connections running at 28.8Kbps or 14.4Kbps.

Finally, MIS managers must face the realities of the user audience and the application itself. In addition to the users at corporate headquarters, the application may be distributed to hundreds or thousands of users nationwide or worldwide. Just as the user

audience expands, the application itself will likely grow as it is upgraded to reflect changes in the business it addresses.

Using traditional client/server technologies to deploy a client/server Windows application in such an enterprise network raises huge cost, management, and performance challenges for MIS.

Challenge One: Cost

The bottom-line cost for deploying an application includes hardware, software, and personnel costs involved with the application's initial rollout as well as its on-going maintenance.

Start with the application's initial rollout. A server must be installed at each remote site that will run the application. In addition to the expense of the required servers, client workstations must often be upgraded with larger hard disk drives or more random access memory to run the new program. In the event that the application requires a processor or operating system upgrade, the rollout costs can grow to include completely new workstations. Finally, an MIS team must either reside at, or travel to, each site to perform any needed hardware installation and upgrades as well as load and configure the application's server and client components.

If the enterprise-wide infrastructure is not in place, the deployment costs will include the purchase, installation and configuration of wide area access devices such as routers, bridges, and modems for all sites that will be part of the enterprise network. Wide area access services must also be procured. Depending on the demands of the application, WAN service might be an asynchronous dial-up connection that costs the same as a long-distance phone call. It might also be a dedicated line running at T1 speeds or higher that costs thousands of dollars a month.

Challenge Two: Management

While hardware and software costs can most often be readily identified, those associated with MIS management of the application deployment are less clear. On the client-side alone, industry analyst firm The Gartner Group estimates that the total cost of PC ownership can surpass $6,000 per year per PC. This figure includes the cost of hardware upgrades required to support increasingly complex applications as well as the requisite MIS staffing expenses associated with the desktop and application maintenance.

In addition to client-side maintenance, the management of an enterprise application covers roll out, version control, scalability, security, and flexibility issues on both the server and the desktop.

The costs noted earlier for an application rollout often prevent MIS from rolling out an application simultaneously to all sites and users enterprise-wide. Instead, MIS must often stagger the rollout to a handful of sites at a time. Depending upon the size of the enterprise and the complexity of the application, that approach means that years could separate the first and the last site that are brought online.

Once the application is in users' hands, the delay encountered during initial rollout is encountered again when the application is upgraded or patched. Here, the problem becomes version control -- more than one version of the application is in use at any one time. At their least damaging, version control problems prevent users running earlier versions of the application from accessing features or data in later versions. At their most damaging, version control problems can crash applications with mismatched client and server code and/or corrupt application data.

Of course, version control problems grow as the application scales, but scalability itself presents other issues for a lean, downsized MIS staff that must put an increasingly popular application in the hands of more users. On the server-side, larger hard disk drives, more random access memory, or faster processors may need to be installed. New users will need the same type of attention that was given to the application's original user base -- potential hardware or software upgrades as well as installation and configuration of the application's client component.

As the application scales to hundreds or thousands of users, security concerns scale with it. Mission-critical applications requires several layers of security, including multiple password and privilege levels, login and data encryption, and for mobile users, roving callbacks. Lack of acceptable security mechanisms within the application may necessitate add-on products. Ideally, application security will be tightly integrated with overall network security to reduce the duplication of security management and administration for MIS.

To an MIS manager, an application deployment model must be flexible enough to work easily within the existing enterprise network. On the desktop, this means deploying an application to a user population that includes 286, 386, 486, Pentium, and Pentium Pro PCs running DOS, Windows, Windows for Workgroups, Windows 95, Windows NT, OS/2, or Solaris. Macintoshes and Unix workstations running AIX, HP/UX, or any other Unix flavor may also reside on the network. At the network level, an application may need to run on NetWare, Windows NT, an intranet, or the Internet while working with any one of several management platforms.

