

DOCUMENT RESUME

ED 368 260

HE 027 276

TITLE Managing Information Technology as a Catalyst of Change. Track I: Leadership during Times of Change.

INSTITUTION CAUSE, Boulder, Colo.

PUB DATE 94

NOTE 86p.; In: Managing Information Technology as a Catalyst of Change. Proceedings of the CAUSE Annual Conference (San Diego, CA, December 7-10, 1993); see HE 027 275.

AVAILABLE FROM CAUSE Exchange Library, 4840 Pearl East Circle, Suite 302E, Boulder, CO 80303 (individual papers available to CAUSE members at cost of reproduction).

PUB TYPE Speeches/Conference Papers (150) -- Reports - Descriptive (141)

EDRS PRICE MF01/PC04 Plus Postage.

DESCRIPTORS *Computer Networks; Computer Uses in Education; Educational Change; Higher Education; *Information Management; *Information Systems; *Information Technology; Integrated Activities; *Leadership; Models

IDENTIFIERS Arizona State University; *CAUSE National Conference; Chief Information Officers; Cornell University NY; *Distributed Computing; Kent State University OH; University of Alberta (Canada)

ABSTRACT

Eight papers are presented from the 1993 CAUSE conference track on leadership challenges facing managers of information technology during times of change at colleges and universities. Papers include: (1) "ASURITE: How To Avoid Creating a Distributed Computing 'Tower of Babel!'" (Neil Armann and others), which discusses the Arizona State University Rational Information Technology Environment; (2) "Integrating Computing Resources: A Shared Distributed Architecture for Academics and Administrators" (Monica Beltrametti), which describes the integrated electronic environment at the University of Alberta (Canada); (3) "Technology Planning for the Nineties: Responding to the Challenge To Do More Work with Fewer Resources" (Ilee Rhimes and Lawrence R. Kelley), which presents the results of an information technology planning study at Kent State University in Ohio; (4) "Lots of Data! No Information! (Why Universities and Colleges Do Not Take Full Advantage of Their Information Systems)" (Bethany M. Baxter), which calls for changes in organizational structure, resource allocation, and approaches to academic and administrative tasks; (5) "Leadership and the Changing Role of the Chief Information Officer in Higher Education" (Gary M. Pitkin); (6) "Strategies for Data Management Leadership: Why, Who, What, and How" (Lore A. Balkan and others); (7) "A Model for Change" (William R. Brunt); and (8) "Implementing Distributed Computing at Cornell University" (Mark Mara). (Some papers contain references.) (JDD)



Managing Information Technology as a Catalyst of Change

Proceedings of the
1993 CAUSE
Annual Conference

TRACK I

LEADERSHIP DURING TIMES OF CHANGE

December 7-10
Sheraton on Harbor Island
San Diego, California

PERMISSION TO REPRODUCE THIS
MATERIAL HAS BEEN GRANTED BY

V. Ryland

CAUSE

TO THE EDUCATIONAL RESOURCES
INFORMATION CENTER (ERIC)™

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

This document has been reproduced as
received from the person or organization
originating it.
 Minor changes have been made to improve
reproduction quality.

• Points of view or opinions stated in this docu-
ment do not necessarily represent official
OERI position or policy.

BEST COPY AVAILABLE

HE 027 276



TRACK I

LEADERSHIP DURING TIMES OF CHANGE

Coordinator: Don C. Wolfe

Universities and colleges have come to accept change as the only constant in their institutional processes. Information technology is and will continue to be in the forefront of those change agents that will distinguish successful institutions from the less fortunate. In this era of rapid change in mission, economics, and technology, what special leadership challenges do managers of information technology face?

ASURITE

How to Avoid Creating a Distributed Computing "Tower of Babel"!

**Neil Armann, L. Dean Conrad, Darrel Huish
Information Technology
Arizona State University
Tempe, AZ 85287-0101**

ABSTRACT

The entire computer industry is racing at breakneck speed to implement client/server applications and begin reaping the benefits of the new distributed computing paradigm. However, there are many ways to implement a client/server environment and institutions must plan their implementations to avoid technological anarchy and keep from creating a "Tower of Babel" in which the pieces do not communicate or cooperate. This paper discusses the distributed computing architecture Arizona State University has established to ensure all the new distributed pieces will actually work together. This architectural approach is called ASURITE, ASU Rational Information Technology Environment. We describe ASURITE and the impact this architecture will have on the academic and administrative computing communities.

I. INTRODUCTION

The Problem

Computers are getting continually more powerful and consistently less expensive. This has encouraged the pervasive use of information technology throughout the University. Despite our historical reliance on mainframe computing, the majority of ASU's investment in technology now sits on the desktop.

When viewed at the department level, this evolution toward distributed computing is a good thing. It allows the department to improve their operation without long delays waiting for external expertise. However, when technology is implemented in piecemeal fashion, units can find themselves unable to accomplish work that crosses organizational boundaries. This is usually because the various departments' computing environments have not been designed to work together. Computer people call this situation "Islands of Technology." For example, a faculty member may have a useful computing "island" for tracking records and grades of his students, yet find it impossible to send final grades to the University's student information system.

When faced with the obstacles presented by these isolated islands of technology, many organizations attempt to centrally control or coordinate technology implementation. Some organizations try to settle on a single vendor to insure that all pieces of technology will operate together. This approach is not practical for ASU because we have a high degree of autonomy within our various units; we have already made a large investment in a de facto heterogeneous computing environment; and even if we could all agree on a homogeneous set of technical products, we lack the massive budget required to replace significant portions of our technology.

So, we need a better strategy to address this problem. Departments need to answer three basic questions:

If I'm buying technology and want to maximize my ability to work cooperatively with others, what should I buy?

If I have limited budget but want to improve my technological situation incrementally, how can I evolve in the same direction as everyone else?

If some functions are going to be handled centrally for efficiency's sake, what tasks will be done for me and what tasks should I prepare to do for myself?

Approach

In a simplistic sense, the way to guarantee that we can cooperatively share technology is to identify standards. This approach has worked well in the area of audio music. We can all buy cassettes and expect them to work on any cassette player made by any manufacturer. The same is true for compact disks. One can plug a Sony amplifier into a Technics receiver with Bose speakers and easily construct a viable audio system. This is because the manufacturers adhere to standards.

Unfortunately, we don't yet have widely accepted standards for most aspects of computer technology. There are emerging standards that hold promise for improving our situation, but today there is no "plug and play" ability among all vendors. A significant obstacle to the ultimate success of any standard is the rapid pace of change among computer technologies.

ASU's Rational Information Technology Environment (ASURITE) is an information technology architecture that positions the university to take advantage of emerging standards. It also recognizes the need to accommodate budget constraints, moderate the pace of change and preserve the autonomy of the individual departments.

ASURITE describes a distributed style of computing that is constructed of modules. Each module performs a specific function and can be thought of as analogous to an audio component. The components are frequently called "servers." So, the computing environment at ASU will have several data servers which store and retrieve data. Several print servers will produce output at various locations. Mail servers will store and forward mail throughout the university, etc.

Some modules lend themselves to being implemented and supported centrally. For example, security can be best maintained by allowing a single method to gain access to computing resources. Customers would typically identify themselves during their first interaction with any computing resource and then be granted authority to all valid and appropriate services. One server can be maintained centrally rather than have multiple security checks with multiple procedures and passwords.

Departmental implementation and support is more appropriate for other modules, such as a database used only by a single department. But that departmental database may need to obtain some of its data from a central database, so the ability to interact with central services must be maintained.

ASURITE is an architectural framework which describes how all supported components will interact within such an environment. The architecture encompasses various styles of computing including client/server, distributed computing, cooperative processing and object orientation. It is intended to help ASU achieve flexibility, adaptability and efficiency in information technology, by putting processes on the right platforms, in the right location, and in a consistent manner.

ASURITE treats the individual as the focal point of a series of software "services" supporting the individual's dual role as both **data producer** and **information consumer**. In general, services can exist on any combination of hardware and physical locations deployed in an "intelligent" network. It is primarily the desktop which invokes the services where and when needed to satisfy an individual's need.

The desktop will become the focal point of the individual's interaction with enterprise systems and data, as well as with collaborative groups, research data bases, etc. Data, voice/sound, graphics/images, and live video will converge on the desktop as the common denominator for synthesizing information from data. All applications will reside in a robust, intelligent network which presents the information consumer with a single system image. ASU systems and links to the external world will appear to be a single network composed of services and data, invoked by name regardless of the physical locations and technology used to provide those services and data.

The diagram in Appendix A provides an overview of the ASURITE as it will be implemented over time.

II. GENERAL CHARACTERISTICS

In order to achieve the overall ASURITE architectural objectives, the following qualitative characteristics are established:

Adaptability - change as national & industry standards evolve, so we can enhance and incorporate new ways of doing essentially the same business function without major developmental impact.

Manageability - centrally manage or coordinate and monitor, including the orderly planning for capacity changes of various essential services.

Reliability - remain in continuous operation even if part of the system suffers failure, needs maintenance or upgrading, or is destroyed or damaged by a disaster.

Securability - provide different access to individuals based on the classification of data and the user's business function. This will require that all basic ASURITE services use standard (ASURITE) authentication and authorization services.

Extensibility - easily add new kinds of functionality to existing processes without major impact.

Scalability - increase or decrease size or capability in cost-effective increments without software impact or "spikes" in the unit cost of operations due to step functions in procuring additional resources.

Performance - fast response and high throughput.

Connectability - communications access to a variety of area, national, and international networks.

Consistency - relative stability of the person/machine interface over time.

Accessibility - university community members should be able to access and use ASURITE services wherever they are, provided that they have a properly configured "client" workstation.

III. CHARACTERISTICS OF INDIVIDUAL SERVICES

The following qualitative objectives are established for each individual service offered within ASURITE:

Like an extension to the desktop - how information is presented to, manipulated by, or provided by the user needs to be consistent across all applications no matter where the application is actually running -- on the desktop or on a network server. The user interface can be made more consistent by making it appear as if the information is completely under the control of the workstation software with which the individual user is already familiar. Just as the user can tailor the workstation software to satisfy her needs, so should she be able to tailor the interface for information from and to external systems.

Interoperable - (1) any supported service is available to any supported client no matter the particular brand of server or client hardware and software and (2) the interaction between clients and servers is transparent to the client, e.g., the client does not need to know where the service is coming from.

Incorruptible (virus-free) and as secure as practical - computer viruses are detected and prevented from spreading to servers and clients, and data and computer systems are protected from unauthorized use and tampering. Absolute guarantees, however, of virus prevention and security are not feasible.

Fault-tolerant - reduce the impact of hardware and software failures. Highly critical servers might have redundant processors and databases so recovery from a failure would be immediate and transparent to the user; other, less critical servers could have backup servers that could be put into operation within a few hours.

Disaster-tolerant - restore services in a timely manner when a disaster, such as a fire, destroys equipment.

Expandable - additional capacity can be added to meet the demands of more users or increased functionality without modification to user procedures.

Peer to the client - no master/slave relationship should exist between client and any server - the client is not controlled by the server. A client makes a request of a server and is prepared to receive a response (or a request to supply more data to the server) but is free to do other processes in the meantime.

Restricted in access as needed - since all services are technically accessible to any user on the network, individual service providers may limit access to their services if necessary; e.g., a departmental printer may be restricted to use by members of the department.

Non-interfering and non-conflicting with other services - any user can use any service or combination of services concurrently.

Appropriately interactive with the client - the client can monitor and alter its requests for services. The server should make the status of a request for service available to the client so the client/user can cancel or modify the request if needed; e.g., it should be possible to determine that a data base query is retrieving much more data than anticipated and to cancel that query if desired.

Optional - local (to the client) services are allowed; e.g. a local printer may be completely under the control of the local client.

IV. SPECIFIC SERVICES THAT WILL BE PART OF ASURITE

- A. The following services are considered basic to ASURITE, and must be in place before other services can be added:

Authentication and authorization - authentication is the verification process that confirms the identity of a person requesting service; this process must be done in a secure manner to prevent others from determining the method of verification. Authorization is the process that permits only those who have been granted permission to use a particular service to actually use that service.

Finder/navigator - permits users, clients, and servers to reference services, devices, and people by name rather than by physical location or network address. These services also permit transparent relocation of devices, servers, etc.

Time - synchronize date and time of day on all the servers and clients so that time-dependent processes are coordinated.

File management - provides for the storage, access, and security of data (e.g., text, images, and voice) particularly to facilitate the sharing, interchange and security of data. File management services include

backup and recovery services make duplicate copies of data in case the working copies are damaged and provide procedures to restore lost data from the backup copies, and

archiving services provide facilities to store and retrieve seldom used data on low-cost media, such as tape.

Print - provides for the transmission, temporary storage, and production of paper output of data (including text, plots, and images) from clients and other servers.

Configuration management - set of services to coordinate the software and hardware on the servers and clients and includes

notification of changes,

update by subscription,

coordination of non-optional upgrade of software, and

verification of hardware compatibility.

Network management data base and status - maintains data concerning the network configuration and operations.

- B. The following are examples of services or classes of services that will be added to the basic ASURITE services over time:

Collaboration support is a set of services that facilitate human communication between individuals, within a group working on a common effort, or among groups interested in a particular topic. These services include messaging, computer-facilitated conferencing, electronic mail, voice mail, calendaring, and groupware.

Enterprise applications are those computer applications that support the major administrative functions of the university.

Object catalog contains data about enterprise data (e.g., names, descriptions, and usage rules) and common processes using enterprise data (e.g., a complex query that extracts data from several data bases and puts them in a spreadsheet). An object catalog is an extension of the data dictionary concept.

On-line consulting is a repository of information and previously asked questions and answers to help users and support personnel solve hardware and software problems.

Computation servers handle resource-intensive calculations that are inappropriate for running on a local workstation.

Software library services provide for the distribution of shareware or site-licensed software and the lending of software for trial use.

Databases services make information available to any client and are provided by commercial, research, administrative, and other sources (e.g., the administrative data warehouse).

Scheduling of tasks services control processes that do not need to be run immediately, e.g. long reports or database backups.

External access services permit use of the servers from locations outside of the university-operated network, e.g. from Internet sites or via dial-in from home.

Problem reporting and tracking services receive and facilitate the resolution of system problems.

Approval and signature services permit the electronic (i.e., sans hardcopy) authorization of official documents, e.g., purchase requisitions, grades, and payroll.

V. PRODUCT ARCHITECTURE

The intent of this section is to list some specific products and standards that currently appear to comply with the general characteristics and services that form the core of ASURITE. The list of products and standards is incomplete because in some cases no product or standard exists that conforms to the ASURITE architecture. However, current products and standards will evolve and new ones will emerge to fill in the gaps.

An example of emerging distributed computing standards are the Distributed Computing Environment (DCE) and Distributed Management Environment (DME) standards provided by the Open Software Foundation (OSF). OSF is an independent company formed by a coalition of computer and network product suppliers. Its goal is to provide a set of open industry standards for distributed computing. Manufacturers that conform to these standards are assured of interoperable products. Because DCE and DME also provide extensive coverage of the services listed earlier, ASURITE will be relying heavily on products that use them.

Users view ASURITE from their desktop workstations, each with the individual's own preferred method of interaction. In acknowledgment of this, ASURITE will provide support to the general community for desktop Operating System-Graphical User Interface (OS-GUI) combinations that will interact with appropriate servers. The initial set will be:

DOS 5 - Windows 3.1
Mac System 7.1
UNIX-Motif.

These OS-GUI are expected to evolve or be replaced over time. Possible successors include Windows NT and Windows for Workgroups for DOS/Windows and AOCF for the Macintosh.

The primary network communications protocol set required to obtain ASURITE services will be Ethernet and TCP/IP. However, protocols in addition to TCP/IP, e.g., Appletalk, IPX, and DECNET, will also be supported on the university backbone for use by individual work clusters and departmental local area networks for the near term. It is expected that additional network capacity will be needed at ASU, particularly on the University Backbone, to support all of the new services, so ASURITE will be evolving to new, higher speed protocols as needed and funds allow. Examples of new network standards being considered are FDDI and ATM. In the near future higher speed protocols will be used primarily on the backbone and to connect high volume servers and workstations, but the vast majority of workstations will be connected via Ethernet.

With the large number of LAN's at ASU, it will be important that ASURITE coexist with the LAN's and interoperate with LAN services such as file sharing, local mail, and print services. Client workstation connectivity via Ethernet provides access to both LAN services and ASURITE services even though the LAN might be using a non-TCP/IP protocol. The ASURITE file service, initially, will use the Andrew File System (AFS) which at present can only be directly used by UNIX clients. However, Network File System (NFS) software can be installed on Mac and DOS/Windows workstations and access AFS files via a NFS/AFS translator provided as part of the ASURITE file service. To the user the AFS files appear in the directory as if they resided on the workstation and can be moved to a LAN server or the workstation by simple drag and drop of the file icons. LAN vendors are expected to incorporate interoperability functions in the future so intermediary services, such as the NFS/AFS translator, are not required.

The following tables summarize the ASURITE standards and related products for basic services, client workstations, and communications. These form the components of ASURITE, but, of course, must interact with each other to form a cohesive architecture.

Table 1: Basic Services Product Architecture

Service	Future Standard	Initial Product	Future Product
Authentication	OSF/DCE	Kerberos v. 4	Kerberos v. 5
Authorization	OSF/DCE	native OS access control	DCE Access Control List
File	OSF/DCE	Andrew File System (Transarc)	DCE Distributed File System
Time	OSF/DCE	Internet, Network Time Protocol	DCE Distributed Time Service
Finder/Navigator	OSF/DCE X.500	X.500 Internet Domain Name Service	X.500
Print	OSF/DME	TCP/IP LPR/LPD Palladium	DME Print Service
Data Base Management	SQL Access Group	Sybase Informix Oracle	Sybase Informix Oracle
Data Base Transaction Management	SQL Access Group X/Open-XA	Encina	
Configuration Management	OSF/DME		
Network Management Database	OSF/DME		
Email	X.400 SMTP	SMTP/IMAP	X.400 SMTP/IMAP
Calendaring	None yet		

Table 2: Client Product Architecture

Function	Initial Product	Future Product
Operating System / Graphical User Interface	DOS 5.0/Windows 3.1, Mac System 7.1, UNIX/Motif	Future versions compatible with initial products
Data Base Access	Sybase Open Client, Microsoft ODBC, Borland IDAPI	SQL Access Group standard
Communications card	Ethernet	Ethernet; FDDI on high volume workstations
Communications software	TCP/IP	TCP/IP

Table 3: Network Architecture

Function	Initial Product	Future Product
Wiring	Broadband coax with some fiber optic campus backbone to building router, Broadband coax to floor closet, Twisted pair to workstation	Fiber optic backbone and to high volume workstations, Twisted pair to most workstations
Low-level communication protocol	Ethernet with limited FDDI on backbone	FDDI backbone and to high volume workstations in short term, ATM or other long term; Ethernet to most workstations,
Transmission protocol	TCP/IP (see Note)	TCP/IP

Note: The backbone network also supports DECNET, IPX, and Appletalk for use by departmental LAN's, but access to the ASURITE services requires TCP/IP.

Vi. Impact on Academic Computing

While ASU struggles with the challenge of implementing distributed computing for our administrative systems, much of academic computing is already distributed. Starting in the mid-1980's, the administration began a series of "faculty computer infusion" programs where the goal was to place a computer on the desktop of every faculty member that desired one. Most of these were PC and Mac systems. This was followed a couple years ago with a "networking infusion" program where the goal was to get the faculty connected.

The faculty infusion programs were very successful. However, as we began to look at implementing the ASURITE architecture, we realized we would have to do something about all the older, obsolete technology that is the legacy from the earlier programs. The architecture provides the promise of much additional function and flexibility, but it depends on a minimum configuration that will run Windows 3.1 or Mac System 7. Meanwhile, we had literally hundreds of faculty members with PC-XT's, early Macs, or worse on their desktops. Furthermore, we had many old, obsolete systems in our student computing sites. It would require a mainframe-level expenditure to resolve this problem.

Fortunately, we were able to enlist the support of our Provost and Senior Vice President. At the end of FY93, the Provost funded a new infusion program that allowed ~700 faculty to upgrade their systems, which addresses roughly half of the obsolete systems. A similar infusion program is anticipated for FY94 to address the remaining half. In addition, Information Technology (IT) also received funding that allowed us to install ~350 new systems in the student sites. The base systems purchased were 486 and Centris 610 configurations. These upgrades will allow virtually every faculty member and many students to participate in the ASURITE environment.

