Implementation Report of the Subcommittee on Research Computing

Summary and Recommendations

The Computer Planning Board has recommended that Rice University provide a computing environment for education and research consistent with the university's commitment to academic excellence. This report contains the principal recommendations of the Subcommittee on Research Computing aimed at achieving that goal. The individual recommendations are discussed in more detail in the following sections of this report. Cost estimates are made for each recommendation and these are summarized in Tables 1 and 2 at the end of this section.

Some of the recommendations made in this report go well beyond what is currently happening at Rice. None of the recommendations, however, go beyond what is currently policy at comparable universities. If the recommendations of this report are all implemented, an excellent research computing environment at Rice will be the result.

What is envisioned is a distributed computing system, with computing being done on the smallest, most convenient computer capable of doing the job. Such a system has the following components.

* Micro-computers and workstations available to do the required computing at the lowest level, and to provide the entry point to more powerful computers via a network.

* Local mid-range computers for focussed research groups.

* Large central computers to provide the CPU power to handle large problems that are too computer-intensive to be done on available local machines, to provide the memory space and high speed I/O that is needed for data-intensive problems, and to provide the mid-range computing power for those users who have no other computer available.

* Supercomputers to provide the state of the art computing capability to deal with problems which are too large to be handled by the smaller computers available.

* A high speed network to connect all the components and provide easy transmission of data as well as easy access to remote computers.

* Provision for the maintenance of all of the components of the system.

* A support staff adequate to ensure the efficient usage of the entire system.

The micro-computer category includes standard IBM-compatible PC's and Macintoshes. By work-stations are meant the low end of the scientific and engineering workstations manufactured by Sun, Apollo, etc. The designation of supercomputer refers to the largest and fastest computers currently available, such as the Cray XMP or Cray II. The category of mid-range computer, then, refers to everything in between these poles.

Equipment

In considering computing at the level of a micro-computer or workstation, the subcommittee decided not to limit its discussion to research needs. This was done because it was recognized that such computers are general productivity tools, and the same machines are frequently used for administration and education as well as research. Recognizing how indispensable such computers have become over the past few years, the subcommittee makes the following recommendation.

> **Every faculty member and student should have access to a microcomputer/workstation environment meeting at least specified minimum standards.**

The minimal environment should include a computer in the office of every faculty member who can demonstrate a need. The subcommittee fully expects that most faculty will be able to demonstrate such a need. Clusters of computers should be available to students who are involved in research. The latter is to some extent already in place. The provision of computers for faculty lags behind what is available at comparable universities.

In the middle range of computing, there are two rather distinct modes. The first is a well focussed research group with specific computational needs that can best be met by a computer dedicated to that purpose. These machines have usually been purchased by funds from research grants, with some matching funds being provided by the university. This method has worked well for the University and its faculty in the past, and should be continued in the future.

> **Rice University should continue to encourage the purchase of local mid-range computers by providing matching funds to supplement research grant funds.**

The other mode of mid-range computing is done on powerful computers available to users across the campus. No attempt was made to quantify the exact needs in this area, but enough research was done to determine that there are significant unmet needs. Furthermore the current campus mainframe was installed in 1982 and is now out of date.

> **The campus-wide mid-range computing capability should be substantially upgraded. This upgrade should be done in such a way as to provide for the disparate needs of data-intensive and computation-intensive computing.**

There are increasing numbers of researchers on campus for whom the use of a supercomputer is essential. Given its size, Rice will probably not be in a position to purchase its own supercomputer(s) as an increasing number of larger universities are doing. Instead the university should arrange with various off campus supercomputer centers to make CPU time easily available to those who need it.

> **Every faculty research group which can establish a need should be provided with a minimum of 10 hours of CPU time on a supercomputer. A small number of projects should be provided with up to 50 hours.**

Computer graphics capabilities and visualization processes have rapidly evolved to the point of being an integral tool for all types and all phases of research. In recognition of this fact the university has started a central Advanced Visualization Lab. It is to be expected that the use of computer graphics and visualization will increase dramatically over the next five years in almost every computer environment on campus. The central laboratory will be essential to lead and facilitate this growth.