Challenge Three: Performance

Establishing acceptable performance is paramount to an application deployment. If the application is mission critical and must be used, poor performance will frustrate users with slow response times. In this case, users are unlikely to take full advantage of the

application's features because users want to spend as little time working with it as possible. If users have a viable alternative to the slow application, the application will soon be ignored. Overall, business inevitably suffers from unacceptably slow response times.

Application performance is usually not a problem on the local area network. Today, connections between clients and servers average 10Mbps, and 100Mbps connections are becoming increasingly common. Traditional client/server applications, in turn, are designed and optimized to run over these multi-Mbps local links. The network can support the traffic the application generates.

When MIS deploys the same application to the enterprise at large, however, performance becomes problematic. In contrast to local speeds, wide area connections between networks average 56Kbps to 1.5Mbps while the dial-up connections of remote users average 28.8Kbps or less. Application response times are easily throttled by the meager connections that can degenerate remote performance by several orders of magnitude when compared to local performance.

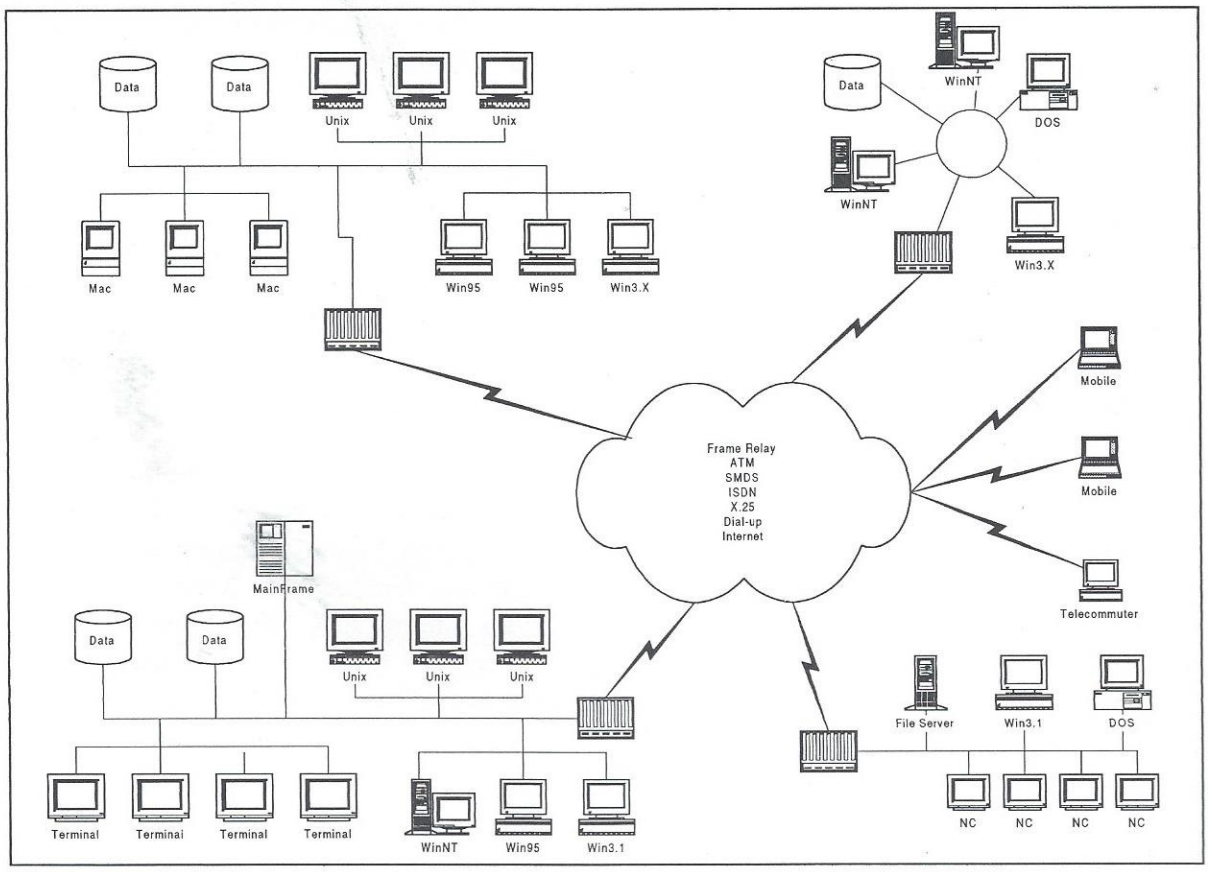

MIS rarely has the luxury of deploying mission-critical applications in a homogeneous environment, let alone from a centralized location. Instead, the enterprise network usually includes a widely-dispersed variety of servers, client workstations and operating systems. A variety of wide area connections connects remote office LANs throughout the nation or the world. The user base can include from dozens to thousands of local, remote, mobile, telecommuting users.

Client/Server Complications -- Application Architectures