Over this next year, IT will be focusing on deploying the ASURITE basic services. At the same time, several projects are being initiated or extended -- by IT as well as other units on campus -- that will utilize and build upon these basic services. These include:

- o a distributed client/server e-mail system (including forms routing and calendaring) to replace the present mainframe-based environment -- this will be extended to all faculty, staff, and students as opposed to the present limited availability;
- o a "self-subscription" service for e-mail and other popular services that will allow us to automate the accounts process for the massive number of accounts that will need to be created;
- o transitioning customers off the academic mainframes by deploying adequate statistical, compute, and file servers as well as migrating many faculty primarily to their desktops;
- o a dial-in Ethernet service for access to all ASURITE-compliant services from virtually anywhere;
- o a server for downloading software products including the clients that will be needed to access the many servers that are being deployed, e.g., e-mail, Library CD-ROM, and GOPHER clients;

- o a Netnews/Usenet server;
- o docking stations in the student computing sites to allow students to bring in their own laptops and plug into the network;
- o an on-line CD-ROM data base server in the Library;
- o a high quality color print/copy server; and
- o a slide maker server for presentation (deferred to FY95).

A robust distributed academic computing environment promises to provide powerful new tools and capabilities to the faculty. However, careful planning will be needed to deliver on this promise. We have initiated a joint planning effort between IT and the colleges to help ensure we fully understand the needs of each college and that faculty and students experience the minimum amount of discomfort as we begin the evolution to this new environment.

VII. Impact on Administrative Computing

ASURITE has had a significant impact on ASU's administrative computing. The technical architecture is now widely accepted as an important part of any decision making about technical projects. There are at least three reasons for this. First is because ASURITE offers guidelines not mandates. In a political environment where autonomy is highly valued, people respond better to a suggestion than a requirement. Second, departments are going to spend funds on some form of technology anyway so they appreciate published information about how to make things work together. This is related to the third reason which is that ASURITE provides some insulation from blame. We have had administrative units call us and ask how they could get their project "certified" as ASURITE compliant. We have resisted the idea of a certification process as too bureaucratic but some customers have been persistent. We have done some informal consulting and recommending about purchases, but don't yet certify projects as architecturally sound.

So, what's really happening?

It isn't enough to have an architecture. We must also build applications that will demonstrate the value of distributed computing. To that end, ASU has acquired an application development environment that will build client/server applications from an enterprise model. We will be using the IEF from Texas Instruments as the basis for automating new student information processes.

ASU is also implementing a server-based data warehouse. Initially populated with student data, we expect to follow soon with financial and human resource data. Customers will form queries and manipulate the results from their workstations.

We are also seeing a strong growth in gopher style services. One project may make admissions information available to high schools while another is as simple as electronically publishing the minutes of the advisory committee meetings.

There is also an emerging need to interoperate between new distributed style applications and legacy IDMS applications. We have a financial aid downsizing project underway that will have data residing on a server as well as using data on our old IDMS database.

Challenges

There are, of course, challenges. One is a feeling of false security that can come with being aligned with the technical architecture. Some departments may ignore other important issues in distributed computing. For example, the need for departmental disaster recovery planning is still very real. In addition, we need to have applications that are not only technically compliant but also consistent with the university's business model and production standards.

Another issue is funding. Who pays for the new costs associated with this style of computing? ASU has managed to provide additional funding to place more workstations on faculty desks, but we haven't yet been able to do the same for administrative staff. Another more concrete example is that we've had customers resist using data access servers to print their own mailing labels because they would then need to buy labels. Under our centralized model, the labels had been provided for "free." From the

institutional perspective this should be a non-issue, but for a struggling department all expenditures are significant.

VII. Summary

We face many critical issues as we deploy ASURITE-compliant services. The following are the most immediate; resolution of these will be crucial to achieving our vision of a distributed computing future:

- o funding for servers: after the past three years of budget cuts, IT no longer has the funds to provide all campus-wide computing services;
- o funding for maintenance: for the environment to work smoothly and be reliable, we will need to carefully coordinate release levels for critical software components on clients and servers alike -- however, it's very difficult to dictate how other units spend their budget dollars;
- o network management: problem and change coordination for critical components of the network is already a problem on our campus and the situation will only get worse; and
- o systems management: the faculty do not have time to become systems managers, yet software will need to be upgraded, system backups taken, and files archived/restored -- we are actively pursuing solutions with various vendors, but not all of the necessary pieces appear to be available as yet.

Establishment of an architecture will make addressing these challenges possible and will ensure the resulting systems environment will deliver on the functions and flexibility that a distributed computing paradigm promises. Perhaps the fundamental issue is the client/server paradigm itself. Can all this stuff work? We hold frequent meetings under the heading of "Client/Server Reality Check." We focus on what products and services will be needed as compared to actual product availability. The frequent conclusion to these reality checks is "yes, we'll get there--but it won't be easy and it won't be soon." The key to ASURITE is keeping in mind that it will evolve and will never be "completed." It is a journey, but not a destination.

Without an architectural framework we risk computing anarchy and creation of a distributed computing "Tower of Babel!"

REFERENCES

- Apple Computer, Inc. *VITAL Fundamentals*. Apple Computer, 1993.
- Berson, Alex. *Client/Server Architecture*. McGraw-Hill, 1992
- Digital Equipment Corporation, ed. *Project Athena: The First Five Years*. Digital Equipment Corporation, 1988
- Open Software Foundation. *Distributed Computing Environment Rational*. Open Software Foundation, May 1990.
- Rosenberry, Ward, David Kenney, and Gerry Fisher. *Understanding DCE*. Sebastopol, CA: O'Reilly & Associates, 1993.
- Tapscott, Don, and Art Caston. *Paradigm Shift: The New Promise of Information Technology*. New York: McGraw-Hill, 1993

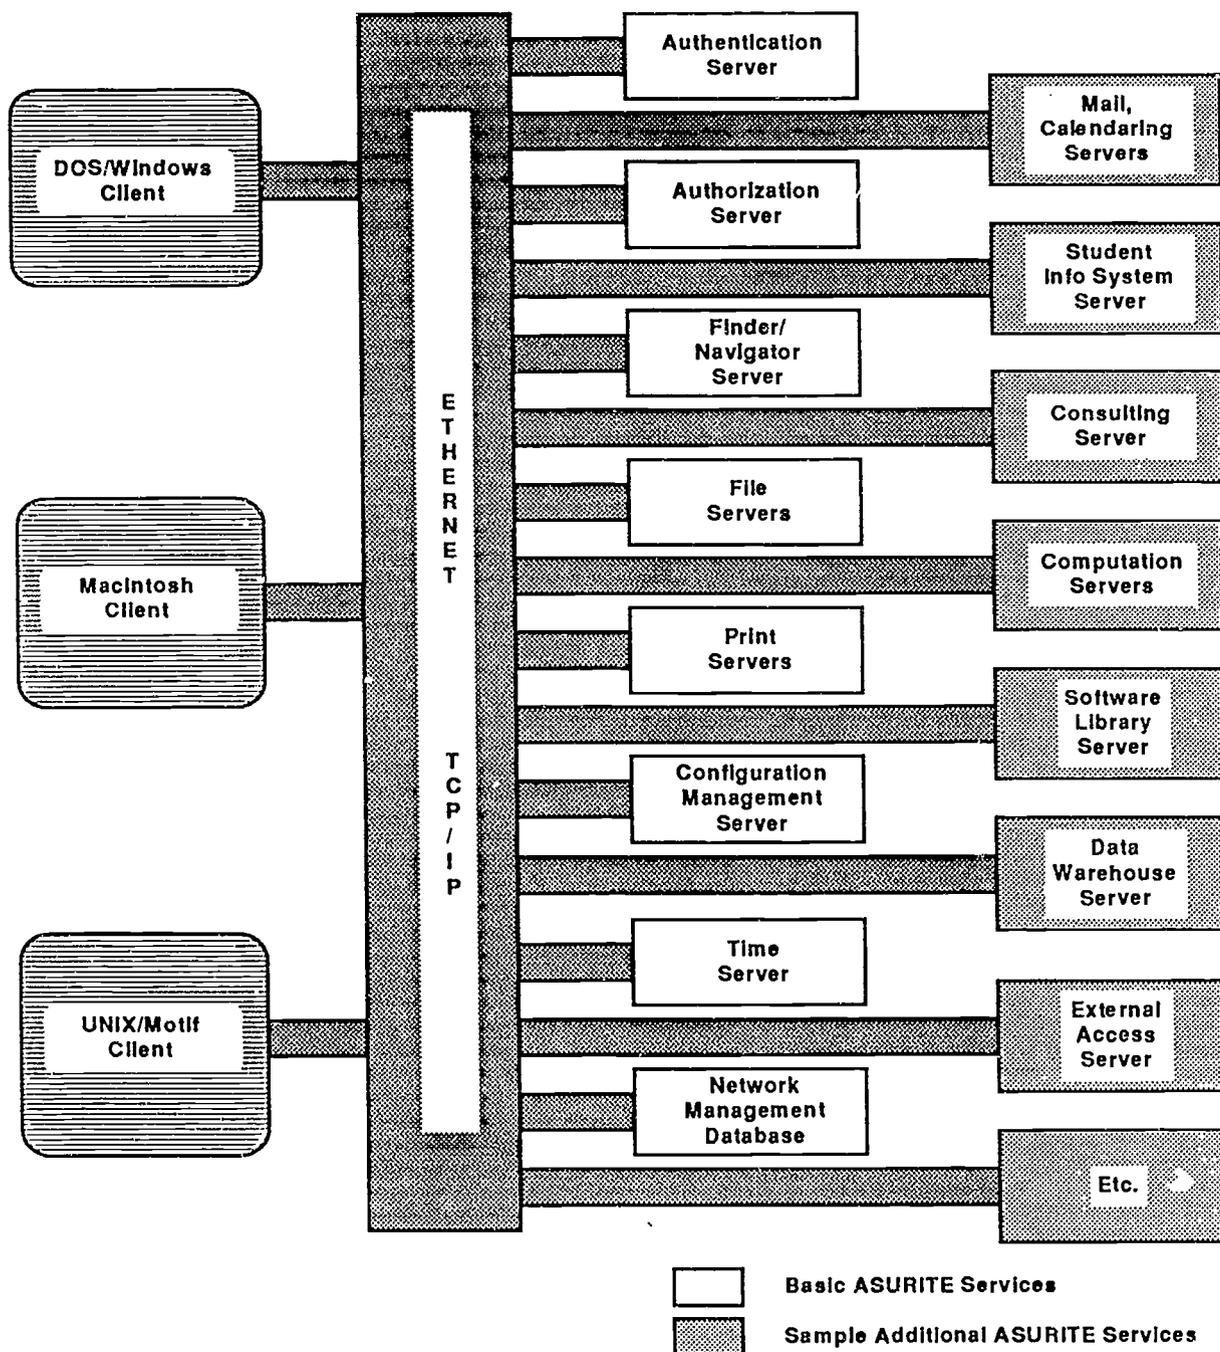
ABOUT THE AUTHORS

Mr. Armann is Director of the Computer Center at Arizona State University (ASU) and can be contacted at (602)965-5263 or via e-mail at narmann@asu.edu. Mr. Conrad is Director of Computing and Network Consulting Services (Academic Computing) at ASU and can be contacted at (602)965-5620 or via e-mail at larry.conrad@asu.edu. Mr. Huish is Director of Administrative Information Technology (Administrative Computing) at ASU and can be contacted at (602)965-5674 or via e-mail at darrel.huish@asu.edu.

ACKNOWLEDGEMENT

The development of ASURITE was a team effort involving many people at ASU besides the authors of this presentation. In particular, Connie McNeill, Dr. William Lewis, Dr. Bruce Millard, Gerry Samchuck, and Jack Hsu provided the leadership and made significant contributions.

APPENDIX A
ASURITE Overview



Integrating Computing Resources: A Shared Distributed Architecture for Academics and Administrators

**Dr. Monica Beltrametti, Director
Computing and Network Services
University of Alberta
Edmonton
Alberta**

Abstract

At the University of Alberta we have defined a computing architecture that enables the University community to share computing resources distributed across campus. This architecture provides users with an integrated electronic environment, with easy and transparent access to academic and institutional data, high performance computing, printers, and other peripherals. The sharing stretches beyond departmental boundaries. On demand, researchers can, for example, share each others' workstations, and take advantage of idle cycles on servers that are otherwise managing institutional data.

We explain how the vision of shared distributed computing was developed and how it is being implemented: we give a high-level description of the architecture and the users' view of the electronic environment; we show the budgetary advantages of the project; we explain the technical, and especially the managerial challenges behind and ahead of us; and we list the campus-wide human infrastructures necessary to manage such an integrated electronic environment.

Some Historical Background

During the mid-eighties computing at the University of Alberta underwent some dramatic changes. As the economic boom of the oil industry declined, funds to the University shrank quickly. Consequently capital funds for computing were reduced from \$8 to \$1 million a year. This prevented the University's central computing organization from keeping up with technology. The reduction of funds happened at the same time as workstations became affordable. Researchers and University staff did not hesitate to take advantage of this new commodity and started realizing their own computing solutions using their own funds, such as research grants or re-directing other sources of money. This led to a peculiar environment: at the end of the eighties the University computing organization was managing second or third-hand computing equipment, while the campus (having 7,000 staff members and 35,000 students) had bought state-of-the-art desktop computers by the thousands (see Table 1).

In the early nineties, users started to perceive the burden of managing their own computing equipment and actively demanded that the central computing organization change in nature to provide support for distributed computing. The users also demanded that computing at the University be raised in profile to match the standards of computing that the University had enjoyed during the seventies when funds were generous and computing at the University could compete with the top universities in North America.

This is why in 1991 Computing and Network Services (CNS), the computing department at the University of Alberta, formulated, in collaboration with senior University administration and the user community, a strategic plan in an attempt to define a road map that would restore the state-of-the-art computing and support that the University had enjoyed earlier (see "Computing Services Planning, Downsizing, and Re-organization at the University of Alberta" proceedings of the CAUSE conference, Dallas, December 1992; and "Networking Computers and People on Campus and Beyond" Computing and Network Services strategic plan, University of Alberta, February 1992.).

While we formulated the strategic plan we knew, however, that one important parameter had changed: the level of funding that had been enjoyed in the past would not be available again. The University President, Dr. Paul Davenport, recommended that the computing organization transfer salary funds to capital funds at a rate of \$200,000 a year for five years, which would result in the department having an additional \$1 million a year for capital allocation. Although painful to staff, this recommendation helped improve our financial position.

The capital funds we could count on were, however, still not large and it was clear that raising the profile of computing despite this meant that we had to develop a new way of thinking. Besides financial considerations, several other factors contributed to conceiving a new computing environment for the campus.

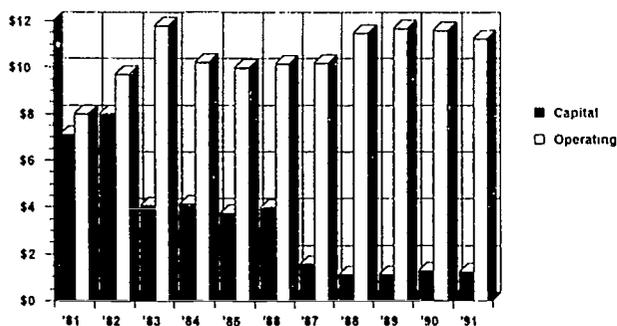


Figure 1: Retrospective of CNS Capital and Operating Funds (millions of dollars)

Central Mainframes	
Amdahl 5870	MIS
Amdahl 5880	MVS
IBM 4381-R14	VM/XA
Convex C210	UNIX

Campus Workstations	
Microcomputers (PCs, MAC's, Clones, etc)	10,000 units
Workstations (RS-6000s, SUNs, SGIs, etc)	1,000 units
Operating Systems	UNIX, DOS, etc

Table 1: Central Computers and User Environments at the University of Alberta, 1993

Rationale for Defining a Shared Distributed Computing Architecture

Implementing the Strategic Plan

The first shift in our thinking occurred when we started implementing the strategic plan. The plan contains fifteen strategies. They call for:

- various aspects of distributed computing support, such as installation of a high-bandwidth campus network, services to help clients with the management of their own computing equipment, and the establishment of guidelines for purchasing equipment to ensure adequate support and interoperability
- increased computing capacity for both researchers and administrators
- integration and modernization of administrative applications
- enhanced support for instructional computing (instructional labs and computer-assisted instruction tools).

Originally, we distributed the strategies among the CNS management team with the mandate to prepare operational and financial plans for each strategy and then provide the resources and project management for implementation. We soon found out that this was not the optimal way of implementing the strategic plan. Several difficulties arose.

Need to Abolish the Distinction between Administrative and Academic Computing Support Staff

No strategy could be implemented within the jurisdiction of one single manager. For example, integrating and modernizing administrative applications using a new client/server architecture could not be implemented by the Information Systems group alone, which in the past had primarily dealt with mainframes. Expertise was needed from staff with networking and workstation experience. Similarly, implementing a distributed open system environment for research could not be implemented by staff with only workstation background; expertise in system software and networking was required. We were very fortunate to be able to address this problem easily, thanks to some organizational changes that had occurred previously.

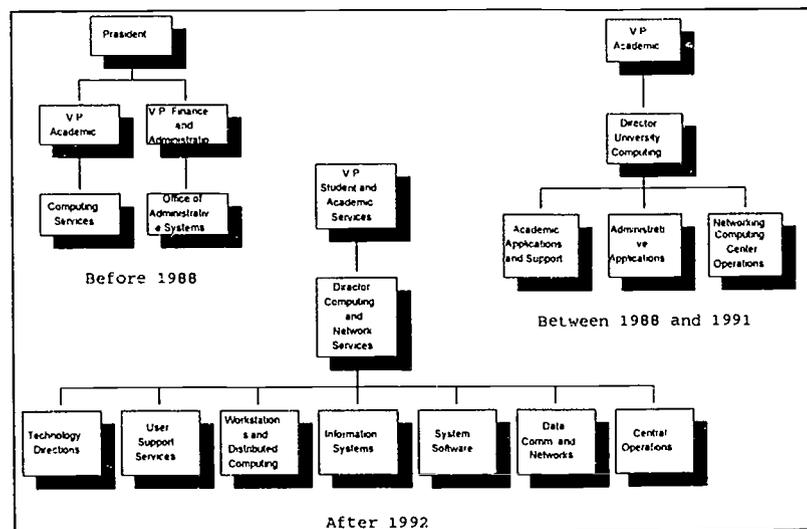


Figure 2. Retrospective of Computing Reporting Structure

The University of Alberta used to have two computing departments: the Office of Administrative Systems and Computing Services. The former department reported to the V.P. Finance and Administration and the latter to the V.P. Academic. Following suggestions put forward as early as 1985 by some visionary University staff, in the spring 1988 the two departments were put under a single director reporting to the V.P. Academic. At that time the two computing rooms were amalgamated; the rest of the organization, however, continued to operate as two independent entities. Administrators and researchers used different equipment and sought support from different staff. This led to substantial duplication of services and in 1992 we carried the organizational change a step forward by eliminating in the computing organization the distinction between academic and administrative computing.

reorganizing along technology lines. Even though the new organization was efficient to give day-to-day support to clients, it was still too rigid to implement the strategic plan effectively. Thanks, however, to a new way of thinking that had developed over the years it was easy to break departmental barriers even further. We established task forces with staff from various CNS groups and technical backgrounds. Each task force has a leader and has the mandate to implement one strategy in the plan. By mixing staff with different backgrounds new ideas soon started to flourish.

Need for Architectural Cohesiveness and Integrity

It was good to generate new ideas, but we were left with the problem of organizing and selecting them to ensure effective and compatible implementation of the strategies defined in the plan. Filtering ideas was difficult since within each task force we had staff with different technology backgrounds and beliefs: some were wedded to mainframes, while others had embraced workstation technology to the extreme. Also, due to the intricacy of distributed computing some strategies had overlapping requirements. Since each group was working without an overall technological guideline, groups tended to find different technology solutions for the same requirements. For example, implementation of electronic forms and general purpose communication required e-mail solutions. These solutions were being looked at separately without an attempt to coordinate.