The Rice Advanced Visualization Laboratory should be provided with state-of-the-art equipment to enable it to provide needed service to the entire campus.

Networking

A network connecting all components of the distributed computing system is absolutely essential to its efficient operation. Its importance is emphasized in each of the discussion sections that follow.

A high speed network should be installed to connect all research computers on campus and provide access to regional and national networks.

The computers in separate units on campus should be connected in local area networks (LAN's), and these networks should be connected, via gateways, to the university backbone. Higher speed networks should be provided where high band-width data is transferred.

Support

The maintenance and support of the computer services on campus is of critical importance. At every level the need for increased support staff is already being felt. The support staff presently available is inadequate for current needs.

It is recommended that the total support staff be increased by approximately 20 full time positions to support the needs of the enhanced computer environment recommended in this report.

A distributed computing environment will be best served by a distributed support staff. It is recommended that much of the increased staff have primary responsibility to departments or groups of departments. This is especially true for staff supporting micro-computers, workstations, and local midrange computers. Such a distributed support staff will be more familiar with the needs of the users and will be able to respond more quickly and more adequately to these needs.

On the other hand, there are some advantages to be obtained from a degree of central coordination. Through a centrally coordinated staff, the university could better achieve standardization and the resulting savings. In addition, researchers should have a central source of information about machine capabilities and costs. Activities such as networking will need the continued support of a central staff. Consequently, the need for the Computer Resource Center will continue, but its responsibilities should be changed to include only those activities which are campus wide in scope.

Finally there is one area in which a new centralized support facility is required. Good research in almost all experimental work in the physical sciences and engineering requires computer-based data acquisition and analysis (CBDAA). In fact, CBDAA is now so widely available, and the enhancement of capabilities that it provides is so substantial, that hardly anyone can claim to be doing competent experimental work unless he or she is using it. In addition, CBDAA is becoming an essential feature both in graduate and undergraduate education in the physical sciences and engineering. It has become essential because most Rice graduates in these disciplines will be employed in environments where they must use CBDAA or manage projects whose success depends on its application. The widespread use of these techniques on campus indicates that a small, centrally administered, group of support personnel could provide increased efficiency.

A central support staff should be provided to support computer based data acquisition and analysis.

This group should provide instruction in CBDAA (both to users and undergraduates), diagnose and fix difficult CBDAA system bugs, write software, install equipment, provide sound advice, and implement informal standards.

Space

The proposals submitted here involve substantial additions of personnel and equipment and therefore require a substantial increase of space. In particular, there will be a need for a secure equipment room in or near every office building that houses servers for local area networks, and space will have to be provided for printers in these buildings. Office space will be required for the proposed additional support staff, and since much of this staff will be distributed, space will be needed in many buildings on campus. In some cases space already exists or can be easily provided, but other buildings are so crowded that such space will be hard to find. A systematic survey will have to be made in order to assess the additional space need created by this proposal.

Costs

The costs in the tables below are estimates. It is assumed here that the university will invest in two campus wide mid-range computers. The first is a campus mainframe which is modeled on an IBM 3090 model 180S, intended to replace the existing AS /9000. The second is a mini-supercomputer which is modeled on a Convex 220. It deserves to be repeated that the subcommittee did not make a precise determination of the needs for campus wide mid-range computing.

These estimates do not take into account the space requirements for either the equipment or the support personnel recommended here.

Capital Costs

Microcomputers/workstations	$1,500,000
Campus mainframe	3,200,000
Mini-super computer	2,000,000
Rice Advanced Visualization Laboratory	600,000
Networking	700,000
Total initial costs	$8,000,000

Table 1.

The costs in the following Table in FY 90 dollars are incremental costs over what is currently being spent by Rice University. They represent steady state costs after the equipment has been installed.

Annual Costs

	Maintenance	Support	Replacement	Totals
Microcomputers/workstations	$150,000	$400,000	$300,000	$850,000
Networking	50,000	100,000	160,000	310,000
Local midrange computers		100,000	100,000	200,000
Campus mainframe			640,000	640,000
Mini-supercomputer	200,000	50,000	400,000	650,000
Supercomputing		100,000		100,000
Visualization Laboratory	90,000	150,000	120,000	360,000
Laboratory/Computer Workshop		150,000		150,000
Totals	$490,000	$1,050,000	$1,720,000	$3,260,000

Table 2.