A number of client/server application architectures have emerged to meet the growing demands of large-scale, enterprise-wide deployment. However, when evaluated by the challenges -- cost, management, performance -- these architectures regularly fail to meet MIS's needs.

To analyze these different application architectures, the components of a program must be divided into three basic functions. The data access function stores and retrieves data on the file server or other secondary storage device. The transaction processing component, what is commonly referred to as the business logic, deals with the computational aspects of an application. The presentation services function handles the user interface information such as keyboard and mouse input as well as the data presentation on the user's screen.

Shared file applications

Many legacy, LAN-based applications have a shared file architecture in which the entire application is loaded and executed on the user's PC. Other than opening and closing files and records in a shared manner, the network is transparent, and the application design is identical for running both on a standalone PC or a PC attached to a LAN.

This application architecture can require extensive client-side investments in hardware and maintenance. Network bandwidth is also taxed, and performance over dial-up or WAN connections can suffer.

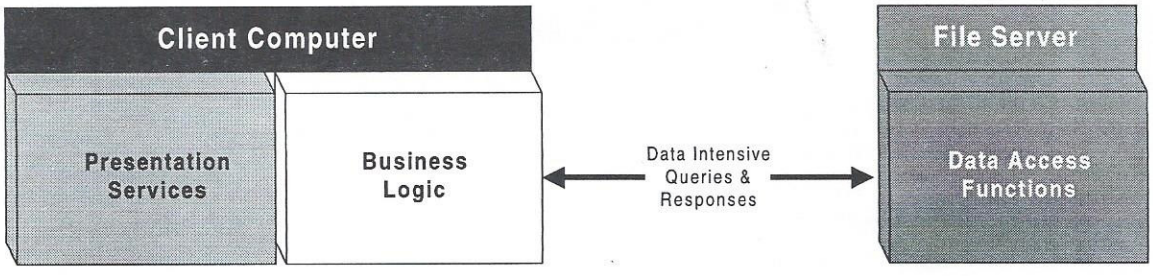

Two-tier client/server applications

The predominant architecture today, two-tier client/server divides the applications into two parts. The presentation services and the business logic functions execute at the client PC, while data access functions are handled by a database server on the network. Two-tier client/server applications are not optimized for dial-up or WAN connections, so response

times are often unacceptable for remote users. Application upgrades require software and potential hardware upgrades to all client PCs, resulting in potential version control problems.