We therefore felt the need to define clear architectural directions for computing and networking, which could guide CNS and the users in their decisions. Will English, who had helped pinpoint these problems, was appointed manager, Technology Directions, with the mandate to define a computing and network architecture for the campus, taking advantage of as much staff wisdom as possible. To cause minimal disruption to the work of the task forces, he decided to formulate and publish the architecture plan a chapter at a time, starting with the most urgent issues that needed coordination between the task forces.

Guiding Principles

The campus computing and network architecture is being defined with some basic principles in mind:

- *Do Not Resist the Distributed Computing Trend:* Clients had been frustrated for many years by the obsolescence of the central computing power and by the backlog of work in administrative applications development. As a result many clients had already found alternative solutions using workstations. It was clear that distributed computing was happening no matter what. New solutions had to acknowledge this and had to accommodate the desire of clients to choose their own solutions while still benefiting from a central coordinating infrastructure.
- *Avoid Paying Mainframe Bills:* Our budget was too tight to be able to think about mainframe upgrades. We had to find a way to increase and modernize the computational power using cheaper technology.
- *Capitalize on Existing University Investments:* Our clients had already invested heavily in new technology by purchasing workstations by the hundreds. We had to find a way to capitalize on this investment.
- *Create an Integrated Electronic Environment:* It was clear that we could add value to the investments that already existed on campus by tying together the distributed computing resources; for example, by providing easy access to academic and institutional data and to special devices such as high speed printers or high performance computers.

Breaking the Barriers between Academic and Administrative Computing

As the architecture was being developed an idea emerged which we think is fundamental to changing the way computing resources are used on campus. For decades academics and administrators used different computer systems. With the way technology has evolved lately, this is no longer necessary, nor desirable.

A decade ago when many scientists started moving away from aging mainframes, seeking faster response times from workstations, manipulation of large amounts of institutional data could still only be handled by mainframes. Today, the same workstations being used for scientific work can also handle large data bases. If used wisely, they can entirely substitute mainframes for that use. Furthermore, in the past, UNIX was primarily used for scientific computing. Today UNIX, if used in connection with other tools, can also handle secure manipulation of data.

It seems that today many ingredients are available to define a uniform architecture for both academic and administrative computing. Pushed by the desire to make use of economies of scale and by making full use of campus computing resources, the CNS staff defined an architecture that allows academics and administrators to share resources across campus. Even though the sharing comes with considerable management challenges, there is strong support on campus to make the shared environment work. We describe next this shared distributed architecture.

Description of the U of A Shared Distributed Architecture

The Network Backbone

Fundamental to a distributed architecture is a solid network. Distributed resources on campus are harnessed via a fast network backbone. The backbone being installed is a 100 million bits per second fibre optic FDDI network installed along the University tunnels in a figure of eight (see Figure 3). The network is being connected to concentrators at preferred ring node locations (denoted with letters in the figure). The concentrators are used to extend the fibre from each ring node location to the machine room of adjacent buildings. The fibre optic cable has 24 strands. Twelve strands have been reserved to implement the distributed architecture described below. The rest of the strands have been reserved for future applications, such as energy and utility management.

The network is being installed in phases. The north route is functioning, the rest of the ring will be ready by 1994. The ring is being financed by converting salary funds into capital.

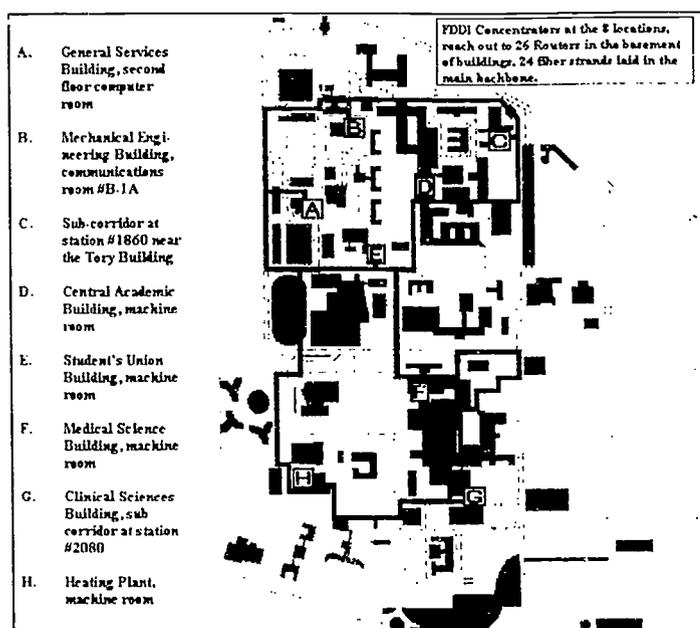


Figure 3 University of Alberta FDDI Backbone Network

The Campus Ring

Two strands of the fibre optic backbone are used to form a general purpose Campus Ring. The Campus Ring will harness most computers on campus. This is done by connecting the fibre in the machine room of each building to routers which, in turn, connect to the departmental local area networks (LANs) (see Figure 4).

The Campus Ring is being used to provide general purpose services, such as e-mail, file transfers, distributed printing, and manipulation of institutional data. Institutional data will reside on six enterprise servers placed at strategic locations on campus and connected to the FDDI backbone network. To provide maximum performance,

each Enterprise Server will be located near the local area network of the administrative unit that the server primarily supports. Information and applications of interest to only one administrative department will be handled within the department's local area network. The servers will handle information and applications that are of interest for all campus users. The Servers will support enterprise-wide applications, such as:

- enterprise data used by many applications across the campus
- enterprise file libraries housing the programs and files needed to support the users in accessing the enterprise data
- enterprise services supporting cross-platform applications, such as electronic University form processing or enterprise document management.

We are currently installing new information system tools which will eventually handle this client/server architecture. Some of the tools will have to evolve further before they will be able to handle our desired environment.

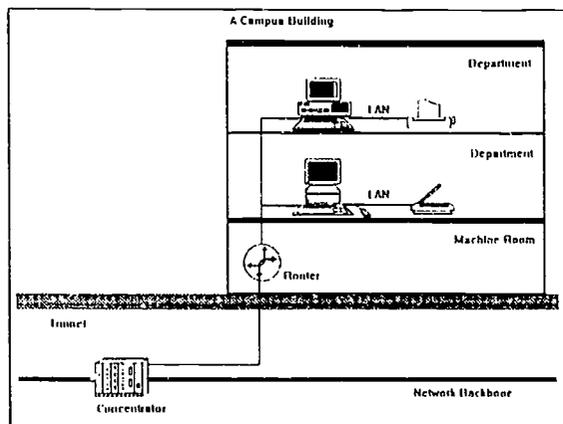


Figure 4 Connecting Department LANs to the Network Backbone

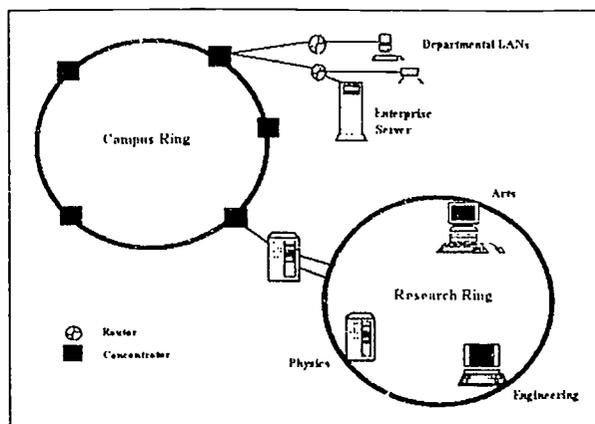


Figure 5 Campus and Research Rings

Research Ring

Two strands of the backbone will be used to network research workstations across campus. For each workstation, four fibre strands will be extended from the nearest concentrator to the workstation. The goal is to better utilize the campus investment in workstations by allowing researchers to execute their jobs on idle workstations. The ring can contain workstations and larger systems from different vendors. The workstations connected to the ring can be on researchers' desks or in instructional labs. The jobs submitted to the ring can be sequential or parallel:

- *Sequential Program Execution:* Researchers will be able to utilize workstation cycles on the ring to execute their sequential jobs. Each workstation on the ring will run an NQS-like software package (Network Queuing System). The software allows users to submit a job to a queue on a particular machine on the ring or to submit a job to a "ring queue" which will then select the best available machine on the ring.
- *Parallel Program Execution:* Researchers will be able to utilize workstations on the ring for parallel processing. The workstations will run the Parallel Virtual Machine (PVM) software package from Oakridge National Laboratory, a tool which provides an environment for designing, testing, and debugging parallel applications.

Both software packages mentioned above allow the research ring to be formed at particular times of day. Researchers wishing to use their workstation alone can disconnect from the ring. They can connect to the ring when they desire. There are two primary reasons for connecting:

- if researchers need more computational power than their personal workstations can handle, they can connect to the ring to ship some of their programs to idle workstations on campus. At night, for example, computers in instructional labs can thus be easily used for production runs.

- if researchers anticipate that their personal workstations will not be fully utilized for some time, they can donate some of their cycles to colleagues.

"Foreign" programs run on a workstation are always "niced," so that researchers always have first call on their workstations.

Many other institutions have already networked workstations in this manner. The novelty here, is that the workstations do not all reside in one machine room and are not centrally owned; they belong to researchers or instructional labs dispersed across campus. Obviously, there are various management and security issues that come along with the research ring. We describe them later in the paper.

Sharing of Resources

The Campus Ring and the Research Ring will be connected to allow researchers to have full usage of the Campus Ring. We also anticipate that the Enterprise Servers will not be used at full capacity all the time. Researchers will be able to use the Enterprise Servers, especially at night, for the execution of their programs.

CNS Workstation Farm

CNS is also building a "farm" of workstations to be located in CNS machine room. The farm will run the same tools as the Research Ring, and more. Researchers will be able to off-load their programs to the farm as well.

Open System Environment

The Campus Ring and the Research Ring will be managed under the auspices of the Open System Environment Project. This project is dealing with:

- authentication services to ensure data security, such as Kerberos
- a user ID service that uniquely identifies campus users, such as UNIQUENAME
- a campus-wide network file system, such as AFS
- monitoring tools to measure computing resource usage in a distributed open system environment
- system administration guidelines for managing distributed resources in a collaborative manner with faculties and departments

The Demise of Mainframes

The computing resources on the Campus Ring and the Research Ring will slowly replace our mainframes. We anticipate that by 1998 CNS will not host any mainframes.

The Users' View of an Integrated Electronic Environment

Many devices and utilities are already attached to the network backbone and their number will grow over time. It is essential that we present to the user an easy interface to the network environment, hiding from the user the complexity of the operation. The goal is to provide to users (students, faculty, staff, and the community) access to an integrated electronic environment from any workstation on campus, be it located on the user's desk, in an instructional lab, or in the library. Access will be provided to academic data (library, special data bases), to institutional data (student records, classes, forms, etc.), and to special devices (workstations, printers, publishing devices, high performance computers). These resources might be situated on campus, directly attached to the backbone, or they might be somewhere on the Internet.

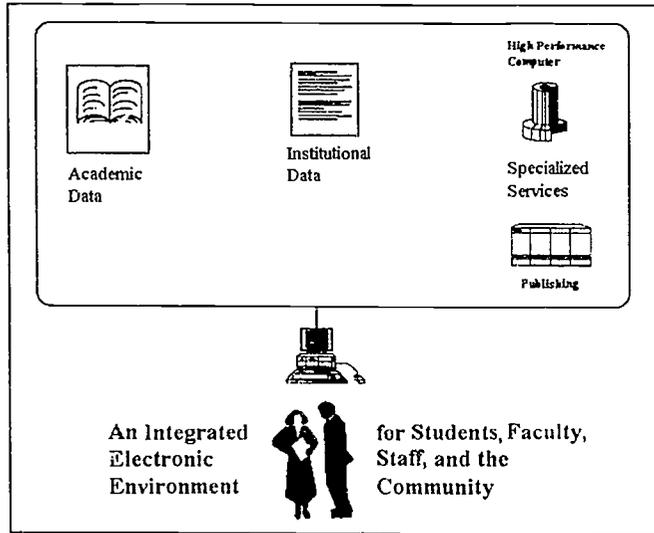


Figure 6: An Integrated Electronic Environment

Budgetary Gymnastics to Move from Mainframes to a Distributed Architecture

The shared distributed computing environment we want to establish will be much cheaper than the current mainframe environment (see first graph in Figure 7). However, while mainframe maintenance constitutes a large drain on our budget, mainframes are also responsible for very substantial revenues for CNS (see second graph in Figure 7).

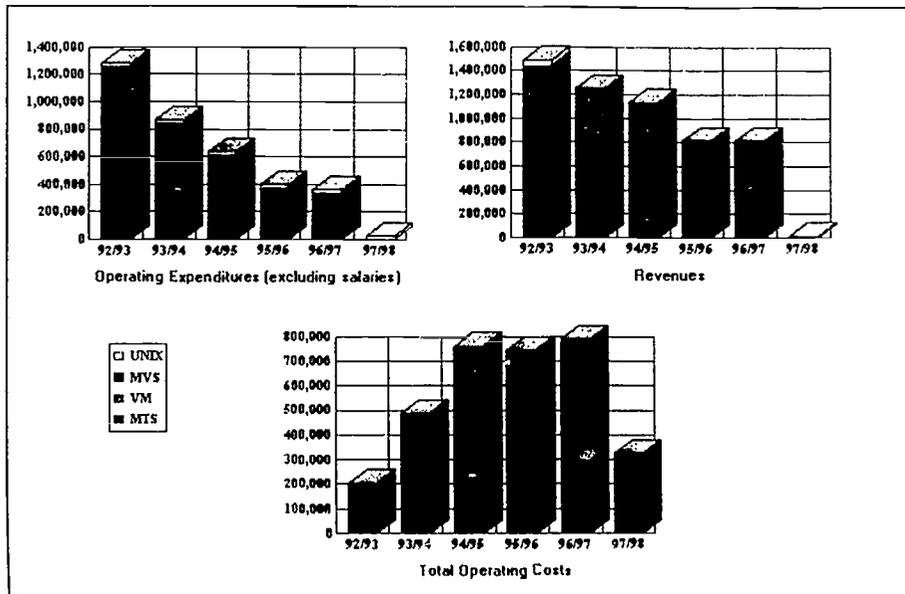


Figure 7: Projected Expenditures, Revenues, and Costs during the Transition Period from Central to Distributed Computing (Costs are Expenditures minus Revenue)

As the mainframes disappear, the revenue will disappear with them. We have agreed with senior University administration that we would not seek other sources of revenue, once the mainframes are decommissioned. It was felt that internal University accounting is expensive, and that computing should eventually be regarded as a utility.

We all believe this is a very desirable environment to establish for users, but it leaves us with an interesting budgetary challenge:

- Users will move off mainframes faster than we can transfer mainframe applications (such as administrative applications) to a distributed client/server environment. Therefore, as users move off mainframes, the decline in revenues will be faster than the decline in mainframe maintenance costs. The result is a net loss in CNS revenue.
- During the transition period from central to distributed computing, we will have to maintain both types of systems. Our expenses will therefore temporarily grow.

The last graph in Figure 7 shows the overall projected operating costs of moving from mainframes to the shared distributed computing environment. We are currently devising with senior University administration various methods of how to bridge the transition period. We are extremely fortunate to have a senior University administration that understands the issues, and is supportive of initiatives (such as ours) aimed at improving University operations.

Supporting Human Infrastructures - Managing an Extended Enterprise

Implementing and supporting a distributed computing environment comes with various managerial challenges. We believe that these challenges far outweigh the technical challenges. In the past, central computing organizations dictated computing solutions to the campus. Today, both computing resources and computing expertise have spread on campus and management of a computing infrastructure has become an endeavor to be shared with the users.

Advisory Committees

There are various tasks which we are trying to share and coordinate with user input and participation. Besides the University Computing Advisory Group (UCAG) which reports to the V.P., Student and Academic Services, and which has existed for many years, we have established other advisory committees. Officially these committees report to the same V.P., but it is CNS that is managing them. All these committees have representatives from faculty and administration. Some of the committees also have student representatives. The committees help gather user requests, provide a sounding board for CNS plans, help establish policies and procedures, and help us coordinate campus activities related to information technology:

- *Network Advisory Committee (NAC)*: The committee was used as a sounding board for the FDDI network backbone architecture. Current challenges comprise educating the campus on how to plan for and connect departmental LANs to the network; managing a presidential fund to cover the cost of connecting departmental LANs to the network; and teaching users on Internet access and potentials.
- *Numerically Intensive Computing Committee (NICC)*: The committee advises on directions for high-performance research computing on and off campus. The challenge ahead is management of the Research Ring. A special user committee will be set up to establish policies and procedures and to coordinate collaboration. Other challenges relate to connecting the Research Ring to high-performance computers in Alberta and elsewhere in North America.
- *Technology and Standards Committee (TASC)*: The committee is used as a sounding board for the information technology directions formulated by CNS. It is also being used to help determine which information technology tools are recommended to the campus (and thus supported by CNS).
- *Information Systems Advisory Committee (ISAC)*: The committee advises on information systems directions and plans. The committee will be pivotal to help us migrate institutional data from the MVS mainframe to Oracle running on the enterprise servers and client workstations distributed across campus. Tasks ahead of us for which we will seek advice from the committee are:
 - establishment of a training and support program for customers to learn about the new Oracle tools
 - establishment of criteria to select some pilot projects to test the new Oracle tools
 - establishment of University-wide criteria for selection of industry applications to run with Oracle, such as financial or student record systems.
 - CNS will define a data dictionary and the data management infrastructure, deciding, for example, which data is to reside on the enterprise servers and which data is to be stored within departmental LANs.

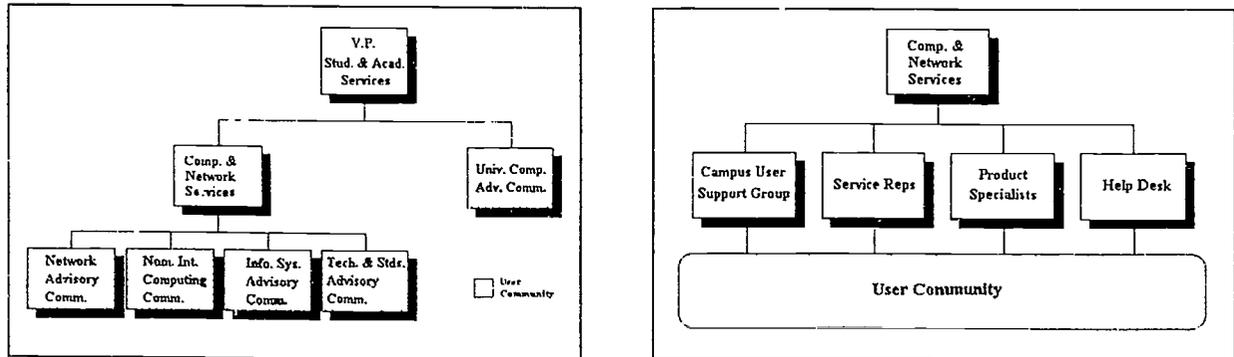


Figure 8: Managing an Extended Enterprise: Advisory Committees (left) and CNS Outreach Program (right)

CNS Outreach Program

Besides advisory groups, CNS has established additional outreach programs:

- *Campus Computing Support Group*: Departments with large computer needs have appointed local computing support staff. The primary task of these staff is to manage the departmental LANs. In the past, this departmental computing support staff has worked in isolation. CNS is now starting a program to establish an extended user support enterprise on campus. By exchanging information, coordinating training and user documentation, CNS will collaborate with the departmental support staff to provide computing support on campus.
- *Service Representatives*: CNS has also appointed staff to help faculties and departments plan for computing and networking. The task of a Service Rep, a CNS staff member assigned to a faculty, is to understand the way computers are used in that unit, guide users in high-level plans, disseminate information on CNS directions, and keep CNS apprised of issues and user wishes.
- *Product Support Specialists*: While Service Representatives are assigned to faculties, Product Specialists help users with the use of specific products. Since different faculties may in many cases use the same products, Product Specialists transcend faculty and departmental boundaries. Product Specialists keep track of departmental expertise and experience and establish a network among users with similar technology interests.
- *Help Desk*: The Help Desk collaborates closely with the Campus Computing Support Group, the Service Representatives, and the Product Specialists to establish and maintain a network of users with various expertise and responsibilities in information technology.