Microcomputer/workstation requirements

Introduction

In considering computing at the level of a microcomputer or a workstation, the committee decided not to limit its discussion to research needs. This was done because it was recognized that such computers are general productivity tools and the same machines are frequently used for administration and education as well as research.

Personal microcomputers or more powerful independent workstations are now indispensable tools for research and teaching in every academic field. The trend at first rate colleges and universities is to recognize this situation by providing small scale computers to each faculty member. In order for Rice to remain competitive with comparable institutions, it is imperative that the minimal computing needs of all faculty members be met through a carefully enacted policy of equipment- and software-acquisition, maintenance, upgrading, replacement, and support.

While no attempt was made to poll other universities systematically, the following examples came to the attention of the committee. Notre Dame has begun to implement a program of providing every faculty member with a microcomputer or workstation. At Princeton, a faculty member who can demonstrate a need will be provided with such equipment. Lehigh provided each of its faculty members with a Zenith microcomputer. Brown and Harvard also provide their faculty with microcomputer/workstation capability.

Experience at Rice and other universities has shown that faculty with access to microcomputers and workstations are more productive. Using computers, an increasing number of faculty are able to prepare their own exams, course syllabi, course instructions, manuscripts, letters, memos, proposals, reports and budgets, rapidly update and maintain lists of references and image libraries, develop research databases, perform library searches, communicate with colleagues across the country, analyze digitized graphic images, and even compose music. Several illustrations of the improved quality of research have been discussed in Appendix A of the Report of this subcommittee in Computing Goals for Rice University.

The use of computers increases personal efficiency and improves teaching through faster and better class preparation, including preparation of illustrations, and simulating and improving communication and rapport with students. At some universities, bringing laptop computers to class and turning in exams, papers and thesis drafts on floppy disks has become common practice. The instructor comments directly -- and often more extensively -- on the disk, and the student easily incorporates the teacher's suggestions or criticisms into subsequent versions.

In light of these gains, we believe the cost of equipping the entire faculty with adequate computing facilities will pay long-term dividends far in excess of the investment itself.

Personal computers are now so commonplace that some government granting agencies treat them as office equipment, a category of equipment that must be purchased with either university or personal funds. Small computers to carry out essential university tasks are thus becoming more and more difficult to acquire under research grants, even for faculty members in fields where research is heavily subsidized through external funds.

For faculty members in the humanities and in some of the social sciences - where equipment grants have always been scarce - there is almost no chance for an individual faculty member to obtain even minimal computational equipment through non-university sources. The net result has been that, in these fields, faculty members have often paid for microcomputers themselves. This has led to a proliferation of incompatible software and hardware. Further, attracting and keeping junior faculty members is becoming more difficult. Departments and divisions must, therefore, be able to assure talented young researchers and teachers that Rice will meet the computing "packages" offered to them as a matter of course at comparable institutions.

For all these reasons, we propose that all Rice University faculty members who can demonstrate the need should have direct, immediate, and routine access to the computational hardware, software, and communication facilities essential for carrying out the teaching, research, administrative, and scholarly activities that are expected of them.

Equipment Requirements

The needs of individual faculty members will depend on their particular situations. However, it is important that the equipment and software be as compatible as possible.

Standard hardware and software packages are more economical to acquire and support, and much more convenient to use. Planning for a variety of needs, while maintaining a basic policy of standardization, requires a plan with flexibility as well as basic standards for equipment. To accomplish this, we have assembled a list of minimal microcomputer equipment- and software-packages. We have also planned for faculty whose minimal needs are for a more computationally intensive environment than is currently available in personal computers, for example, faculty members engaged in symbolic, numerical, or digitized-image computing, applications in which speed is important or color graphics are needed.

In order to accommodate the range of the faculty's basic computer needs, we have provided two sets of hardware requirements. The Standard Package addresses the most essential needs of the microcomputer user; the Extended Package addresses microcomputer upgrades and workstation configurations.