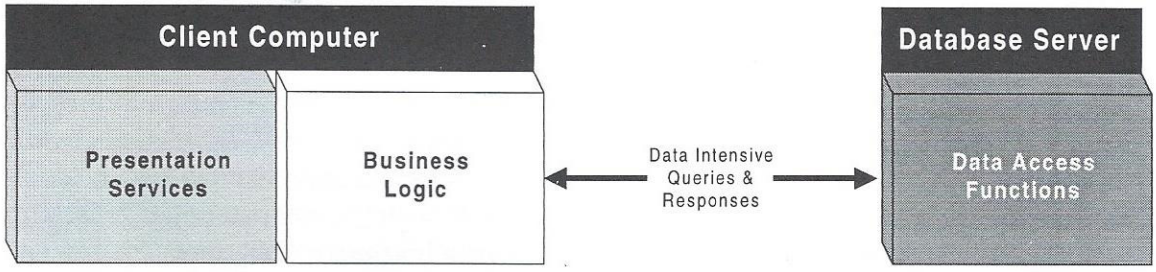

N-tier client/server applications

N-tier client/server architecture requires the restructuring of an application into three or more parts. Presentation services functions execute on the client PC, which may also perform a portion of the business logic functions. The rest of the business logic is executed by one or more intermediate servers, or application servers. Data access is handled by one or more backend database servers. The N-tier architecture can be optimized for dial-up and WAN connections, since clients may send and receive only low-bandwidth screen updates. N-tier applications, however, require complex application design and high-level application development skills.

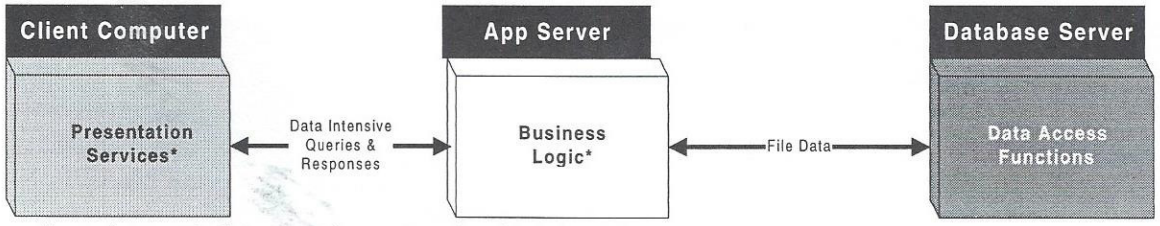

Complications Continued -- WAN/LAN Architectures

In addition to the various application architectures, MIS has several traditional WAN/LAN architectures at its disposal. These architectures have been designed primarily to accommodate the shared file and two-tier client/server applications that make up the bulk of enterprise network applications. By leaving much of the application processing to the client, these models often fail to deliver a satisfactory deployment solution for MIS when graded on cost, management, and performance standards.

Distributed LANs

Distributed LANs create a deployment model that puts a copy of the enterprise application on the local area network of each remote site. Remote sites, in turn, typically connect to a central LAN at the organization's headquarters via WAN connection such as frame relay, SMDS, or ATM running at fractional T1 speeds or faster. Applications data from the remote sites is then sent back to the central site at daily, weekly, or other regular intervals. The data from all remote sites is then consolidated manually or automatically to produce enterprise-wide results.

Cost assessment: Unsatisfactory. MIS staff is needed to deploy the application's server and client components at each site. Server hardware and maintenance is required at each site. Client-side hardware upgrades may be required. Distributed clients require on-going maintenance by MIS at each site.

Management assessment: Unsatisfactory. Rollout delays and version control issues may arise, depending on the size and availability of local MIS staff relative to the size of the enterprise network. Scalability relies on additional network installations, increasing demands on MIS. Tight integration between application and network security is not assured. Flexibility depends on individual application's support of network services and client and server platforms.

Performance assessment: Excellent. Users access applications via fast, local connections.

Remote Control

Remote control software enhances LAN deployment with a means of granting application access to mobile users, telecommuters, and other remote users not directly connected to the LAN. Remote control users may access one of the many individual LANs that comprise the enterprise, or they may all access a single LAN that acts as a central application hub. In either scenario, the remote control software lets users dial in and take control of a dedicated PC located on the LAN.