CNS is still learning how to manage these outreach programs. Initial feedback from the user community is good, but we feel that there is still room for improving how we capitalize on the information we gather from the user community and how we use it to constantly improve services.

Technology Planning for the Nineties: Responding to the Challenge to Do More Work With Fewer Resources.

**Ilee Rhimes and Lawrence R. Kelley
Kent State University
Kent, Ohio**

Abstract:

Kent, like most universities, is being challenged to do more work with fewer resources. Information technology is expected to play a more important and strategic role as the University responds to this challenge. Kent developed and implemented a plan for information technology that builds on current strengths, promotes broad commitment and ownership, and helps ensure that scarce technology resources are allocated to strategic activities.

Using IBM's Information Planning Study (IPS) Methodology, Kent organized a nine person Study Team and charged them with developing an information technology plan.

The objective of this paper is to share the results of Kent's study.

Background:

Kent State University, located in northeastern Ohio, is a state-assisted institution with eight campuses, i.e., the Kent Campus and seven Regional Campuses. Kent, was established as the Kent Normal School in 1910, renamed Kent State College in 1929 and became Kent State University in 1935. The University, which is the 28th largest in America, serves approximately 33,000 students with 23,000 on the Kent Campus. The Kent Campus, located in Kent, is 35 miles southeast of Cleveland, and 11 miles east of Akron. The Regional Campuses names and locations are Ashtabula (in Ashtabula), East Liverpool (in East Liverpool), Geauga (in Burton Township), Salem (in Salem), Stark (in Canton), Trumbull (in Warren), and Tuscarawas (in New Philadelphia). Kent's total full-time and part-time faculty, staff and student employees is approximately 6,000. Kent's total annual budget is approximately one quarter billion dollars.

Undergraduate students can pursue studies in over 135 majors and 111 concentrations. Graduate students can pursue 13 different graduate-level degrees in 118 masters majors and 62 doctoral majors. As a member of the Northeastern Ohio Universities College of Medicine, (a consortium of Kent State University, the University of Akron, and Youngstown State University), Kent offers students the opportunity to earn combined bachelor of science and doctorate of medicine degrees in six years. Kent State University's academic programs in Liquid Crystals, Fashion Design, Psychology, and Education among others are known internationally for their quality. The Kent State University Library has been designated as an official depository of U.S. documents - one of only four such libraries in Ohio.

Kent State University encompasses 2,300 acres, including 824 at the Kent Campus and 663 for the Regional Campuses. Other acreage includes an airport, golf course and numerous conservation areas. The physical plant on the Kent Campus includes 102 buildings with 30 residence halls and 230 family-housing apartments capable of housing 7,000 students.

Kent State University is organized into the following divisions: Academic and Student Affairs; Business and Finance; Institutional Advancement; and Human Resources. The Information Services Department, which supports both academic and administrative computing needs, is in the Division of Business and Finance. The Director of Information Services reports to the Associate Vice President for Business and Finance. Information Services is organized into five major functional areas: Administrative and Clerical Support; Application Services; Technical Support and Operations; Network Services; and Computer Equipment Services. Information Services has 69 full-time staff members and approximately 32 student staff members.

The central computing facility consists of an IBM 3090-200S mainframe that runs under the MVS/ESA operating system, and an IBM 4381-R24 mainframe that runs under the VM operating system. The two systems share 134 gigabytes of random access disk storage. Other peripherals include IBM 7171 protocol converters, IBM communication controllers, fiber optic channel extenders, local and remote laser and line printers, cartridge tape drives, and an automatic tape library robotics system. The IBM 3090 is used primarily for administrative and library support while the IBM 4381-R24 is used primarily for academic support.

Complimenting the Computer Center are numerous resources in various other offices. They include network-based and free-standing work stations that support day-to-day functions such as word processing, spreadsheet, record keeping, local data processing, etc. Also, there are over 40 local area networks. Academic departments provide students and faculty with computing support in shared "public" and restricted-access facilities in over 23 separate buildings and 50 different sites. Computing resources in support of academic programs consist of a variety of equipment ranging from PC to RISC networks in the departments of Mathematics and Computer Science, Physics, Geology, Architecture, Technology, and the Scientific Computing Laboratory in the Liquid Crystals Institute. Direct access to computing resources is made via both fiber and coax trunk lines with unshielded twisted pair and coax connections to individual work stations and terminals.

Rationale:

Kent, like most universities, is being challenged to do more work with fewer resources. Along with responding to this challenge, Kent State University also must continue to improve the quality of service and support to its diverse group of stakeholders in order to remain competitive with peer institutions. In her state of the University address, delivered in November 1992, President Carol A. Cartwright stated:

"We're at a point in higher education when we can either create a few temporary safety nets to bring us through this crisis, or we can seize the moment to bring about meaningful and fundamental change. Quite frankly, I'm eager to get on with the next phase of development and provide leadership for these changes..."

Computing and information technology is expected to play an important and strategic role as Kent State University responds to this challenge by developing and implementing new and creative approaches to sustain academic programs and support operations.

Over the years, Kent State University has made a significant investment in computing and information technology to sustain its academic programs and support operations. As a result, Kent students, faculty and staff have become increasingly active in the use of technology to fulfill their scholarship, research and support needs. However, like most universities, Kent is being challenged in its effort to maintain a sufficient level of resources to meet demands. Furthermore, this situation is exacerbated by the kind of budget constraints that colleges and universities currently are experiencing all over the country.

Thus, Kent is faced with the need to expand its use of technology at the same time that it faces budget constraints. Kent concluded that it needed to develop, implement and maintain a strategic computing and information technology plan - with input and consensus from its primary stakeholders. The purpose of this plan would be to provide a blueprint for allocating scarce technological resources to the "right things", i.e., allotting them to activities that are high priority (from an institutional perspective) and consistent with the University's mission. This

analysis and planning effort would build on previous plans; and define computing and technology issues, strengths, recommendations, benefits and implementation priorities.

IBM offered to assist Kent with this planning effort by providing the use of its Information Planning Study (IPS) methodology as well as a consultant/facilitator to assist a small working committee with completing this project. IPS (formerly referred to as Application Transfer Study or ATS) is an IBM developed methodology that was first developed in 1978 to assist customers with reviewing their computing and information technology operations and preparing strategic plans. IBM was one of many vendors that Kent established a closer association with and received support from during this period.

The primary goal of the IPS methodology is to facilitate the development of a document that represents a consensus on the direction and priorities for computing and information technology. In addition, the process of going through the IPS methodology encourages involvement and input from a variety of students, faculty and staff, and promotes broad commitment and ownership for plans and recommendations.

In a preliminary meeting to discuss the IPS methodology, IBM cautioned against making the scope of the study too broad and recommended that it be conducted in two parts. Based on recent feedback from students, faculty and staff, the primary areas of concern were campus-wide communication networking and information systems support. Thus, Kent decided to focus the initial IPS study on Network and Information Systems. Also, it was decided that the second part of the study would focus on the needs of Academic Computing and Instructional Technology.

The Study Team, appointed and charged by the President, Vice President for Business and Finance and Provost consisted of nine individuals representing the following areas: Provost's Office, Regional Campuses, Student Affairs, Library, Faculty, Financial Affairs, Information Services, Business Services, and Business and Finance. The Associate Vice President for Business and Finance was appointed chair of the committee because computing and information technology is one of his areas of responsibility.

Goal and Charge:

The nine-person Study Team was asked to build on previous studies and conduct a University-wide analysis of communication network and information systems with:

- the goal of developing a plan for establishing an integrated infrastructure to facilitate the access and use of computing and information technology by students, faculty and staff.

The charge to the Study Team was:

- to assess network and information system needs, in part, by interviewing and gathering information from Kent students, faculty and staff from all divisions of the University;

- to develop a plan that defines communication network and institutional information systems requirements, objectives, issues, recommendations and implementation priorities.

Methodology:

Using the IBM IPS methodology, the Study Team completed the following major tasks:

- **Planning and Team Orientation.** A one-day planning meeting was held in September, 1992 to brief senior management and the Study Team on the methodology, finalize the project work plan, begin developing the interview schedule, and prepare a draft of the instruments that would be used to survey students, faculty and staff. Also, during the orientation session, it was emphasized that technology would be used during every step of the IPS process to facilitate the capture and evaluation of data.
- **Interviews and Questionnaires.** After finalizing the interview schedule and survey instruments, the Study Team solicited information from students, faculty and staff concerning their experience and opinions about current information systems and network facilities.

Twenty interview sessions were held (including one open session) and approximately 100 individuals were interviewed. Interview comments and notes were captured after each session via a lap-top computer. Also, some interviewees submitted written reports.

Over 3,000 questionnaires were mailed to all full-time faculty and staff and a random sample to 300 students. In addition, questionnaires were circulated to students utilizing the computer laboratory sponsored by undergraduate student senate located in the student center, and more were circulated by leaders of the undergraduate and graduate student governments. A total of 905 completed questionnaires were returned with 285 from faculty, 76 from students, 294 from classified staff and 249 from administrative staff. The questionnaire results were entered into a personal computer and analyzed.

- **Identification of Problems and Strengths.** The interview sessions resulted in over eighty pages of notes summarizing the discussions. The Study Team analyzed the information received in the interviews, in the reports prepared by interview participants and in the questionnaires. The Team then prepared strength and problem statements. Note that strength statements were prepared to provide balance to the overall document and to ensure that strengths were considered in the recommendation development process. A lap-top computer linked to an over-head projector was used by the Team to expedite the process of synthesizing data and developing problem and strength statements.
- **Description of Current Environment.** Concurrent with the identification of problem and strength statements, a chapter outlining the current environment was developed. This chapter was prepared by Study Team members representing the Library, Information Services, Business Services and the Regional Campuses.

- **Study Team Education.** The Study Team also was updated on computing and information technology trends and directions in the University environment. Study Team education included presentations on administrative systems, enrollment management, network communications and local area network-based electronic mail. IBM sponsored and coordinated this education by bringing in several of its business partners with expertise in the above mentioned areas to conduct presentations.
- **Solutions, Alternatives and Recommendations.** In preparation for developing recommendations, the Study Team reviewed problem statements, supporting interview documentation, and questionnaires. In addition, they considered the strengths of the current computing and information technology environment at Kent State University as well as current trends in technology. Furthermore, the Study Team reviewed the results of a previous study on information systems and a recently completed University-wide Network Communications Plan. After the Study Team completed this analysis process, it deliberated over alternative solutions and prepared a series of forty-seven recommendations.
- **Benefits.** During this step, the Study Team identified the potential numerous benefits that would be realized if the recommendations were implemented. The determination of potential benefits was based on interview comments, questionnaire comments, and individual opinions. Benefits were categorized by student, faculty, staff and institution.
- **Implementation Plan.** The final step in the Study Team's deliberations was the preparation of an action-oriented, five-year implementation plan.

Key Recommendations:

Following is a summary of the top fifteen recommendations:

- **Mission.** The University should develop a statement that reflects the strategic importance of information technology. Furthermore, this statement should articulate the role that information technology plays in support of the achievement of the University's mission and strategic plans for the future.
- **Peer Institutions.** The University should identify a peer group of institutions which would serve as a benchmark against which comparisons could be made on information technology and network communications development.
- **Technology Policy Advisory Committee.** A broad-based committee should be established to advise the Vice President for Business and Finance on policy matters related to information technology and network planning, standards and priorities. The committee should include representatives from faculty and staff.

- **University-wide Planning Process.** A University-wide plan for information technology and network communications should be maintained and updated on a biennial basis. Therefore, plans for all units (including Regional Campuses) should be shared biannually and integrated with University-wide plans.
- **Data Access/Authorization.** The University needs to develop and implement a more flexible and uniform approach to authorizing data access. In the context of migration toward an integrated administrative data base, common standards for data access are needed. Procedures and policies should be developed which facilitate appropriate levels of authorized access to needed information.
- **University-wide Network.** The University should recognize that a University-wide network is becoming a required utility (like water or heat) for the survival of a modern campus. The network should be ubiquitous, reliable, and provide an adequate capacity to support voice, data and interactive video communications.

The existing network plan should be enhanced to include single-mode fiber in the backbone and from the backbone to the Office of Teleproductions. This enhancement will allow interactive video transmissions from Teleproductions to locations along the backbone. Also, this enhancement will provide the network infrastructure to support the establishment of fully mediated classrooms and distance learning programs.

The network should permit evolution to integrated campus-wide access by students, faculty and staff to University facilities such as the Library and other information resources, computing resources and student information systems. This would include improved access from student laboratories and dormitories, faculty and staff work areas and off-campus locations.

- **Electronic Mail.** Convenient access to an easy-to-use electronic mail system is of strategic importance to the University. This should allow all users to communicate within and beyond the campus in a seamless manner.
- **University-wide Systems Implementation Perspective.** The University should establish an approach to systems implementation that considers both institutional and departmental needs, while at the same time, supporting database distribution and systems integration. Student information should be implemented as one integrated system in order to enhance competitiveness and better support the achievement of enrollment goals.
- **System Data Infrastructure.** The database infrastructure for an integrated information system should be installed. It should consist of a relational database management system with a data dictionary and user-friendly query and reporting capabilities; and integrated and seamless support for distributed computing and document imaging technologies.

- **Need for Training.** The University should provide better overall coordination and systematic strategies to meet existing and future training needs. The University must recognize the added importance of appropriate training as it migrates into a more integrated information systems and network communications environment.
- **Document Handling and Process Flow Evaluation.** Technology alone will not provide the break-through solutions that are envisioned for administrative operations. Therefore, the University should review existing office work flow processes and determine where document handling and duplicating can be streamlined, integrated or possibly eliminated.
- **University-wide Information System.** The University should implement user-friendly access points such as kiosks or touch-tone phones to University information systems. The purpose would be to allow various users to gain direct access to general information about the University and to allow students to access specific information about themselves.
- **On-Going Hardware/Software Resources.** The University needs to plan for on-going budget resources to address the need to support, upgrade, maintain and replace hardware and software. The plan should consider that the appropriate level of technology varies widely from department to department. In addition, the plan should incorporate the migration of equipment from users with high-end work station needs to users with low-end work station needs. Note that there is a base level below which it is not cost effective to retain old equipment because of operating and maintenance costs.
- **Central Computer Upgrade.** The University should plan to upgrade its central computing capacity to support the data base, integration and access requirements of the proposed student information system and network. This upgrade plan should consider replacing the existing "water cooled" mainframe computer system with a more energy efficient "air cooled" system. An analysis of capacity requirements should be performed and used to determine the appropriate processor capacity needed to serve current and future needs (i.e., three to five years). If current technology trends continue, this may be the last mainframe processor that the University will have to acquire because administrative application systems are being re-engineered to run on smaller down-sized and network-based computers. These new application systems are expected to be widely available to universities in the latter part of this decade.
- **Academic Computing Plan.** The University should develop an academic computing plan. With the exception of the network infrastructure, the scope of this study did not include academic computing. Members of the Study Team believe that the full potential benefits of this study will not be realized without an institutional plan for academic computing and an overall instructional and classroom support plan which integrates research and instructional technology with the overall strategic infrastructure proposed here. Further, the plan should include an evaluation of the concept of a student computer usage fee to help fund academic computing requirements on an on-going basis.

Key Anticipated Benefits

The Study Team identified numerous benefits that would be realized as a result of implementing the aforementioned recommendations. Some of the key benefits are as follows:

- A statement that articulates the role that computing and information technology should play in supporting the achievement of the University mission will help provide consistency and direction.
- Identifying institutions that the University considers peers or competitors will allow Kent to identify the level of services that it must maintain to remain competitive in the recruitment and retention of students, faculty and staff.
- Establishing a Technology Policy Advisory Committee will help ensure broad-based participation in computing and information technology plans, priorities and standards.
- University computer equipment will be able to be connected to an integrated University network. The network will provide University-wide access to various services such as E-mail (including BITNET and INTERNET), mainframe databases and services, library databases and services, selected departmental systems and others.
- Electronic interaction and collaboration among students, faculty and staff across the Kent Campus, the seven Regional Campuses and other universities will be significantly improved.
- The network infrastructure will provide faculty with the means of developing mediated classrooms which will enhance the University's image with students, as well as, prospective faculty, employers and funding agencies.
- The student recruitment process will be improved because of better capabilities for identifying, contacting and responding to potential student requests with timely and accurate information; thus, enhancing the yield rate and maximizing enrollment.
- Students will be able to access their own financial aid, academic record, current class schedule and billing information as well as other general information such as directories, etc. This will reduce the time that students currently wait in line to get general information and free up staff to work with students who have more complex issues.
- Student retention may be enhanced because of better access to computing resources and the availability of more integrated information system support for advising and other services.
- There will be a significantly improved technological environment for supporting administrative services, student support services, instruction and research.

- The integration of information systems will provide data which will be both easily and widely accessible. Similar reports containing conflicting information will be reduced. Flexibility will be enhanced as well as the capacity to expand to meet future needs.
- A more effective training program will help ensure that existing hardware/software capacity is better utilized by faculty and staff. More effective and accessible training also should increase productivity, computing skills and job satisfaction.

Implementation Plan:

The final step in the Study Team's deliberations was the preparation of a draft of an action-oriented implementation plan. Each action item in the plan was analyzed in the context of the recommendations and the benefits to be derived. The plan included an initial schedule for completing each item. An individual or a team was assigned the responsibility for developing the necessary tactical plan to complete each item. The Associate Vice President for Business and Finance was assigned the responsibility for monitoring the plan implementation. Examples of key action items currently in progress include: the academic computing study; the integrated student information system; and the upgrade and expansion of the University-wide network.

Conclusion:

The principal content of this study was the result of the Study Team's analysis of interview comments, questionnaires, and written statements provided by members of the Kent State University community. The process of going through the study promoted involvement and input from a variety of students, faculty and staff, and resulted in a broad commitment and ownership for plans and recommendations.

The completion of the University-wide network was the number one priority because it was consistently raised as the top priority issue by the University community in interviews and questionnaires. The need for an integrated student information system was a close second.

The Study Team concluded that technology alone would not provide the break-through solutions that are envisioned for administrative operations. Existing office work flow processes and document handling procedures also need to be reviewed and re-engineered (where necessary) to make more effective use of new technologies.

The recommendations in this study provide a strategy and blueprint that Kent State University plans to utilize as it develops and implements the integrated systems and network infrastructure that it needs to take full advantage of new technologies in the nineties and beyond. With this infrastructure in place, the University will be in a better position to enhance the quality and productivity of its instruction, research and support programs and, therefore, to provide better service to its stakeholders as well as remain competitive with peer institutions in the recruitment and retention of students, faculty and staff.

LOTS OF DATA! NO INFORMATION!
(Why Universities and Colleges Do Not Take Full
Advantage of Their Information Systems)

Bethany M. Baxter
IBM Academic Consulting and Services
Louisville
Kentucky

Over the past five years, IBM Academic Consulting and Services Group has worked with over 200 colleges and universities to establish strategic plans for academic and administrative information systems. From these studies, a common set of problems has emerged.

Most schools originally automated paper systems without making the necessary changes to take full advantage of their investment. The primary reason that these changes did not and do not occur is an unresolved power/authority struggle. Most often, the authority to automate is delegated to director(s) of information systems, yet power to change the approach to tasks resides with administrators whose knowledge of technology varies. Furthermore, many essential changes require collaborative efforts across work groups unaccustomed to collaborating.

This paper will examine three key changes which need to occur for schools to fully utilize their information systems. These include changes in: organizational structure; resource allocation; and approaches to academic and administrative tasks.

LOTS OF DATA! NO INFORMATION!
**(Why Universities and Colleges Do Not Take Full Advantage of
 Their Information Systems)**

During the past five years, IBM's Academic Consulting and Services Group has worked with over 200 colleges and universities to establish strategic plans for information systems. Interviews have been conducted with over 20,000 faculty, staff, administrators and students, and over 100,000 faculty, staff, administrators, and students have responded to questionnaires. As a result of these studies a common set of problems has been identified across all types of schools, regardless of size, orientation, and mission.

The key to these problems is the fact that initially most schools simply automated their paper systems without making appropriate changes to take full advantage of the investment they made in information systems. In fact, if schools purchased administrative systems they frequently modified these newly-acquired information systems in order to match their paper systems, thereby reducing the efficiency already present in the purchased systems.