Remote control software consists of a host and a remote client program. The remote client program runs on the remote PC and is used to dial into the network and control the local PC. The host program runs locally on a dedicated PC attached to the LAN and is used to perform all application processing, essentially turning the remote PC into a long-distance dumb terminal. Applications and data are stored on the host PC, rather than being downloaded to the remote PC. The only data transmitted between the remote and host PCs are the screen updates of the applications and the keystrokes sent by the remote user.

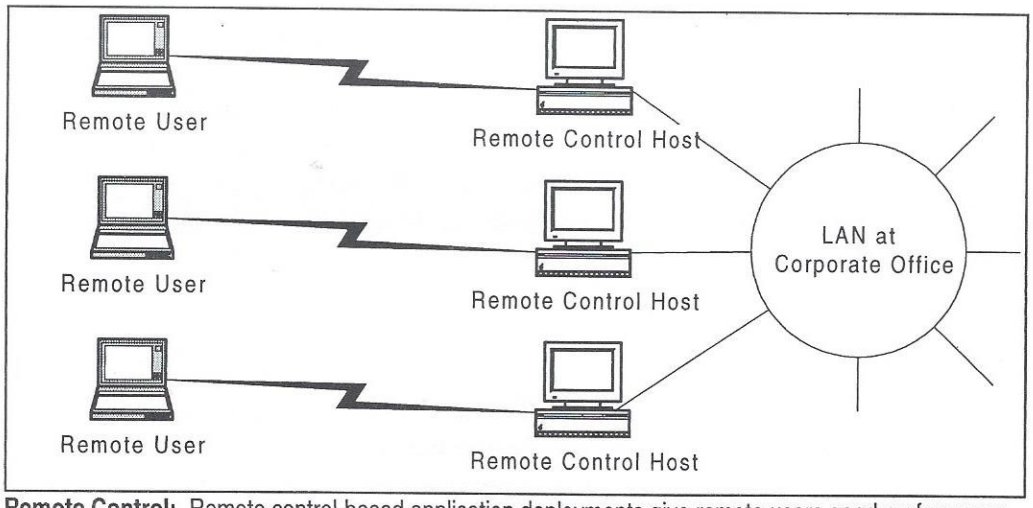

Remote Control: Remote control-based application deployments give remote users good performance since only screen updates are passed between the remote user's workstation and the local host computer. Unfortunately, the additional layer of hardware increases deployment costs and hinders scalability: Each new, concurrent remote user needs a local computer to control. MIS winds up managing local host computers in addition to local servers and remote users.

Cost assessment: Unsatisfactory. A dedicated, local PC is required for every dial-in line supported. MIS management duties grow to include an additional hardware layer -- local host PCs -- on top of remote user PCs and servers. Upgrades for remote user PCs depend on support for remote client program. Upgrades for local host PCs depend on application requirements.

Management assessment: Unsatisfactory. Scalability is hampered by the need to add and manage a dedicated, local PC to support each new, concurrent remote control user. Security and flexibility are application-dependent as in the distributed LAN model. Rollout delays are mitigated by consolidating local host PCs. Application upgrades, however, may require upgrades to local host PCs and introduce version control problems.

Performance assessment: Good. Remote users send and receive low-bandwidth screen updates to boost application performance.

Remote Node

Remote node software, too, gives remote users access to LAN-based applications. Like

their remote control cousins, remote node users may either access applications that have been distributed to the company's many LANs, or all users may access a single LAN.