Our studies have indicated that the primary reason appropriate changes were not and are not made is an unresolved power/authority struggle. On most campuses, the authority to implement information systems has been clearly delegated to the director(s) of those systems. However, the power to change the approach to tasks has remained with the departmental administrators whose knowledge of information systems varies from campus to campus. In fact, the experience of the Consulting and Services Group suggests that administrators' attitudes toward information systems are a far better indicator of the efficiencies to be gained by the implementation of such systems than are the skill sets of the computing center directors. Furthermore, many of the changes that need to occur require collaborative efforts across work groups that are not used to collaborating.

There are three key areas in which changes need to occur. These are changes in organizational structure, changes in the allocation of resources, and changes in approaches to tasks, including both the administrative tasks of handling data and the academic tasks of teaching, research/scholarly activity, and service. This paper will examine each of these areas.

Organizational Structures

Even though automation has significantly changed job responsibilities, on many campuses reporting structures have remained unchanged. As a result, multiple officers frequently manage information systems at various levels, and multiple

offices may be charged with similar tasks. For example, it is not uncommon to find a satellite dish funded by an engineering college, the cabling from it the responsibility of Telecommunications, the storage and delivery of taped material under the auspices of the library, the facility for uplinking programs through the satellite in Continuing or Distance Education and the development of video materials in the Drama Department. Multiple groups may be wiring campuses for academic networks, administrative networks, television, security, telephone, Electronic Data Interchange (EDI) and facilities management.

Similarly, the committees which are in place to accomplish collaborative tasks may not be designed or charged to address the concerns of today's technological world. Such committees include long-standing committees and more recently appointed technology advisory committees. For example, the Director of Telecommunications may not be a permanent members of the Physical Facilities Committee, and the Director of Information Systems may not be a member of the Strategic Planning Committee. With increasing frequency, changes on campuses require information resources. Unless directors of systems are included in the initial stages of planning, both the costs of the technology required and the efficiencies to be gained by using technology are not identified. Furthermore, some strategic planning time cycles do not identify the service orientation of information systems providers. It is not uncommon for such offices to be required to submit yearly plans at the same time as the plans of other departments on campus, before they know what will be required of them.

On the other hand, newly appointed technology advisory groups may not be appropriately representative of the applications or the users on campus. Too often the criteria used to select members of these committees is technological knowledge and skill. This approach often results in advisory committees who advise on technology only. Such committees often become deadlocked over issues such as "UNIX versus VM or VMS" or "fiber versus ATM." Meanwhile, other important issues, such as the prioritization of projects, standardization of tools, and cross-discipline sharing of resources, remain unaddressed. Users in these environments have no identifiable representation and may become frustrated. Frustrated users may attempt to meet their own needs by installing their own "shadow" systems. Ultimately, this approach may result in the degradation of the integrity of the data on the entire campus as individual units begin to rely on their own databases with a resulting disregard for the accuracy of the central system.

Finally, there are committees that are not established on many campuses. For example, the current trend for faculty, staff and students to use their campus identification cards as Smart Cards, Debit Cards and for electronic building access, automated

library check out, and other functions has resulted in many groups planning for the use of these cards with little sharing of information between the groups involved.

Resource Allocation

At this time, as funding is being reduced for many colleges and universities, improving the efficiency of operations is critical. In many areas a great deal of efficiency can be gained by having well trained personnel use appropriate technology. However, the criteria for the distribution of resources and the training of personnel are often based on premises that are no longer relevant. This is particularly true in state universities, although this problem is prevalent on many other campuses.

On many campuses, requirements for positions and criteria for promotion and salary increases do not reflect the skills needed in today's world. In many institutions, job descriptions of support staff still require typing rather than word processing. Similarly, programmers are often only rewarded for longevity or for promotion to management positions rather than for the acquisition of needed new skills.

Regulations regarding the acquisition of equipment are similarly flawed. In many states and on many campuses, capital and operational expenses remain separate, and the ability to purchase equipment to offset the need to hire additional staff or to replace staff is very difficult. In fact, the purchase of equipment is often more easily linked to the renovation or construction of a building than to efficiencies that could be gained in operations. Furthermore, many procurement manuals are unclear about which purchasing categories address certain new elements of technology, such as network design, and purchasing departments may make determinations that complicate the procurement process.

There are also many areas which are not addressed in current financial structures. Equipment procurement is usually seen as cyclical rather than ongoing. Therefore, there are frequently no funds available for maintenance and upgrade of acquired equipment. Similarly, the predetermined length of time that equipment must be kept before it can be replaced may not reflect the rapid advances of technology.

Even without cumbersome regulations, many universities and colleges have separated hardware and software in the procurement process. In effect, they have separated applications from equipment. It is not uncommon to find universities where hardware is provided by the department and software is provided by the university or vice versa. Such separation makes the process of prioritizing less meaningful and precludes the sharing of resources across functional groups. This approach may result

in unexpected demands on central resources or the proliferation of underutilized departmental equipment.

Underlying the formal procurement policies, most members of institutions have a preconceived notion, based on experience, of how much resource a given area can request at a given time, whether the request is for programmer's time or for equipment. Therefore, many requests are curtailed to meet the perceived "appropriate" profile. Equipment may be purchased, even though the purchaser knows that it has a short life span, because the price falls within the perceived resource niche. Similarly, requests for services are frequently presented piecemeal so that the requester does not appear to be asking for more resource than (s)he should at a given time. This makes it difficult to understand the real needs of the campus and to address them efficiently.

Approaches to Tasks

Reorganization, new procurement policies, and well trained personnel will not produce the efficiency desired unless universities and colleges change the way they operate. They must look at the way they administer, teach, produce research and provide service to the community in the light of new technology available. Instead, as areas are automated, strong emphasis is often placed on automating all processes exactly as they are carried out manually.

The most common defense against changing is, "We are unique!" Homegrown software has been shown to be as much as four times as expensive as purchasing off-the-shelf administrative packages, but many universities believe that their uniqueness demands this approach. These systems are often designed as separate files that matched the file cabinets in other times. The collectors of data are still seen, not only as the owners of the data, but also as the determiners of what should be collected, how data is stored, who has the ability to access the data, and in what form access may occur. Such narrow definitions may severely limit the ability of faculty, staff and administrators to use data effectively.

As industry moves to business process reengineering, it is imperative that universities also look at the processes they use to accomplish their administrative tasks, identify steps that can be eliminated, and automate steps that can be automated. However, the leadership in many areas does not have the necessary understanding of what technology can accomplish to utilize it well.

The responses to a survey used by the consulting group in a broad range of colleges and universities consistently show that 60% to 80% of the faculty and 30% to 34% of the support staff have computers at home. However, administrators with computers in their homes vary from 20% to 80% on any given campus. Our

experience has shown that this particular statistic is a good indicator of the approaches to change that are occurring in administrative areas. The presence of a computer in the home seems to lift it from the status of "glorified typewriter" to tool, regardless of who in the home uses it. Where a computer is seen as a tool, administrators try to identify areas where it can be useful and provide leadership for the reengineering process. On the other hand, if the computer is seen as a "glorified typewriter," there are concerns about who should have one and who should be trained. "Will middle management be doing secretarial tasks if they have a computer?" is a frequently asked question in such environments.

The need for information from administrative systems is a frequent concern of those in academic areas. They particularly desire information that will assist them in advising students and managing budgets. Yet, in too many institutions, pertinent administrative information is not shared because of security concerns or interpretations of the Buckley Amendment. As a result, students may be placed in failure situations, spend extra time and dollars to receive a degree, or have their health needs unidentified in an emergency. Deans and chairs may struggle to manage their resources with limited information. Ultimately, such attitudes may further increase the development of "shadow" systems and further erode the integrity of the institutional data.

Using technology for teaching is too often seen as training students to use computers. To accomplish this, most colleges and universities have developed some "computer literacy" criteria that a student must complete before graduation. The relevance of this criteria to the needs of the workplace or to the needs of the educational institution is often not regularly reevaluated. There is also little recognition in these requirements of the experiences many traditional students are now having in K-12 or that nontraditional students have had in the workplace. Requiring literacy requirements to be fulfilled before graduation means that faculty cannot assume that students have any skills during their time on campus. This keeps many faculty from utilizing many beneficial educational tools.

Support for researchers also varies from campus to campus. Campuses that centrally provide information about funding that is available, identify research interests of other faculty, provide access to appropriate databases, and provide a broad variety of statistical tools have often seen an increase in scarce research dollars. However, in many schools, support for research still occurs on a departmental or individual basis.

Conclusion

In conclusion, although the actions needed to solve these problems vary from campus to campus, most of the solutions must

be based on collaborative efforts of those who do not have a history of collaborating. Such efforts must include the collaboration of collectors of data with each other and the users, the collaboration of faculty to determine curriculum that meets the needs of the student as a learner today and as a member of the work force tomorrow, the collaboration of researchers to identify and share scarce resources, and the collaboration of the providers of information systems to provide the most cost effective approach to technology.

Leadership and the Changing Role of the Chief Information Officer in Higher Education

Gary M. Pitkin
Dean of University Libraries
University of Northern Colorado
Greeley, Colorado

This presentation reports on the results of a questionnaire, with an 83 % return rate, designed to obtain information on the leadership role of the Chief Information Officer in higher education. The research instrument facilitated the identification of role, perception of role as compared to the business sector, role in the transformation process, comparison of demographic characteristics and institutional practices to previous studies, and the identification of differences in the leadership role of the CIO according to type of higher education institution.

During the Summer of 1992, a questionnaire designed to obtain information on the leadership role of the Chief Information Officer in higher education was sent to 200 individuals identified as serving in that capacity. The return rate was 83 percent. The information provided by the respondents was analyzed descriptively and employed to address the following research questions:

- 1) What is the role of the CIO in higher education?
- 2) What do CIOs in higher education perceive their roles to be as compared to roles identified for CIOs in the business sector?
- 3) What is the role of the CIO in higher education in the process of transformation?
- 4) Are there differences in the leadership role of the CIO according to type of higher education institution?
- 5) Are the demographic characteristics and institutional practices identified for CIOs in studies conducted by Woodsworth (1991) in 1986 and by Penrod, Dolence, and Douglas (1990) in 1989 continuing today or have they changed consequently affecting the role of the CIO?

In examining the Chief Information Officer position, this presentation emphasizes role in terms of higher education, comparison to the business sector, and the transformation process. In terms of demographics, respondents in the current research affirm the trends cited by Penrod, Dolence, and Douglas (1990) with regard to CIOs in higher education being 45 to 49 years of age, male, and Caucasian, who do not have academic tenure and primarily come from an administrative background. Respondents in the current study refute the remaining trends and institutional practices cited by Woodsworth (1991) and Penrod, Dolence, and Douglas (1990) by reporting that most CIOs in higher education do not refer to themselves as Chief Information Officers, were previously and remain directors instead of vice presidents or associate vice presidents, do not report to the president of their institution, do not have academic rank, and do not have the doctorate as the highest earned degree.

Role of the Chief Information Officer in Higher Education

The literature on the role of the Chief Information Officer emphasized executive officer status for the position. Brumm (1989) reported that 73% of the business sector CIOs who responded to her survey had the title of Vice President or higher. Synnott (1987) stated that, in the business sector, the position was effective only if the incumbent had organization-wide authority for information management and the associated technologies. This argument was supported by Rockart (1988) and Gantz (1985). Weiss (1987), Penrod and Dolence (1988), Sherron (1988), Ryland (1989), Dillman and Hicks (1990), and Temares (1991) transferred this argument to higher education by stating that the

CIO position could be effective only if an executive officer at the institution-wide level.

Executive officer status was defined by these same authors as regularly attending board of trustees/regents meetings, serving as an active participant in those meetings, and being involved in the decision-making process as it affects all aspects of the institution including planning, budgeting, academic issues, human resources, facilities, and curriculum. They concluded that as an executive officer responsible for all information units, the Chief Information Officer should have responsibility for academic computing, administrative computing, copying/reprographic services, data communications, institutional research, library services, mail services, media services, planning, printing, television services, and voice communications. These issues were examined in terms of their applicability to the Chief Information Officers contacted for this study. The responses were compared to those received by Penrod, Dolence, and Douglas (1990).

As shown in Table 1, the majority (53.6%) of the respondents in the current study do not consider themselves to be executive officers of their respective institutions. This contrasts with the respondents of the Penrod, Dolence, and Douglas (1990) research in which the majority (58.6%) were executive officers. Based on these results, Penrod, Dolence, and Douglas (1990) claimed that Chief Information Officers were executive officers and indicated that this trend would continue.

Table 1

Executive Officer Status of Respondents in All Types of Institutions

Executive Officer Status	Penrod, Dolence, and Douglas (1990)		Pitkin	
	n	%	n	%
Yes	34	58.6	71	46.4
No	24	41.4	82	53.6
	<u>N</u> = 58		<u>N</u> = 153	

They also reported that as an executive officer, the Chief Information Officer regularly attended board meetings as an active participant. Tables 2 and 3 provide data on attendance and level of participation. As shown in Table 2, 48.3% of the respondents to the Penrod, Dolence, and Douglas (1990) study attended board meetings. The authors surmised, however, that this would

continue. The results of the current study, however, refute this allegation; the majority (63.4%) do not attend board meetings.

In addition, as shown in Table 3, a higher number of respondents to the Penrod, Dolence, and Douglas (1990) study stated that their reason for attending board meetings was to serve as a resource person (17.3%) as opposed to being an active participant (15.5%). Respondents to the current study stated their primary reason for attendance (42.5%) was to be a resource person, meaning that they attended only when called upon for specific information.

In examining the responses to the current study, the conclusion can be drawn that the majority of responding Chief Information Officers are not executive officers of their institutions. They do not attend board meetings on a regular basis as active participants. When they do attend, they do so as resource persons.

Synnott and Gruber (1981), Gantz (1985), and Synnott (1987) argued that the Chief Information Officer's status as an executive officer is exemplified through involvement in executive decisions at the institution-wide level. Penrod, Dolence, and Douglas (1990) substantiated this argument (see Table 4); the majority (65.5%) of their respondents indicated that they were involved in that decision-making process. Somewhat fewer (50.3%) of the respondents in the current study reported to have that level of responsibility. This difference of over 15% in three years indicates that the executive-level responsibility of the Chief Information Officer may be decreasing.

Table 2

Board Meeting Attendance of Respondents in All Types of Institutions

Attendance on a Regular Basis	Penrod, Dolence, and Douglas (1990)		Pitkin	
	n	%	n	%
No	30	51.7	97	63.4
Yes	28	48.3	56	36.6
	<u>N = 58</u>		<u>N = 153</u>	

Table 3

Purpose of Board Meeting Attendance of Respondents in All Types of Institutions

Purpose of Attendance	Penrod, Dolence, and Douglas (1990)		Pitkin	
	n	%	n	%
Do Not Attend	30	51.8	36	23.5
Resource Person	10	17.3	65	42.5
Participant	9	15.5	37	24.2
Combination	6	10.3	0	0.0
Observer	2	3.4	12	7.8
Did Not Respond	1	1.7	3	2.0
	<u>N</u> = 58		<u>N</u> = 153	

Penrod, Dolence, and Douglas (1990) discussed in their study the Chief Information Officer's role in dealing with information management. Table 5 presents findings for the Penrod, Dolence, and Douglas (1990) study and the current research. In academic computing, 86.2% of the Penrod, Dolence, and Douglas (1990) respondents and 80.4% of the current study respondents reported responsibility. For administrative computing, 89.7% of the Penrod, Dolence, and Douglas (1990) respondents and 91.5% of the current investigation respondents reported responsibility. For data communications, 96.6% of the Penrod, Dolence, and Douglas (1990) respondents and 92.2% of the current study respondents reported responsibility. For voice communications, 69% of the Penrod, Dolence, and Douglas (1990) respondents and 62.7% of the current study respondents reported responsibility.

Table 4

Participation in Executive Decisions of Respondents in All Types of Institutions

Participation on a Regular Basis	Penrod, Dolence, and Douglas (1990)		Pitkin	
	n	%	n	%
Yes	38	65.5	77	50.3
No	20	34.5	76	49.7
	<u>N</u> = 58		<u>N</u> = 153	

Percentages were much less for Penrod, Dolence, and Douglas (1991) and the current research for all remaining areas of information management: copying/reprographic services, institutional research, library services, mail services, media services, planning, printing, and television services. This indicates that the level of responsibility has not changed for the Chief Information Officer across units associated with information management.

In summary, respondents in the current study indicated that the role of the Chief Information Officer in higher education is not consistent with that reported by Penrod, Dolence, and Douglas (1990). Unlike the respondents in that study, the Chief Information Officers in the current research do not consider themselves to be executive officers of their institutions and they do not regularly attend board meetings. Respondents in the current study (50.3%) reported a decrease, as compared to the Penrod, Dolence, and Douglas (1990) respondents (65.5%), in their participation in executive decisions. As reported earlier, however, the respondents in the current study do agree with the respondents in the two previous investigations concerning areas of information management over which they have responsibility.

Table 5

Areas of Responsibility of Respondents in All Types of Institutions

Area of Responsibility	Penrod, Dolence, and Douglas (1990)		Pitkin	
	n	%	n	%
Academic Computing	50	86.2	123	80.4
Administrative Computing	52	89.7	140	91.5
Copying/Reprographic Services	10	17.2	23	15.1
Data Communications	56	96.6	141	92.2
Institutional Research	11	19	34	22.2
Library Services	9	15.5	17	11.1
Mail Services	10	17.2	22	14.4
Media Services	9	25.5	29	19
Planning	19	32.8	29	19
Printing	10	17.2	20	13.1
Television Services	16	27.6	42	27.5
Voice Communications	40	69	96	62.7
	<u>N = 58</u>		<u>N = 153</u>	

Synnott and Gruber's (1981) Roles of the Chief Information Officer in the Business Sector

In writing about the Chief Information Officer concept in the business sector, Synnott and Gruber (1981) stated that "EDP [Electronic Data Processing] managers are . . . taking on new and broader responsibilities" (p. 45). As Chief Information Officers in executive-level positions, they are "responsible for establishing corporate information policies, standards and procedures" through providing "centralized management and control over information processing and utilization, even though it may be distributed geographically and functionally throughout the organization." (pp. 66-68)

Synnott and Gruber (1981) identified nine roles of the Chief Information Officer as essential for the position to be effective as an executive officer at the organization-wide level. These roles are strategic planner, change agent, proactivist, politician, integrator, information controller, staff professional, manager, and futurist. The roles were confirmed in studies of the Chief Information Officer in the business sector by Passino and Severance (1988), Brumm (1989), and Adler and Ferdows (1990). In addition to identifying the prevalence of these role characteristics, these authors further agreed that the CIO is an executive officer with institution-wide decision-making authority.

The research instrument was configured to help identify whether the Chief Information Officer in higher education carries out responsibilities identified by Synnott and Gruber (1981) for the CIO in the business sector. The Chief Information Officers who received the research instrument were asked to indicate the level of Synnott and Gruber's (1981) roles as being incorporated into their responsibilities. Of the nine roles only three are "always" present: manager (74.5%), futurist (61.4%), and strategic planner (54.2%). Synnott and Gruber (1981) claimed that for the Chief Information Officer to be effective as an executive officer involved with the management of the entire organization, all nine roles must operate at the "always" level. The respondents indicated the remaining six roles (change agent, proactivist, politician, integrator, information controller, staff professional) are less than "always" present.

Chief Information Officers in higher education do not consistently carry out two-thirds of the roles identified by Synnott and Gruber (1981) as being necessary to be effective in that position. Since Synnott and Gruber (1981) argued that their roles were a necessary facet of the executive-level role, and of the identification of organization-wide responsibility, the conclusion can be drawn that the Chief Information Officer in higher education is not an executive officer and does not have institution-wide responsibilities. This conclusion is consistent with the responses that indicate that the Chief Information Officer in higher education does not refer to himself or herself as a Chief Information Officer, does not have an executive-level title, does not consider himself or herself to be an executive officer, is not involved in making executive decisions, and does not have responsibility for most units associated with information management (see prior discussion).

In summary, Chief Information Officers in higher education, as represented by respondents in the current study, emulate the CIO in the business sector in that they are strategic planners, managers, and futurists. They do not, however, carry out the remaining six roles identified by Synnott and Gruber (1981) as

necessary to be effective as an executive officer of the organization. The conclusion is consequently drawn that most of the respondents to the current study do not emulate the CIO in the business sector, and, as concluded earlier, do not function as executive officers.

The Transformation Process and the Role of the Chief Information Officer

Several authors (Hopper, 1990; McFarlan, 1990; Penrod & Dolence, 1990, 1991; Peterson, 1990) argued that the Chief Information Officer position in both the business sector and higher education is exemplary of the transformation process. Theorists of organizational structure and change (Kanter, 1983, Morgan, 1989; 1989; Peters, 1987; Waterman, 1987) referred to the period of time during which the Chief Information Officer position was created in the business sector as the era of transformation. The characteristics needed for success in transformation are similar to the roles required of the Chief Information Officer as identified by Synnott and Gruber (1981).

Nolan (1990) and Penrod and Dolence (1991) stated that transformation must take place in higher education institutions. They emphasized, as did Hopper (1990) and Hammer (1990), that information technology as directed by a Chief Information Officer, is the driving force behind the transformation process. They contended that transformation cannot take place without the application of information technologies. Morgan (1989), Kanter (1983, 1989), Peters (1987), and Waterman (1987) contended that the organization must be entrepreneurial, incorporate transformation into its overall organizational goals, and conduct change through the transformation process for positions like the Chief Information Officer to be effective.

The research instrument obtained perceptions about the level of transformation taking place at the institutions represented by the respondents. To obtain information on the level of transformation taking place at institutions of higher education, Chief Information Officers receiving the research instrument were asked to identify the stage of progress which best depicted their institution's movement toward transformation. Five stages, based on the continuum of achievement developed by Rogers (1971), were defined for respondents. These stages are: "awareness," having knowledge of transformation; "interest," actively seeking additional information; "evaluation," deciding whether or not to implement transformation; "trial," applying transformation on a small scale for analysis purposes; and "adoption," applying transformation throughout the institution as a matter of policy.

Respondents were first directed to apply this continuum to define the stage of transformation taking place at their institution. The definition of transformation, "new ways of perceiving, thinking, and behaving by all members of the organization," was based on the definition of corporate transformation compiled by Kilmann and Covin (1988). Respondents' reports indicated that the transformation process is not dominant at any stage along the continuum from "awareness" through "adoption". The responses also indicated a bimodal distribution in that 47% of the respondents are at the early stages of "awareness" and "interest" and that 42% are at the advanced stages of "trial" and "adoption."

Within the definition of transformation, Beckhard (1988) identified four types of activities that are transformational. These are change from a product-driven mission to a market-driven mission; change from a centralized to a decentralized management structure; change from low-technology to high-technology approaches to accomplishing responsibilities; and changes in organizational culture from traditional to creative structures. Respondents to the research instrument indicated at what stage, using Rogers' (1971) continuum, their institutions were with each transformation activity. The responses are similar to those received concerning the stage of achievement for the definition of transformation. No stage, "awareness" through "adoption," reached the majority level (50.0% or higher) for any of the four transformation activities. There was no commonly identified stage along the continuum for the transformation activities concerning mission and management structure. In fact, the range of percentage from the top rated stage to the lowest was only 12.5% for mission and 9.1% for management structure. The responses for those two transformation activities were spread somewhat evenly across the stages of progression.

This is not the case, however, in changes concerning technology and organizational culture. "Adoption," the most advanced level of achievement along the continuum, was the most prevalent stage (48.3%) in the change from low- to high-technology approaches to accomplishing responsibilities. "Awareness," the first stage of transformation, was the most prevalent (40.5%) in the movement from traditional to creative organizational structures. These higher education institutions appear to be at an advanced point in only one area of transformation activity (technology) identified by Beckhard (1988).

Respondents were also asked to indicate which one of six decision-making/communication processes best described his/her institution. The six processes, or models, were originally identified by Morgan (1989). He contended that organizations in the process of transformation were migrating from Model 1 to Model 6. Respondents would consequently provide a good indication of their institution's progress in the transformation process by indicating which model was prevalent.

In the progression from Model 1 toward Model 6, respondents indicated that the common models are the second and third. Manager to subordinates with the flexibility of being authoritarian or participative was identified most often (40.5%). The involvement of project teams with decisions still made at the top of the organization received the second highest number of responses (30.0%). The remaining four models, or processes, were not identified frequently (ranging from 16.3% to 0.0%). The indication here is that the decision-making/communication process still reflects a high level of bureaucracy in that decisions are still made at the top of the organization. This conclusion is consistent with the responses received concerning the definition of transformation and the level of achievement in transformational activities. The institutions represented by the respondents appear not to be actively involved in the process of transformation.

Types of Institutions

An exception to the general conclusion that that the Chief Information Officer is not an executive officer, does not emulate the CIO in the business sector and is not involved in management processes that support the transformation proces appears to be the private undergraduate institutions. Respondents from those institutions have the second highest frequency of reported executive officer status (50.0%), the highest frequency of reported board meeting attendance (62.5%) at the "always" level, the highest frequency of reported attendance at board meetings as an active participant (37.5%), and the highest frequency of reported participation in executive decisions (75%) at the "always" and "often" levels. Within the executive decision-making process, they have the highest frequency of reported participation in every category: planning (87.5%), budgeting (87.5%), academic issues (50.0%), human resources (50.0%), facilities (87.5%), and curriculum issues (25.0%). They also have the highest frequency of reported responsibility in areas associated with information management: copying/reprographic services (37.5%), mail services (37.5%), institutional research (62.5%), planning (50%), and printing services (25.0%).

The private undergraduate respondents come closest to emulating the business sector CIO. They most often reported Synnott and Gruber's (1981) role characteristics at the "always" and "often" levels for change agent (100.0%), proactivist (100.0%), politician (100.0%), integrator (100.0%), information controller (75.0%), and staff professional (75.0%). More respondents on a percentage basis from the private undergraduate institutions (37.5%) reached the fourth level (project teams with leaders involved in final decisions) of Morgan's (1989) models of decision-making/communication processes, than from any other type of institution.

References

- Adler, P. S., & Ferdows, K. (1990). The chief technology officer. California Management Review, 32(3), 55-62.
- Beckhard, R. (1988). The executive management of transformational change. In R. H. Kilmann & T. J. Covin (Eds.), Corporate transformation: Revitalizing organizations for a competitive world (pp. 89-101). San Francisco: Jossey-Bass.
- Brumm, E. K. (1989). Chief information officers in service and industrial organizations. Unpublished manuscript.
- Dillman, H. L., & Hicks, M. A. (1990). Reorganizing for information technology management on campus. Cause/Effect, 13(3), 4-6.
- Gantz, J. (1985, October). Telecommunications management: Who's in charge? Telecommunication Products + Technology, pp. 17-37.
- Hammer, M. (1990). Reengineering work: Don't automate, obliterate. Harvard Business Review, 68(4), 104-112.
- Hopper, M. D. (1990). Rattling SABRE--new ways to compete on information. Harvard Business Review, 68(3), 118-124.
- Kanter, R. M. (1983). The change masters. NY: Simon and Schuster.
- Kanter, R. M. (1989). When giants learn to dance. NY: Simon and Schuster.

- Kilmann, R. H., & Covin, T. J. (1988). Preface. In R. H. Kilmann & T. J. Covin (Eds.), Corporate transformation: Revitalizing organizations for a competitive world (pp. xiii-xviii). San Francisco: Jossey-Bass.
- McFarlan, F. W. (1990). Letter to the editor. Harvard Business Review, 68(4), 176.
- Morgan, G. (1989). Creative organization theory: A resourcebook. Newbury Park, CA: Sage.
- Nolan, R. L. (1990). Too many executives today just don't get it! Cause/Effect, 13(4), 5-11.
- Passino, J. H., & Severance, D. G. (1988). The changing role of the chief information officer. Planning Review, 16(5), 38-42.
- Penrod, J. I., & Dolence, M. G. (1988). IRM a short lived concept? Proceedings of the 1987 CAUSE National Conference, 173-183.
- Penrod, J. I., & Dolence, M. G. (1990). Managing information technology: Facing the issues. Proceedings of the 1989 CAUSE National Conference, 71-81.
- Penrod, J. I., & Dolence, M. G. (1991). Concepts for reengineering in higher education. Cause/Effect, 14(2), 10-17.
- Penrod, J. I., Dolence, M. G., & Douglas, J. V. (1990). The chief information officer in higher education (Professional Paper Series No. 4). Boulder, CO: CAUSE.
- Peters, T. (1987). Thriving on chaos: Handbook for a management revolution. NY: Alfred A. Knopf.
- Peterson, D. J. (1990). Letter to the editor. Harvard Business Review, 68(4), 180-181.
- Rockart, J. F. (1988). The line takes the leadership--IS management in a wired society. Sloan Management Review, 29(4), 57-64.
- Rogers, E. M. (1971). Communication of Innovations: A Cross-cultural Approach (2nd ed.). NY: Free Press.
- Ryland, J. (1989). Organizing and managing information technology in higher education: A historical perspective. In B. L. Hawkins (Ed.), Organizing and managing information resources on campus (pp. 17-32). McKinney, TX: Academic Computing Publications.
- Sherron, G. T. (1988). Organizing to manage information resources. Proceedings of the 1987 CAUSE National Conference, 185-195.
- Synnott, W. R. (1987). The information weapon: Winning customers and markets with technology. NY: John Wiley & Sons.
- Synnott, W. R., & Gruber, W. H. (1981). Information resource management: Opportunities and strategies for the 1980s. NY: John Wiley & Sons.
- Temares, M. L. (1991). Future directions in higher education: A CIO's perspective. Proceedings of the 1990 CAUSE National Conference, 47-57.
- Waterman, R. H. (1987). The renewal factor: How the best get and keep the competitive edge. NY: Bantam.
- Weiss, M. (1987, September/October). Transformers. CIO Magazine, pp. 37-41.
- Woodsworth, A. (1991). Patterns and options for managing information technology on campus. Chicago: American Library Association.

Strategies for Data Management Leadership:

Why, Who, What, and How

Lore A. Balkan, Gerry W. McLaughlin, and Richard D. Howard

The critical issue is not one of tools and systems but involvement in the quality efforts of the business units.¹

Overview

A recent advertising campaign for IBM calls for "new ideas for new challenges". This has three elements of relevant to those of us who work with information. First there is a strong theme of change. *New* is the norm and *change* is the standard. The second theme is that of survival in the face of challenges. The way we do everything is subject to question because our world has changed and promises to change at an even greater rate in the future. This theme acknowledges the ever present threat that if we fail to adjust, we and our services will become obsolete. The third theme is less apparent, but provides the keys to open new doors as we close the old. It is quite simply one of creating new vision from ideas generated at all levels of our organizations. Unlike days gone by when demonstrated knowledge and skill were required before ideas were solicited, today we ask the people at the pulse of every activity to interact and generate innovative ideas. The necessity and the ability to acquire new knowledge and skills is accepted as the rule. The road to success is paved with ideas. These ideas are then molded into plans and resources are reallocated to support training on the new and improved practices and products.

The key challenge for organizations is to make better decisions and provide better support for their stake-holders. Information support must enable: 1) recognition of the need to make a decision, 2) identification of alternatives, 3) a call to act, and 4) the ability to verify and defend the action. The successful use of information to this end depends on an information support structure that assures the quality and availability of relevant data and that restructures the data for the decision makers. Not surprisingly, the old ways of doing information support cannot support our new ways of doing business. The disappearance of the middle management of the down-sized organization, the advent of intelligent devices, and the refinement of strategic management, all are the components of the new challenge. To meet this new challenge, we need new

¹ Radding, Alan, Quality is Job #1. In *Datamation*, October 1, 1992, p. 100.

ideas about providing information support for the entirety of our organizations.

Current Culture

The last several years have seen rapid changes in the technology with which we manage our organizational data. Complementary changes now need to occur in our organizations and in our concepts about information systems. We are witnessing a simultaneous push by technology and pull by management to deliver useful information. It is naive to assume that all of the required management information in fact resides in our historical legacy systems, and that the challenge is to simply implement new technology to deliver data on demand to an expanding clientele. What is too often overlooked is that management's requirements are driven by new needs to perform analyses related to both long-term and short-term decisions and then to evaluate the results of those decisions. The operational systems that perform the day-to-day transactions have not been designed to support this activity and quite likely do not contain sufficiently standardized or normalized data for the integrated, re-combined, and longitudinal views required for analysis.

Existing internal and external reporting functions often use informal procedures to obtain data and to interpret the variables. There is not a predominance of historical files, or census files, or standardized data. Furthermore, there is often no formal assignment of responsibility for data management, nor policies governing the process by which those involved with the operational systems capture, store, define, secure, or provide the data to management or do reports for distributed decision makers. Typically when a management call for information fans out to a variety of operational areas each responds by providing data from the perspective of its functional activity. The result is that 1) the executive will be drowning in data with no options to apply analysis and transform the data to useful information, 2) the executive will receive multiple, biased, and conflicting informational answers from a variety of functional areas, 3) the information provided will be based on incomplete assumptions about the desired analysis and incomplete data in terms of integration with other functional areas, and 4) the failure to properly obtain the information will produce organizational overhead and require extra resources for unnecessary complexity. Nevertheless, we see several trends that offset this rather gloomy picture and set the stage for creating a data management culture that can effectively respond to management information requirements.

The first trend is a focused movement in our organizations toward greater efficiency and competitiveness. With this comes a willingness to consider a variety of strategic alternatives. These strategies include changes in support structures, processes, and responsibilities and often produce a rational analysis of data needs. Second, most business and institutions are undergoing migrations of at least some of their major

data bases to new operational systems, often moving to increasingly distributed modes of operation. New development tools, as well as off-the-shelf software, include structured methodology for clarifying an enterprise's data architecture, data definitions, and standard code usage. Since these methodologies provide the foundation for any executive information system, the migration projects present the opportunity to consider management's information requirements as part of the analysis. Third, there is recognition by small and large firms alike that there is a need for more participative management.² This compounds the numbers of managers who need to access and use data from across the organization. Thus, we see a momentum to also reach into other operational systems that are not immediate candidates for re-engineering. All of this leads to a realization that there are on-going data management activities that must be done to assure the availability of high quality information from all operational systems to support pro-active decision making.

There is much discussion in trade magazines and journals that couch these trends with such labels as re-engineering, right-sizing, and total quality management. While each of these movements points toward improvement of the information support functions and spurs an interest in what can be done to quicken and maintain progress, the lack of a structured process for coordinating the various data management activities and relationships across the organization limits and jeopardizes the excitement and potential momentum. Still, there remains a multitude of potent opportunities within the existing culture to use traditional relationships and friendships to develop prototypes that push management's "hot buttons" and thereby create support for ongoing quality data management work. Since culture is dynamic and changing, it is imperative to proceed with a sense of urgency to produce quick results that demonstrate value to those who can benefit within the current culture. By producing relevant prototype examples that address real current needs, there is increased likelihood that quality data management issues and projects will be included as part of other evolving organizational changes, thus, formalizing the data management function. This paper presents paradigm for a new data management culture and then presents a leadership strategy for achieving that culture.

A Data Management Paradigm

There are three key roles in a data management paradigm that work together to create an organization's data management process. One role, located at the data source point, is the **supplier**. This role is generally played by a variety of functional offices such as finance or human resources, which have operational responsibility for segregated sets of organizational activities. Suppliers are the essential actors insuring the reliability of the data. They must control the random influences in the capture and storage of the

² Harper, Stephen C., Ph.D. The Challenges Facing CEOs: Past, Present, and Future, In *The Executive*, Vol. VI, No. 3, August 1992, pp. 10-11.

data which can destroy the consistency, the stability, and the objective nature of attributes being captured. They are also responsible for documenting the capture and storage. Finally, the supplier must apply standardized recoding as part of creating cyclic extracts of the operational data for inclusion in the organization-wide repository, or data warehouse. These data custodial responsibilities are part of managing the functional area and are frequently delegated to designated system support stewards within the functional area. These distributed data administration activities are the basis of an organization's data management function.

Next is the work of specifying standards that will allow integration from the various functional areas and will support data restructuring and delivery of entities as required by the numerous data users. This coordination is done by a central data management function, which plays the role of **producer** of management information. The producer is typically the custodian of a central data repository which contains important parts of the organizational data extracted cyclically and stored in a standardized form. In general, the central function is a bridging activity and this activity often includes running the organizational data warehouse with varying levels of responsibility for distribution. The person responsible for the producer functions may be referred to as the Data Administrator where there is a very limited distribution responsibility, or the Data Manager where the transformation of data into distributed information is also included. With the restructuring of data in a warehouse and the distribution of information comes responsibility for the internal validity of the data. This is the ability to interpret what the data mean and what the information created via analysis means. To the extent that delivery of information is included as part of the central function, the producers must also have the analytical ability to integrate qualitative and quantitative data and to determine likelihood of events. This may also involve manipulating the data to move it toward the format of the business rules for the customer to use in distributed decision making. Consequently, the producer is involved in quality assurance and implements checking mechanisms to determine consistency of business rules across supplier systems.

The third role of user, or **customer**, involves using the information in some manner to identify the need to make a decision, select and evaluate alternatives, describe the situation, or defend and advocate previous decisions. The customer may want a sub-set of the organizational data selected based on pre-defined or newly identified key indicators. The level of flexibility that can be offered to the customer will depend on the sufficiency, relevancy, and timeliness of the organizational data warehouse. It is imperative that the customer be trained and learn how to evaluate the provided data in terms of sufficiency (some factors may need to be applied from external sources), relevancy (criteria for selecting sub-sets must be carefully determined), and timeliness (a determination must be made whether to use current data or trend data). These

skills for interpreting and effectively using data are necessary for users to fulfill their data management responsibility, that of assuring the external validity. This represents the degree to which the data and information can be appropriately generalized to the specific situation. Since the specific situation, as well as the elements of the decision, are known only to the user, it is clear that the user or customer must carry this responsibility. Similarly, the user is responsible for monitoring and measuring the value of the information. This is often referred to as the construct validity of the data and consideration of the degree to which "as-is" is becoming "to-be." In this regard, it is usually the user who is in the best position to identify changes in business rules that must be reflected by the organization's data. The user also needs to educate the supplier as to the true needs for data and the uses. This empowers the supplier to capture and store the most appropriate data.

It is important to note that each of the three roles (supplier, producer, and customer) bring a different and necessary perspective to data management. The suppliers must be primarily concerned with a specific decentralized function. They have little reason to consider data requirements beyond those that directly support the function unless required to do so by cultural norms or rules. The customers' views may be somewhat generalized in terms of the subject matter of interest, but are most often as narrow as the most recent requirement for information. Customers function in distributed mode and generally interact with others only as necessary. They have limited views of the needs of other customers. The producer's view is of the organization as a whole. Producers cannot effectively assess either the "as-is" or "to-be" infrastructures without significant interfaces with both the suppliers and the customers. Therefore, they must be central and accessible to both the suppliers and the customers. The producers are the central data management function. They coordinate and facilitate between suppliers and customers. Furthermore, they are the interpreters, providing translations to enable communication and transformation to relay understanding.

Characteristics of Decision Making

Information which reduces the uncertainty of the decision making or planning process is the goal of a successful data management function. In order to accomplish this, consideration has to be given as to what the information will be used for, the philosophy within which the information will be used, how the information is generated or created, who are the major players; and, finally, the current state of the campuses competing environment. Peter Ewell³ identified five primary uses of information. Understanding the potential use of information is important as one creates information to help the decision maker reduce

³ Ewell, P. T., 1989., Enhancing Information Use in Decision Making, *NDIR*, No. 64, Jossey-Bass

uncertainty of the decision. These five uses of information--rational decision making, problem identification, context setting, inducing action, and selling decisions--each require a different approach to the way in which the decision maker is supported. It follows then that the decision maker should be involved in the creation of this information. There are four primary decision making philosophies that we can anticipate being represented.

The first is the **political** philosophy where the primary concern is with others' perceptions. In general, this philosophy is found most often with the CEO or president of your institution. His or her primary concerns are often centered around political forces external to the institution (legislatures, other external governing boards, or other community groups), or internal groups (deans, students, or faculty senate). This person usually has many demands on his/her time and will require a very discrete data elements focused on a specific topic. In other words, if meeting the need, a single number is preferable to a detailed accounting of all components making up that number.

The implication is that the CEO has confidence in the decision support system and people to create reliable statistics. Again, as discussed below it will be important to involve him/her in initial discussions about the design of the data management environment.