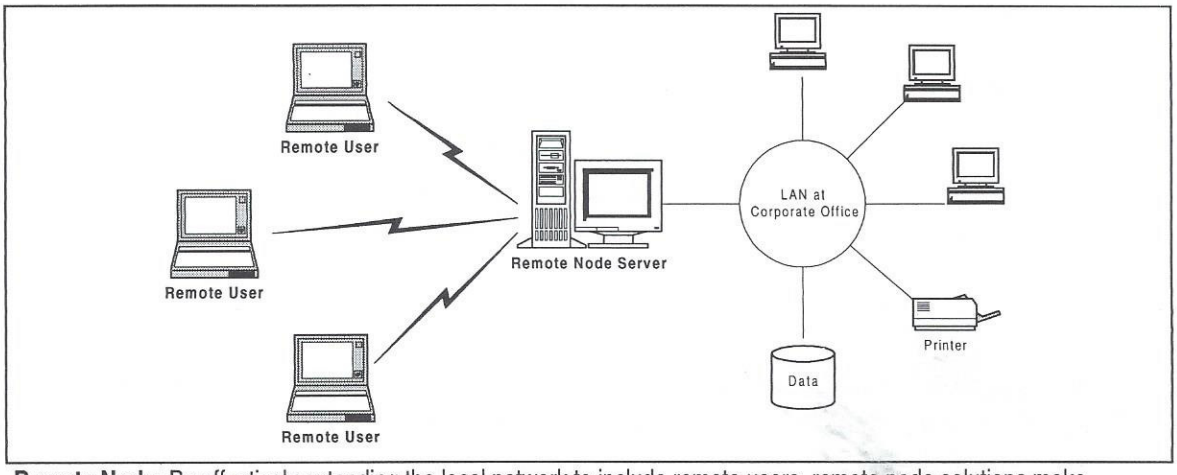

Remote Node: By effectively extending the local network to include remote users, remote node solutions make enterprise-wide application deployments easier to manage and scale. Remote clients are logically identical to their local counterparts, and applications run on local and remote machines in the same manner. Application performance, on the other hand, is stifled by comparatively low-bandwidth WAN and dial-up connections.

Remote node solutions provide access to the LAN by extending the entire underlying network structure out to the remote PC. Once connected, the remote PC operates as if it were directly connected to the LAN. All the application processing occurs on the remote workstation. Logically, a remote node is no different from a local LAN node. Instead of keystrokes and screen updates, the traffic on the remote node's WAN or dial-up line is normal network traffic.

Cost assessment: Unsatisfactory. MIS staff is needed to deploy the remote node software and application's client component for each remote user. Client-side hardware upgrades may be required. Wide distribution of remote clients may exacerbate on-going maintenance challenge.

Management assessment: Fair. Large numbers of remote clients may introduce rollout delays and version control problems. Scalability is acceptable once remote node software and application client software is distributed.

Performance assessment: Unsatisfactory. The transmission of standard network traffic across dial-up and WAN links can slow application response time to unacceptable levels.

Emerging Deployment Options

The fundamental problem with the traditional application architectures and deployment models -- with the exception of certain N-tier applications -- is that they retain the fundamental model established by traditional client/server computing: The user's workstation executes some, if not all, of the application's business logic. In the case of the remote control model, a workstation "surrogate" runs business logic for the remote user. Consequently, these traditional approaches thwart deployment of enterprise-wide applications from several cost, management, and performance points.

To establish a more effective application deployment model, emerging technologies challenge the assumption made by traditional client/server models that client workstations must execute application business logic.

Thin-Client/Server Architecture

Thin-client/server architecture takes an evolutionary approach in solving the application deployment problem. The thin-client/server approach moves all of the business logic off of the client and consolidates it on a multi-user application server.

Multi-user application servers comprise of two parts. The server component extends the operating system to make it a multi-user application server that runs all business logic. Data access takes place on one or more backend database servers.

The client component is a small-footprint program or device that executes only the application's presentation services and is optimized to run over low bandwidth dial-up connections.