A second decision making philosophy is known as **autocratic**. In this case, the decision maker is primarily concerned with a specific and often personal agenda. Information requirements are characterized by a continuous flow of information that is focused on specific issues. In other words, the usefulness of the information you provide to this type of decision maker may be compromised if more than one topic is addressed in a single report. On a university campus, the person that most likely reflects this decision making philosophy is the dean. The dean typically does not have direct responsibility for an academic program or research effort. The primary role is to build a college along an agenda he or she has set.

The **managerial** decision making philosophy is centered around the multitude of characteristics related to the specific program. The department head or chair typically reflects this type of decision making. In this case, the decision maker is interested in a continuous flow of information about any and all topics that are relevant to the programs in the department.

The final decision making philosophy is known as **collegial**. In this philosophy, the context of primary concern is that of correctness of process. Typically, the faculty senate will reflect this type of decision making when working through an issue. The problem will be studied from every angle and analysis often conducted

because of personal interest rather than from a notion that it would provide direct information to guide decision making. The type of information they require is often discrete but covering a broad range of topics.

While in the above certain types of individuals on the campus have been identified as reflecting specific decision making philosophies, it should be remembered that at any given time a president could in fact use a collegial philosophy or any of the others if it suit a particular situation. Decision support must recognize the decision making philosophy that is in play in order to maximize the usefulness of their decision support information.

Data Management Leadership

Successful data management is dependent on advocacy and leadership spread across the roles of supplier, producer, and user/customer. It is predicated on nurturing new working relationships between people who previously did not communicate or were adversaries. Furthermore, there must be input representing all four of the decision making philosophies: political, autocratic, managerial, and collegial. What follows is a discussion of a strategy as a four step cycle (Plan, Do, Check, Act) that can begin to increase the robustness of data quality improvement in the current structure while establishing the foundation for the new formal and informal structures needed to deliver quality information. These steps amount to a strategy for "growing" a data management culture.

Plan

A critical step is to conduct a systematic evaluation of current data management and information support and develop a strategic plan for fortifying the process by which data are managed. Given the "new paradigm for management that encourages highly dispersed decision making,"⁴ it would seem that one means for this activity to occur is to bring together the individuals involved in the development, analysis, and use of the data and to work them through a group process. Breaking the process into parts and involving individuals from various organizational units and perspectives is effective for quickly building consensus on immediate action steps. First, individuals in the group each identify issues. Small groups then cluster these issues, and barriers begin to surface. The next step is to focus on the root problems among the barriers, selecting one as a situation to address. This situation is then described as completely as possible and objectives and strategies to accomplish the objectives are formulated. Using such a process on a periodic basis accomplishes

⁴ Quote by Don Tapscott, author of *Paradigm Shift* (1992, McGraw-Hill., New York).

several things: 1) Awareness of problems and solutions increases and pervades the organization; 2) Do-able tasks can be selected to improve information support and data quality; 3) Cross-functional commitment and accountability are established; 4) Preceding successes are spotlighted in successive group processes. All of this means that incremental improvements are more likely to be sustained and resources allocated to make more progress. The plan must anticipate changing locations of decision making to refocus the type of information as well as the location of delivery.

Do

Actions taken to address the problems that become priorities as a result of the group process should be considered prototypes. This allows pre-established boundaries to be crossed as solutions are proposed. What works well can, and will, become accepted standard operating procedure. What doesn't work so well can be re-worked with a greater chance of success given the new knowledge gained from the prototype.

Although cross-functional support and participation is key to the success of any innovation to improve information support, responsibility for coordination must still be clearly designated. Again, developing a prototype data management function is a good way to begin. For a select sub-set of the organization's data, develop a process to support: 1) reliable data capture and storage at the operational level; 2) internally valid data, extracted and integrated in a central data warehouse; and 3) externally valid data delivered to meet the needs of distributed management. Under the umbrella of centrally coordinated data management are the following functions and tasks which they must either do or coordinate.

1. Information Planning

- a. Work with operational area managers or custodians, system support stewards, and distributed users to maintain and share a model of data elements and their attributes important to the organization.
- b. Help anticipate and respond to users' changing information needs.
- c. Coordinate appropriate policies for population and use of the data warehouse.

2. Standards Administration

- a. Develop and coordinate implementation of standards for data warehouse data elements and codes.
- b. Work with distributed users, operational area managers, and information technology personnel in the establishment and implementation of appropriate policies for data management.
- c. Provide standards and appropriate documents for use by EDP audits.
- d. Maintain inventory of official definitions of codes and values for standardized data.

3. Operational System Management Support

- a. Assist managers of operational systems in extracting, standardizing, and providing data for the data

warehouse.

- b. Assist managers of operational systems in developing and maintaining data elements and definitions.
- c. Provide limited training and advice to operational managers on developing and preparing individuals for data-management.
- d. Help operational managers create and coordinate user groups.

4. Team Building

- a. Coordinate identification of issues and strategies for dealing with data and information management.
- b. Support a quality development process with a baseline of standards and measures of improvement.
- c. Organize and provide administrative support to an information policy steering group which includes personnel from the information technology function, the central data management function, operational area managers, system support stewards, and distributed decision makers or users.
- d. Organize and provide training and administrative support to a group of distributed decision makers who use the data warehouse.
- e. Lead focused cross-functional projects on data quality improvement.

5. Data Administration

- a. Manage the data warehouse, serving as custodian and system steward.
- b. Support creation of a consistent and usable set of data warehouse data elements.
- c. Where necessary, mediate user and operational area custodial concerns on codes and data definitions consistent with organizational authorities.
- d. Distribute information about the data warehouse and its use.
- e. Insure proper archiving and protection of historical data warehouse data.
- f. Act as the official source for the audit trail of changing codes and data element descriptions for elements in the data warehouse.
- g. Maintain currency of all policies on information management.

6. End User Support

- a. Provide examples of using the data warehouse to address concerns and problems across the organization.
- b. Help communicate user's information requirements to management and to the operational areas.
- c. Support user integration of local data with the organizational data in the data warehouse.
- d. Assist with migration of users' local data to the data warehouse when it is relevant to activities, analysis, or reporting that cross functional boundaries.

7. Technical Development

- a. Work with operational area managers or custodians, system support stewards, and end users to clarify technical needs.

- b. Support the information technology department in development, purchase, and use of prototype tools and products for managing and delivering data.
- c. Act as a clearing-house for data management tools and related technology.

Check

For information support and data quality improvements to be sustained, it is critical that on-going stewardship be executed by operational staff in the functional areas. It follows that there must be a mechanism for checking and verifying the adequacy of this stewardship. This amounts to assigned accountability for the following functions as a minimum:

1. **Data standards and documentation.** Work with the central data management function and others to develop and implement standards for selection, compatibility, integration and accessibility of organization-wide data. Incorporate and document implementation of these standards.
2. **Data capture and storage.** Assure reliable data capture, maintain list of allowable values, and archive data at standard cycles.
3. **Data validation and correction.** Implement and document validity checks and documentation in applications that capture, update, or report critical data. Develop and measure data quality and respond to problems by repairing the erroneous data, adjusting the processes that created erroneous data, and notifying impacted users of the corrections.
4. **Data security.** Implement and document access procedures that provide adequate protection as determined by management and monitor violations. Implement back-up and recovery procedures that protect against threats to data integrity from system failure, faulty manipulation, or other disasters.
5. **Data availability.** Provide accessible, meaningful and timely machine readable data which clearly identifies capture and modification dates. Work with central data management to provide training and consulting on data use and solicit input for improving data quality and delivery.

Act

The work of the cross-functional groups and the refinement of prototypes as a result of use and feed-back continually nudge and encourage a data management culture. As pockets of support and commitment began to gel, the time is ripe to pour a firm foundation for data management. It is time to act....to formalize the coordinated data management function. Management of data and information support can become a critical part of the strategic core of the enterprise and those who are involved in these areas can become teachers and leaders for others. The leaders will specify the people, activities, data, and tools that have potential to impact all facets of an organization and its outcomes. Some of the basics include:

1. Develop a data management policy.
2. Develop a criteria for selecting data required by distributed decision makers for inclusion in a data warehouse.
3. Follow the organization's lead, focusing on data from systems or processes that are targeted for re-engineering.
4. Develop a standard for data definitions and descriptions for data warehouse data.
5. Select data management tools such as a data base system for the data warehouse, a complimentary data dictionary, guided user interfaces, report writers, etc.
6. Develop criteria, procedure, and process for refreshing and archiving data warehouse data and definitions.
7. Create cross-functional teams from system migration work-groups to continually address data requirements and quality issues. These groups should include individuals who support and who use the operational transaction systems, individuals who work with data from these operational systems to provide it to decision makers, and individuals who use data provided for analysis and decision making.
8. Create a map of the information management structure with names and job descriptions on the various functions.
9. Develop training at all levels of the information management structure (operational area managers, system stewards, users) on the data management functions, procedures, and tools.
10. Initiate a program of professional development for those involved in data management and information support that sharpens project management, team leadership, and communication skills.
11. Create an environment and structure so the act of integrating improvement identifies the next set of activities where one can plan new ways to remove waste, scrap, and rework.

Culture in Transition

Progress in evolving a data management function or "growing" a data management culture will not always be readily apparent or easy to point out. Behaviors must change before the continuous improvement of data quality becomes a recognized organizational goal. The changes are subtle and incremental. A relevant model for the response of an individual to change is the Kubler-Ross sequence for grieving, i.e., denial, hostility, bargaining, depression, and acceptance.⁵ This model is pertinent because it recognizes that change is first and foremost the loss of the "way we have always done it," and it is always traumatic because comfort zones are threatened. If we do not witness some of the behaviors in this sequence, it is likely that nothing is changing.

⁵ Kubler-Ross, F., 1974. *Questions and Answers on Death and Dying*, Macmillan

There are several cultural indicators we should observe during transition that can alert us to the areas that require particular attention. These indicators can also help us assess whether we have truly evolved to a changed culture, or in the case in point, to a new data management culture.

1. *Affective Orientation* - Do people openly acknowledge and value working relationships with individuals across the organization?
2. *Orientation to Causality* - Do people claim ownership for problems as opposed to blaming others or the system?
3. *Orientation to Hierarchy* - Do people accept working with others at levels above and below them within the organization's hierarchy?
4. *Orientation to Change* - Are people willing to take risks and embark on new ventures?
5. *Orientation to Collaboration* - Are people receptive to others, seeking input and considering a variety of perspectives?
6. *Orientation to Pluralism* - Do people who represent different interest groups make an effort to make contact and communicate with each other?⁶

To the extent that we can answer "yes" to each of the above questions, we are well positioned to evolve a data management culture that indeed embraces "new ideas for new challenges". We must remember that change is a process and subtle attitude changes signal progress. These include shifts from managing to leading, from control to coaching, from quantity to quality, from opinion to information, from resistant-to-change to open-to-change, from people-as-commodities to people-as-resources, from suspicion to trust, from compliance to commitment, from internal focus to customer focus, from individual to team, from detection to prevention.⁷ All of these transitions are part of growing a data management culture. It is important to watch for these transitions and recognize them as gains in terms of awareness and commitment. They signal emerging pockets of support and fertile ground for adjusting priorities and strategy to use new allies as a re-investment in the desired data management culture. Data management can no longer be relegated as a technical issue. There will need to be changes in our organizations. Matrix organizations and cross functional teams focused on specific tasks or problems must become common place. This places the data users and those who need information in direct contact with the operational sources of data. Those who produce information for decision

⁶ Kuh, George D., and Elizabeth j. Whitt, 1988. *The Invisible Tapestry: Culture in american Colleges and Universities. ASHE-ERIC Higher Education Report No.1. Washington, D.C.: Association for the Study of Higher Education. pp. 104-105.*

⁷ Plice, Samuel J. Changing the Culture: Implementing TQM in an IT Organization. In *CAUSE/EFFECT*, Summer 1992, p. 23.

making must shift from the roles of manipulating and analyzing data to roles of facilitator and trainer. Furthermore, those at the operational level will need training on the use of data and tools relevant to the new set of data stewardship activities which they will be expected to perform.

The organization which tries to distribute data without data administration to define it and document it will fail. The organization which does not support the distributed data management function to push out the data and transform it into information will fall behind its more progressive competitors. The organization which does not restructure to be consistent with a revised decision structure will frustrate its employees and produce feudal turf battles. The organization which does not continually train its employees to manage and use information will make questionable decisions.

We must step forward and nurture a data management culture. Deal and Kennedy suggest we must resist the temptation to roll up our sleeves and wade directly into the resolution of the problem as traditional managers if we truly want to encourage lasting change or transformation to a new culture. They suggest as an alternative the notion of a "symbolic manager" who recognizes "that the longer-lasting solution is to rely on the culture to meander its way to a solution of the problem."⁸

"New ideas for new challenges" amounts to creating a web of intelligence from the ideas and vision we have today. The web of intelligence is the new data management culture. It must be anchored by trained and empowered people who perform relevant and timely activities using sophisticated navigational and analytical tools on properly integrated quality data from both internal and external sources. "New ideas for new challenges" means change and the first change is at our own doorstep.

⁸ Deal, Torrence E., and Allen A. Kennedy. *Corporate Cultures: The Rites and Rituals of Corporate Life*, 1982. Addison-Wesley Publishing Co., Inc. p. 154.

A Model for Change
William R. Brunt
Regents College
Albany, New York

This presentation will cover management principles and organizational structures with specific strategies which are both applied and untried to create an IT provider capable of change.

One principle is to take those tasks which have been traditionally assigned to the IT provider and distribute to the user community. For example, a specific strategy is to share the process of technology implementation which is demonstrated, documented and encouraged by the user community. An educated user community which can carefully articulate their needs into report or processing specifications can be achieved through cooperation of provider and user. Users will become familiar with the concepts of files, fields, directories, reports, break points, sorting and filtering. The user community will actually write a large percentage of their own reports.

Introduction

This paper and presentation is intended to present a model for change. This model covers management principles and organization structure to facilitate the implementation of change. As each concept is introduced, actual strategies are described. Some strategies have been used and a discussion of their success follows. Others are untried and it is hoped to implement them in the future.

The following management principles and their specific application are discussed.

1. Work Groups
Specify, Code & Test
Individual Reports
2. Distributed Training/Support
Computer Trainer & Computer Training Group
E-mail SUPPORT
3. Flatter Organizational Structure
Management Team & Technical Pool
4. Functional Specification by the User Community
New Projects initiated with a specification
5. Operational Responsibility Transition
Technology Implementation
Human Resources creates/maintains user accounts
Facilities trained on the cable plant and equipment moves
6. Sharing the Resource Allocation Decision
The Automation Planning Committee and the request queue

Some of these techniques have been used at Regents College in Albany, New York. The College required a structure capable of embracing change as technology introduction had been essentially on hold for 12 years. During the last two and a half years, the College went from having a single person responsible for technology to a staff of 12. During that time, the College went from having approximately 30 stand alone personal computer to 200 in a network configuration and its own administrative data processing system using a commercial product.

Work Groups

Work groups can be utilized in a variety of ways. One involves the production of services. A second use of work groups is in the reporting and task assignment.

All projects, tasks and delivery of services can be completed within the framework of the work group. The minimum size of this work group is two. Consider the creation of a new screen to enter applications from students who already possess significant college level credit. The initial conceptual discussion and final product are all accomplished within the framework of a work group.

In the conceptual discussion, the team approach utilizes a dialogue between the head of a program, the advisors to prospective students and a representative from the computer services department where an outline of the basic functionality is covered. Contrast this approach to the traditional hierarchical structure where the program director indicates the screen requirements in a separate discussion to the representative from computer services in the absence of others.

In the team discussion, the staff from computer services are able to lister to the operational needs as outlined by the staff on the front lines. The program director can react and provide feedback on the larger operational issues. The staff from computer services are able to comment on the functionality which can be provided and the relative costs thereof. The program director can react to the costs of different functions and features and judge the benefits with feedback from those on the front lines.

In a traditional hierarchical approach, this same information will need to be communicated and will take about the same total amount of staff time but will occur over a longer period of time. For example, the program director will work with the front line staff to create a rough functional specification. A separate discussion will then be held between the program director and computing staff. The task will be outlined with little opportunity for any cost benefit analysis of functionality. The computing staff will then begin a specification and subsequent coding. Any questions which arise during the coding will be presented to the director of the program and will require further consultation with the front line staff to resolve.

Once the application has been coded, a team approach will require that testing be performed by another member of the computing staff. This accomplishes several things. The application becomes familiar to at least two people on the technical staff. The testing by someone who did not write the

application ensures a higher quality product and also provides some redundancy for support of said application. Additionally, a formal program review of one person's code by another ensures that new programming techniques are shared and a sense of consistency is continually refined.

Using a work group in the delivery of services shortens the product life cycle and improves quality.

This technique of sharing the development of new applications has been used very successfully at Regents College. It has resulted in higher quality and increased sophistication of the applications.

The reporting and task assignment in a work group means that all members of the group will be aware of the activities of each other. An actual strategy can take the form of individual reports at staff meetings of the tasks since the last meeting and those planned until the next one. This has the benefit of drawing upon the different experiences and talents of all members of the group as plans are developed. It has been my personal experience that these meetings often result in a new better course of action. People must be encouraged to freely comment without personal bias on the work of others. This is perhaps the most difficult to achieve when personnel are at different levels.

It has taken several years for this sort of communication to develop and flow freely. Success in this area will only come in the long term.

Distributed Training/Support

Training can be accomplished in a distributed fashion by having both a dedicated computer trainer and a computer training group. The computer training group are staff from throughout the organization who are interested in computer training. They serve on a committee responsible for developing and providing training. Distributing training among a group accomplishes several things.

As large projects come to completion, a sizable group of trainers is available to bring even larger groups of people on line. The computer trainer can act as an expert providing guidelines and standards to facilitate a structured guided learning process. The amount of training which can be completed over a short period of time is increased. The demand for training will cycle as the IT providers alternate between development and delivery. For example, it may take several months to bring a new application into production. If it is a large application, the demand for training will be great. The period of time to

accomplish this training will be shortened by having a large pool of trainers available.

The computer training group will also serve to critic new lessons as instruction is developed. This process has been termed piloting. It is assumed the computer training group will assist both the dedicated computer trainer and other trainers in the group when piloting new material.

Equally important to distributing training is feedback on the quality of instruction given. Each lesson will require written surveys covering the quality of instruction and relative value of content to expectations.

Flatter Organizational Structure

In computer services, Regents College has managed to flatten the organizational structure. There are 12 positions within the office. Of the twelve, ten are professional staff. The interaction and supervision of professional staff requires significantly more time as the decisions and day to day work are more complex than the route tasks usually assigned to support staff.

The flattening of the organizational structure has been achieved by supervising work on a project by project basis. The director and associate director divide the supervision and the pool of technical staff to work on these projects. Rather than having individuals report on rigid lines of communication, reporting is accomplished on the individual project. Upon completion of the project, the portion of time dedicated to a particular project is put back into the pool and the next task is begun.

On a weekly basis, the director and associate director meet to update one another on each other's progress and to critic and comment. It is assumed communication is complete and sufficient to allow either to assume the other's role.

Functional Specification by the User Community

It has happened where new applications are begun with nothing more than a thumb nail sketch on a napkin over lunch. The project is described as very straight forward and easily completed in several days. The actual result takes many weeks and encompasses issues which were never considered in the initial discussion. There is little appreciation for the issues which arise during the application life cycle. The solution is to require a complete function specification be completed by the user community.

It used to be that concepts such as top down or bottom up design were terms used solely by systems analysts. Today, it is common for people to understand some fundamental design techniques and functional requirements. Training should not only encompass the use of applications but the specifying thereof.

Users are becoming increasingly familiar with terms such as files, fields, directories, reports, break points, sorting and filtering. Several levels of end user functional description can be put to use. For example, IT providers may in fact write the process description and merely require that decision makers within the end user community sign off on the deliverable. This will have mixed results. Some will read in details the proposed system and others will not even glance at the description and place their signature at the bottom. I have found a direct correlation in demand for services and the depth of review. The most sophisticated users will completely describe an entire system and deliverable. This will allow those involved in application development to concentrate their skills on producing deliverables.

Many are wanting to have ad-hoc query tools placed in their own hands. The writing of reports by end users serves to make information more accessible and give a greater appreciation for the work of IT providers.

Regents College has achieved mixed results in these areas. The level of participation is really left to each unit manager. Those who have chosen to become very involved will ask for the most sophisticated application. Additionally, those who are quite involved also require the least re-work as the application is put to use.

Operational Responsibility Transition

Operational responsibility transition involves taking those tasks which have traditionally resided in the IT domain and distributing them to the user community. Examples of this include technology implementation, system account management and hardware relocation.

The implementation of technology has usually fallen on the shoulders of IT providers. The real work in providing technology is to articulate a particular solution and build consensus. This does not require the skills and capabilities of a seasoned data processing individual. On the other hand, it is not a realistic expectation to call upon the end user to fully articulate a solution without some prior experience.

A Model for Change

Partnerships between IT providers and end users will be formed through the successful completion of projects. This will give the end users a better grasp of the steps required to introduce technology. The process will become clear and the role of introduction can transfer from IT provider to end user. The final result will be that the IT provider can concentrate on creating and supplying new technology in concert with the end user community.

It has been the norm for decades that only IT providers could be entrusted with the task system account maintenance. This requires that Human Resources and/or the hiring program notify the IT providers of new employees. The notification requires communication and overhead for a task which is route and can be best accomplished by Human Resources. In other words, have Human Resources create the new accounts on the various systems.

The hiring of new employees always involves providing a physical space. Either facilities or office management are called upon ensure such space is provided. Increasingly, this space will involve a work station of some form as the "PC on every desk" is the norm. Physical facilities would then deliver and install the PC or equivalent to the desk top. Any cross-connection at an intermediate distribution frame or main distribution frame would be handled by facilities as well.

The maintenance of accounts and providing of desk top computers are just two examples where responsibility for providing technology can shift. It is a case of making a single entity responsible for servicing the new and exiting employees.

Sharing the Resource Allocation Decision

The decision of which projects to work on next when provided with a limited amount of resources is best shared among those expecting service. A high level committee entitled the Automation Planning Committee could be charged with examining the projects and determine what is to be done next.

This process would not divide the resources equally based on some pre determined budgeted charge back system. Instead, it would be the committee's charge to prioritize projects in the order which would give the institution the most benefit overall.

The request queue and its handling by a committee has had mixed success. Some new initiatives and programs are started without regard for needing computer support. The automation of the initiative then bypasses the queue and receives immediate and rushed attention. At the time of this writing, I am trying to

create a system which would be responsive to the needs of all and fair in the distribution of services.

I have found the most difficult challenge in creating a model for change is to determine which project will be worked on next. Those who receive attention from the IT providers will consistently use the queue mechanism. Those who do not receive service based on the committee's decisions will seek alternate providers whether they integrate into the developing information systems or not.

The task of actually deciding within the framework of a committee is not easy by any means. Each member of the committee will view the projects with different returns and costs. Some members may require several lists in different categories while others will prefer to work in a single format. As I said earlier, this is one of the biggest challenges in creating a model for change.

Conclusion

I hope that you enjoyed the presentation and this paper. I would be most interested in any comments you may have. I am particularly interested in any techniques which I have not mentioned which facilitate the introduction of changing technology.

The paper is geared towards creating a core of highly capable people within 'IT' to provide services and relying on many operational aspects from the user community. It is most useful when organizations need to play catchup. It is not meant to be utilized when organizations are wanting to down size and consolidate.

Implementing Distributed Computing at Cornell University

Mark Mara
Associate Director, Information Resources
Cornell University
Ithaca
New York

CAUSE 93

Abstract

This paper describes the approach used at Cornell University to enable information technologists and user staff to develop client-server applications. The lofty design principles of open architecture, standards, reusability, scalability, and portability were constrained by limited budgets, legacy data, time to market, security, performance, and the cost of training. Considerable experimentation with prototypes led to the decision to invest in two areas:

- (1) creating a developer tool kit – an integrated set of applications, tools, protocols, and procedures; and
- (2) implementing several key system infrastructure services.

The project that created the tool kit and implemented the infrastructure services was called Project Mandarin. This effort has resulted in an environment that facilitates the rapid creation of client-server applications.

The intent of this paper is to share our experiences. I will explain why we chose the approach we did, briefly describe the Mandarin environment, and describe our experiences implementing several projects using Mandarin technology. These brief case studies will be used to discuss some of the problems, required cultural changes, costs, and benefits of our approach.

Background

In April 1990 Cornell University, with support from Apple Computer, launched a two and a half year project to provide end-user access to data stored in enterprise database systems. Pennsylvania State University and Massachusetts Institute of Technology served on the Steering Committee and have assisted with development. This effort was named Project Mandarin. The original project, Mandarin Release 1.0, is now complete. Mandarin Technology is being used at Cornell to create an ever increasing number of services which are delivering information to students, faculty, and staff. Mandarin technology has empowered colleges and departments to create innovative solutions to their business problems. It also has created many new and in some cases unexpected challenges.

The Business Case

The only justification for any Information Technology (IT) project should be to service the business in some way. Part of the mission of Cornell Information Technologies (CIT) is *to empower people to get the information they need, where, when, and how they need it.* Project Mandarin's role is to provide a technical environment which will enable the rapid development of applications which support this vision. Working from the mission we identified the following business objectives:

Business Objectives

Local empowerment—is embodied in such initiatives as total quality management, legendary service, and self-managed teams. We are moving from hierarchical decision making to distributed decision making. This calls for changes to our traditional information systems and changes in who we consider our customers. Information is power, to support local empowerment we must provide access to information to a large segment of the university community that we have not traditionally considered our customers.

Supporting the "new" information technologies user—is critical to the continued existence of an IT organization. We used to consider a dozen central offices as our users. The definition of user now has changed to include all members of the university community. At Cornell this changes our customer base from a hundred or so central office users to over 20,000 members of the university community. These new customers are not using IT systems every day. They are what we refer to as casual users; users that interact with information systems once a week or maybe a few times a year. They demand high levels of integration and consistent, intuitive interfaces.

Delighting the real customers—is the only way to survive in days of severe fiscal constraints. Cornell University is committed to improving the service provided to its students. We see this initiative as an opportunity to use information technology to assist students in dealing with the administrative requirements of university life. We have the same opportunities with faculty and staff. However, these new customers are not tolerant of time to market measured in years. Their experience with micros has led them to expect instant gratification. They are looking to information technology to reduce the time spent waiting in lines, cut down on paper work, and in general, improve their ability to accomplish their goals.

Saving money—is a very popular objective. Given the reality of today's financial situation, every activity needs to be scrutinized in terms of costs and benefits. For infrastructure and tool building projects such as Mandarin, this is a particularly challenging area. We believe that Mandarin will give us the ability to leverage development across institutions resulting in significant savings in development costs. Distributed technology will allow us to make use of low cost RISC based hardware reducing the ever increasing load on our mainframe. The high level of shareability will reduce the cost of development for any one system and also reduce the maintenance cost for all systems.

Working from the business objectives, we arrived at a set of goals and design principles for the project. The goals set the general direction for the project and the design principles describe in a very general way how we planned on getting there.

Project Goals

Provide access to central, departmental, and local information—One of the most frequent and loudly voiced complaints about Information Technologies organizations is that they do not provide timely and friendly access to information. Once the political hurdles are passed, we must not allow technology to be the excuse for not delivering information to our customers.

Engineer the illusion of simplicity—We cannot expect our new customers to learn arcane rituals to access information. The complexity of the underlying systems must be hidden. We also need to provide a way for data, from several heterogeneous sources, to be displayed to the user as a single data source. We do not have the funds or the desire to convert all of our legacy systems into intuitive user-friendly systems at this time. What we need to do is provide a simple way for people to access information without having to replace our existing systems.

Build infrastructure and develop tools—The key to controlling costs is to create a development and run-time environment in which most of the functions that need to be accomplished are available as services. This allows developers to concentrate on the unique parts of the application and not recreate functionality that already exists in other applications. Using tools to develop applications allows staff to move into distributed computing without extensive retraining; development time is reduced; and a level of consistency is introduced into applications.

Object oriented, modular design—We believe that object orientation is the way to achieve the flexibility and shareability required to accomplish our other goals. We do not see Mandarin as a one-size-fits-all solution. It is, rather, a kit of objects that can be used alone or combined in various ways appropriate to the target institution.

Design Principles

Open Architecture—The architecture provides a framework for assembling various objects into a distributed computing environment. The architecture must not lock us into any one vendor or technology. We must allow for the same functionality to be implemented in different ways.

Standards—The adherence to standards is closely bound to implementing an open architecture. The protocols, interfaces, data encoding, and object definitions must not preclude the use of existing and future vendor supplied products.

Shareability—The key to reducing development time and maintenance costs is the ability to implement a function once and have multiple applications access that functionality. The real win is to isolate common functions, implement them once, and allow objects to access them as needed. This is what we mean by shareability.

Scalability—Solutions must function at various levels. Some applications will have an audience which numbers in the tens of thousands while others may have ten users. We need to support department and office level applications as well as campus and multi-campus applications.

Portability—We are not trying to build a solution for Cornell University. Our intent is to build something that will be useful to other institutions. By focusing on portability we will avoid the trap of relying on some Cornell specific technology which would limit our future options.

Reality Check

After setting our goals and identifying our design principles we reviewed the situation at our institution. This reality check led us to the following conclusions:

Legacy Data—The university has made a major investment in creating and maintaining operational databases. A lot of the information contained in these databases would be of interest to our customers if they could access it conveniently. The problem in providing access is that the databases were not created with ad hoc query or end-user access as design criterion.

Time to market—Time to market is an area where IT organizations are under a lot of pressure. Systems development times are measured in years. By the time a system is implemented, the original requester is no longer in control of the area and the business requirements have changed. It should be no surprise that the delivered system is not perceived as a great success.

Security—The replacement of terminals with microcomputers running terminal emulation programs has resulted in a general reduction of security. The microcomputers are typically connected to a network. Passwords and data are flowing in clear text over the network. To the user, the situation does not seem to have changed from the days when they were using a terminal hard-wired to the mainframe, but in truth security is significantly lower.

Performance—We are running out of capacity on our mainframe. Assuming that demand continues to grow at historical levels and mainframe cost/performance does not radically change, the cost to provide mainframe computing power will be probative. Our users are clamoring for distributed implementations making use of RISC based hardware.

Staff training—Clearly the move to distributed computing will require the retraining of the IT staff. We still have all of our operational responsibilities. The need to cut paychecks, register students, or solicit alumni support will not stop and allow us to focus on a retraining program. Staff are currently being pushed to the limit just to keep up with their workload and stay current with evolutionary changes in technology.

User training—Our new definition of customer could result in the requirement to train thousands of users. Clearly this is going to be very expensive and time consuming. We need to design new applications in such a way that they do not require extensive user training.

Client-Server Investigation

We surveyed the state of client-server computing at the start of the project. In 1990 we found very few examples of successful implementations of distributed computing. The ones we did find had several common characteristics:

Use Relational DBMS and SQL—The most common implementations involved the creation of a data warehouse utilizing a relations DBMS. Access was implemented by the use of SQL or SQL based tools.

Use Shadow Database—Data was extracted from operational systems and loaded into the warehouse.

Put Emphasis On Decision Support—Most of the system were oriented towards decision support.

Retrain Staff—Extensive retraining of the development staff seemed to be common.

Throw Money at it—Most projects involved the purchases of a relational DBMS and SQL based tools. When the cost of staff and user training is added in we were getting estimated total costs measured in the millions of dollars.

Mandarin Choices

This approach did not seem to fit our goals and design principles. In addition to not addressing our immediate concerns, it would cost more money than we had available. As a result we chose to go against the trends. We felt the need to demonstrate some quick wins with as little expense as possible. This led us to the following choices:

Support a variety of database technologies—We did not want to be locked into one vendor's DBMS. We decided not to incur the expense of implementing a new DBMS, but rather, design an architecture which will allow us to support our existing systems and give us the flexibility to access other DBMSs at the same time. We have a staff trained in the use of Natural, a 4GL, and experienced in working with our ADABAS DBMS. We wanted to make use of this expertise. We also wanted to support the hot relational products such as Oracle and Sybase.

Use Operational Data—Although there are many cases where the data warehouse concept is appropriate, we decided to provide access to our operational databases. Since the data which most of our customers are interested in is contained in our these databases we decided to provide direct access to this information.

Provide Real Time Access—We wanted to provide real time access to this data. We saw an opportunity to improve the quality of our operational data by allowing direct update of demographic information by end-users.

Isolate the DBMS from the client applications—We view the client/DBMS interaction as the client invoking a host-based service rather than simply calling a DBMS procedure. Because the operational databases were not designed for end-user access, we need to provide the ability to interpret the high level service request and then access and process the data required to provide the requested information. The use of a Data Access Service isolates the client applications from the DBMS.

Use existing skill where possible—Implementing the servers in the 4GL, which the staff already knew, significantly reduced the training requirement. Using the shareable services concept hides most of the distributed computing aspects and allows the staff to simply extract the data from the database, a task that they know how to accomplish. Providing a graphically oriented client builder minimizes the need for technical training for client developers.

Just Do It—We will try to avoid the temptation to design elegant solutions that cater to the technicians' egos more than the users' needs. We accept the fact that we are not going to get it right the first time. We see the development of distributed computing as an evolutionary process. The most important thing is to get something in place and learn from the successes and failures.

Adopting a Technical Architecture

A technical architecture is a description of design objectives and principles for systems construction and operation. It describes how applications should be constructed and how they are linked to data and services. The adoption of a standard technical architecture is critical in implementing distributed computing. Applications must be able to access system infrastructure services. The services must be able to communicate with other services and various data stores. This type of communication and shareability cannot occur without an overriding technical architecture.

Early in the project, we made an unsuccessful attempt to define a technical architecture. We then decided to construct some prototype applications. After gaining experience building and supporting prototypes, we at least knew what we did not know. Object based design, client-server design, and distributed architectures are not things that are easy to get right the first time. It is necessary to try them a few times to understand the problems that need to be solved. At this point we encountered the VITAL technical architecture.

We found Apple's VITAL (Virtually Integrated Technical Architecture Lifecycle) to be of great value in evaluating our current architecture and planning for our next release. VITAL is the result of thousands of hours dedicated to developing a plan for implementing distributed computing. A basic premise of VITAL is that custom systems need to be constructed of generalized modules that can be reused and shared. This is clearly consistent with our project goals. When we reviewed the VITAL design principles we found a 100% congruence with our design principles. It was comforting to see that our architecture fit within the VITAL framework.

We are building Mandarin using VITAL as a road map, or building code, to help us avoid costly mistakes. Our goal is not to implement VITAL. However, as we evolve Mandarin, VITAL will play a key role in defining our directions. Our product will be consistent with VITAL not because we have mandated the use of VITAL but rather because VITAL makes sense.

Key Services and Tools

Our experience with prototypes convinced us of the necessity to invest in the development of several critical systems infrastructure services. A service is a basic unit of shareability. It is called upon by other applications or other services to perform some action. A service is usually implemented by a process (a server) which may require access to data or metadata and a desktop integration component.

VITAL defines several dozen services. We identified the minimum set of services we needed to implement the functionality required by our first group of users. (It is important to differentiate these services from network-based services such as electronic mail, Gopher servers, or library catalogs.) We selected the following set of services to support our first generation of client applications:

- **Directory Interface**
- **Authentication**
- **Version Control**
- **Software Distribution**
- **Authorization**
- **Authorization Aliasing**
- **Database Access**

The Data Access Service will talk to any authenticated and authorized client application. We have implemented clients in 4th Dimension, HyperCard, and C++. However, to facilitate the rapid development of client applications we created a simple but powerful application creator and associated run-time environment. The application creator is called **ProbeMaker** and the run-time environment is called **ProbeLauncher**.

Case Studies

This section describes some of our experiences using Mandarin technology to solve business problems. Before looking at specific projects it is useful to briefly discuss the criteria for selecting the projects used to introduce this technology.

Picking the projects

The combination of the client building tool and the infrastructure services has resulted in a development environment which allows the rapid creation of client-server applications. These tools are not appropriate for solving every problem. We developed some general guidelines for selecting good candidates for first projects:

- **Project Champion**
- **Quick Win**
- **Not Mission Critical**
- **Accommodating Users**
- **Visibility**
- **Limited Complexity**

Just the Facts

Just the Facts (JTF) was the first application built with Mandarin technology. The project was initiated by senior university management asking us if we could use technology to reduce lines in the student services areas. In our surveying, we found that a high percentage of student inquiries were about academic records (course information, schedules, grades); financial information (CornellCard, bursar bill, financial aid information); and personal information (address, phone number changes). The retrieval and display of this type of information is an ideal application of Mandarin Middleware.

This was clearly not a mission critical application, the students could always go back to waiting in lines. On the other hand, it was definitely a high visibility system. We found the students to be a very accommodating user group.

Just the Facts was developed by a team made up of staff from the VP for Academic Programs office, Project Mandarin, and the programmers who had responsibility for the legacy systems. The result of this project is that Cornell students can check their transcripts and bursar accounts without going to our administrative building and waiting in line for a clerk. Now they can get this and other information from Just The Facts by using their Network ID. When they type in their Network ID and password in the main menu, Just The Facts displays their grades, bursar accounts status, CornellCard status, and addresses on file with Cornell. Just The Facts also lets them view their current course schedule and update most address information. Additionally, they can communicate with the bursar and registrar via e-mail from within Just The Facts.

We learned that technology was only one of the requirements for making distributed computing happen. Cultural changes must occur at all levels in both the user and IT communities. We were fortunate that the time was ripe for change. The university was investing resources in the "total quality improvement" process. Shrinking resources prompted inquiries into how the university could deliver better services for less money.

Organizational structures inhibit change. In many cases the organizational structure is so rigid that any change or innovation will meet resistance, the "If it ain't broke, don't fix it" mentality. An example of this was our attempt to allow students to directly update their own address information. There was much concern that the students would "mess up" the data. This has not proved to be the case. JTF has cut in half the number of manual entry address changes that staff must make (10,000 reduced to 5,000 changes annually). It also reduced the number of errors because the students are inputting their own addresses. The percentage of students who make their own changes continues to increase.

Administrators were generally pleased. JTF greatly reduced the long lines and allowed them to devote more of their time to addressing more significant problems and issues. The number of requests for printed grade reports and transcripts has significantly decreased because students can now access the information themselves and see the information on screens in front of them.

Students were thrilled. In the first ten months, 11,072 unique students accessed JTF in 37,585 sessions. We made it very easy for the students to comment on the services. They have been particularly enthusiastic about the ability to access the information at their convenience as opposed to being limited to when the administrative offices are open. We are constantly getting feedback about their needs, concerns, and wishes. The Mandarin technology allows us to be responsive to their comments (the system is now in its third generation after only two years of being up and running).

Change breeds more change, and change never ends. The first JTF was a prototype - we had no intent of creating anything more. But that quickly changed after its introduction. Students loved it. Don't start anything unless you think you can see it through. If you raise students' level of expectation about service delivery, it is very hard to go back. Once the waiting lines are gone, students are very reluctant to stand in them again.

Bear Access

Bear Access is a CIT sponsored project that provides a consistent interface into a suite of network-based service offerings. These applications provide access to a variety of services over the campus Internet. In addition to providing a common access point for a heterogeneous collection of services, we wanted to be able to administer a group of general university services centrally. We needed to allow for other groupings of services to be managed at department or college level. We also needed to be sure that the software we provided was always current.

We used the Mandarin LaunchPad, Version Control, and Software Distribution services to meet these requirements. Problems arose about who was responsible for what services. There should be some services that are supported centrally and are subject to installation cycles and support requirements. It is equally important to allow the community to readily create and distribute new services. We must educate the users to go to the service owner for support. We also must enhance the Mandarin infrastructure so that it provides the users with information whether a problem is with a service offering or with the infrastructure supporting the service.

Faculty Advisor

The College of Human Ecology, the Office of the Vice President for Academic Programs, and CIT are sponsoring a project to provide the faculty with immediate access to current student data. Faculty members saw JTF and were very interested in being able to have access to information in a similar fashion. But they needed different kinds of information (primarily academic records) and in a slightly different format. This prototype gives faculty information about a student so they can give better service to the student. This project demonstrates the high degree of reusability possible with Mandarin technology.

Just the Facts, Bear Access, and Faculty Advisor are examples of using Mandarin technology to deliver information to the university community. Mandarin technology has allowed these services to be developed very quickly, typically in a matter of weeks. These products are sharing the infrastructure created by Project Mandarin. Each new service does not have to be built from scratch. It can be put together from reusable building blocks. This ability to leverage existing effort is critical in delivering cost effective solutions. We have seen a high level of reusability on these projects which contributed to faster development than we experienced with other technologies.