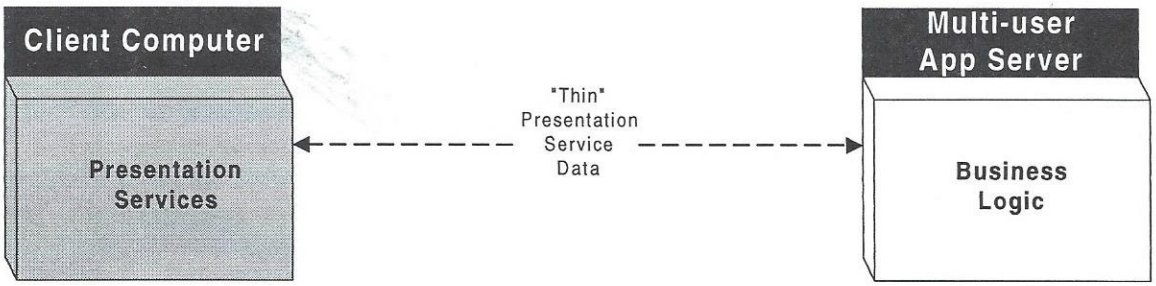

The thin-client/server model functions similarly to an N-tier client/server application that has been partitioned such that all business logic executes on intermediate application servers, leaving client workstations to display the user interface. However, they do differentiate: Thin-client/server architecture can centralize business logic of two-tier

client/server and file sharing applications, as well as N-tier applications. MIS managers familiar with Unix will recognize the approach's similarity to X Windows, which centralizes business logic and data access for Unix applications. In contrast to X Windows, multi-user application servers can run 16- and 32-bit Windows applications.

The thin-client/server model lets MIS deploy a single copy of the mission-critical application at a central site in the enterprise network. Local and remote clients, in turn, connect to the centralized application via LAN, WAN, or dial-up connections.

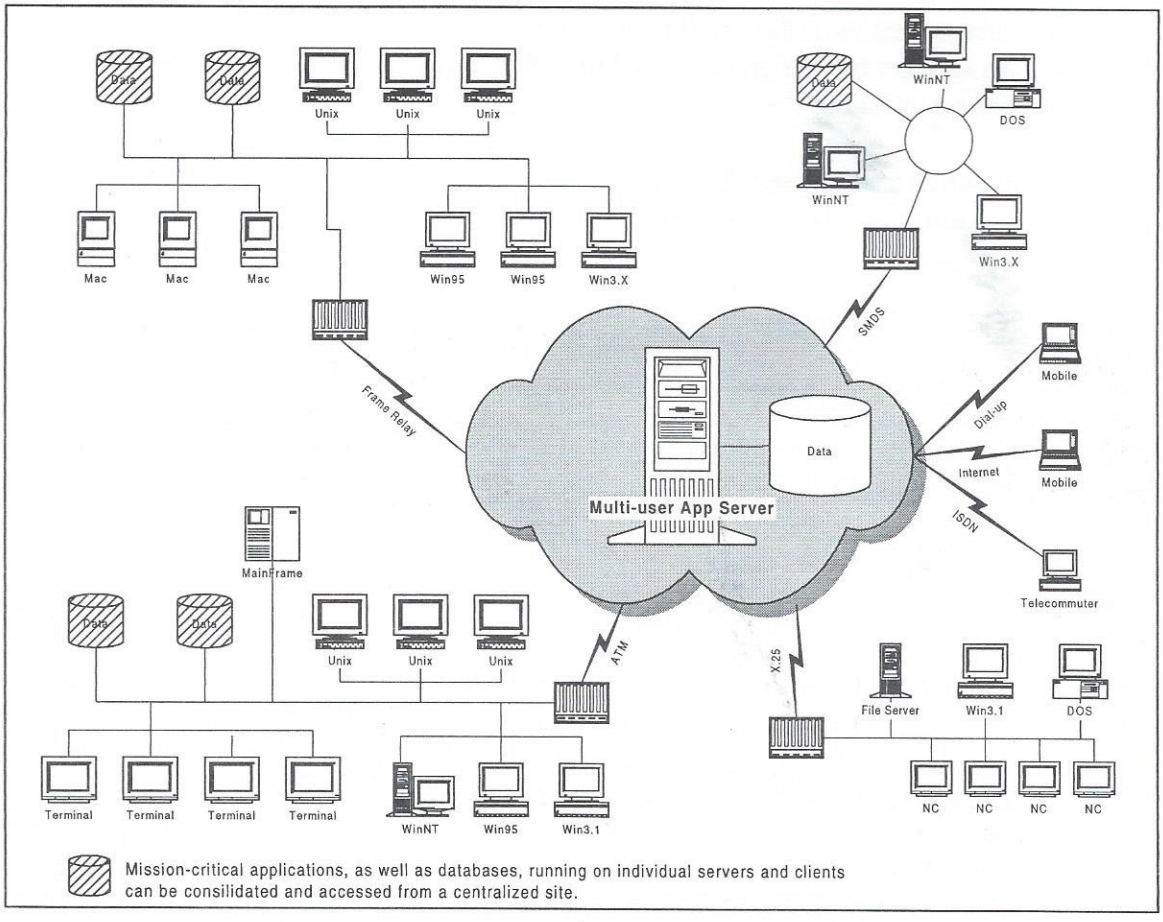

By consolidating the application on centralized servers and limiting the client workstation's job to display duties alone, the thin-client/server model has an assessment profile that differs sharply from the traditional application deployment approaches above.

> **Cost assessment:** Excellent. Centralized application deployment confines application installation, maintenance, and upgrades to one site. Centralization eliminates client-side application maintenance. Heterogeneous client support makes client-side upgrades unnecessary.

Management assessment: Excellent. Centralized deployment allows same-day application rollouts to thousands of users. Version control is likewise established, encouraging frequent, minor application updates. As an extension of the underlying operating system, multi-user application servers will let an application scale to thousands of users if the underlying OS supports such features as SMP server hardware and server clusters. Application security and management, too, will depend on the multi-user application server's underlying operating system. Security may include multi-level passwords and privileges, roving callbacks, encrypted login and data, and file-level security. Likewise, applications will run in any environment supported by the operating system, and they will be managed by any management tool supported by the operating system. Deployment flexibility includes support for remote node and remote control infrastructures as well as the Internet and intranets.

Performance assessment: Excellent. The client component delivers LAN-like response times across WAN links and dial-up lines.

	Cost	Management	Performance
Distributed LAN	Unsatisfactory	Unsatisfactory	Excellent
Remote Control	Unsatisfactory	Unsatisfactory	Good
Remote Node	Unsatisfactory	Satisfactory	Unsatisfactory
Thin-Client/Server	Excellent	Excellent	Excellent

Conclusion

In today's competitive business environments, organizations look to MIS for information tools that are fast, reliable and easy-to-use. Expeditious deployment of mission-critical applications can be a strategic weapon. However, today's dispersed network environments create a number of challenges for MIS staffs who need to deploy mission-critical software.

Using the criteria of cost, management and performance, MIS managers can assess traditional and emerging application deployment options and select the one that best meets their needs, maximizing the company's return on its application investment and delivering a competitive business advantage.

About Citrix

Citrix Systems, Inc. Is the pace-setter and a world leader in thin-client/server software solutions. The company's award-winning WinFrame software, which is based on ICA technology, provides access to virtually any corporate application, across any type of client. WinFrame is a cost-effective and proven solution that provides today's enterprises with centralized management, universal application access, exceptional performance and improved security for business-critical applications.

Independent Computing Architecture (ICA) technology provides the foundation for allowing any device - thin or fat - to serve as the ultimate thin client. It ensures that applications execute 100% on the server. This means that network applications consume as little as one-tenth of their normal network bandwidth over dial-up connections, WANs and LANs. This kind of efficiency enables the latest, most powerful 32-bit client/server applications to be accessed with exceptional performance from exiting PCs, Windows terminals, network computers and a new generation of business and personal information appliances.

CITRIX

6400 NW 6th Way
Fort Lauderdale, FL 33309
954-267-3000
http://www.citrix.com