While it is impossible, at present, to specify the available machine configurations and prices at the time when final purchase decisions are made, we have prepared initial cost estimates based on the following assumptions about hardware configurations. These packages are considered to be typical of the needs of faculty members, but they are used here primarily to estimate costs. It is recognized that what individual faculty members actually need may vary.

1. The Standard Package will consist of a microcomputer with the following configuration:

   a. Computing power equivalent to that of a Macintosh SE (8 MHz), a high-speed IBM PC compatible, or a comparable laptop. The computer should have at least 2.5 megabytes of RAM if it is a Macintosh, or at least 1 megabyte of RAM if it is a DOS machine. It should have at least 40 megabytes of hard drive storage and at least one floppy disk drive. One of these several products currently available will satisfy this minimum requirement: the Macintosh SE; an IBM PS/2 Model 50 or 60, or Compaq PC or compatible of the 386 class; or a comparable laptop.

b. Access to a conveniently located, compatible, laser printer.

c. Connection to the Campus-wide backbone. Cost-equivalent options include an ethernet card, ROLM switch DTI, or 9600 baud modem. Eventually the connection should be via a local area network.

d. Access to hard disk tape back-up facilities located in individual departments.

2. The Extended Package will consist of the following classes of equipment:

a. A powerful workstation such as a SUN, a NeXT, or a similar machine equipped with color monitor, if necessary, and connected to a powerful highspeed fileserver.

or:

An upgrade to a higher performance microcomputer and/or additional storage, etc.

b. Access to Postscript supporting laser printers.

c. Ethernet connection to Campus network.

3. In addition to basic computational equipment, Rice faculty members need convenient access to a variety of state-of-the-art computer technologies for carrying out specialized research, presentation, and class- and course-ware preparation tasks. Such equipment might include:

a. Mass storage add-ons including large hard disks, CD ROMs and WORMs available for checkout from a central facility.

b. Special-purpose boards and disk drives available for checkout.

c. Satellite data antennas, modems, and boards for business, wire service, and other special needs.

d. Portable computers for use on travel, in the field, or in libraries or local archives.

e. Special purpose computers such as the Yamaha C1 Music Computer designed for composing and creating music.

f. Access to high-speed, state-of-the-art, optical text and image scanners (Kurtzweil or equivalent), both hand-held and PC based, through a centralized facility.

g. Digitized video processing in a state-of-the-art centralized facility staffed with technical personnel.

h. Color slide production equipment and services.

i. Computer projection equipment.

Software

The following software should be available as needed:

1. Word processing
2. Spreadsheet
3. Graphics
4. Database
5. Communications software for electronic mail and access to external data bases
6. Statistical package
7. Courseware authoring packages
8. Programming language compiler
9. Special purpose software such as music composition and foreign language applications, business software, desktop publishing with special font capability etc. should be available for easy access.
10. Utility software, including system management, disk backup, etc.

The university should arrange for site-licensing or group purchase of software whenever possible.

Networking

The effectiveness of microcomputers and workstations is greatly enhanced when they are integrated into an electronic network. As soon as possible, every department should be connected in a local area network (LAN) which is connected to the university backbone with access to the Rice mainframe, centralized file servers, and national computer networks.

Furthermore, Rice should adopt a standard for departmental LANs to provide a commonality in support and maintenance.

Hardware Maintenance and Software Upgrading

The university should provide facilities, staff, and a permanent fund for servicing the system of computers through service agreements (time and materials) and/or on-campus maintenance facilities as appropriate. The current cost of maintaining a microcomputer averages about $200 per year. Service for a workstation will average about $500 per year. These estimates include maintenance for the LAN. An additional $300 per unit per year is required for upgrades to the various microcomputer software packages, and $500 for the workstation software packages.

To the extent possible the university should rely on external funding to maintain computers. Since this will not be possible in many situations, however, the university will have to assume a share of the burden. There is also the troublesome middle ground where external funding is available to purchase computers, but funding for maintenance is not available. In these situations, the university should be aware that it is accepting the responsibility for the costs of maintenance at the time that the machines are obtained.

Support

Initial purchase and installation of basic hardware and software is the first step towards meeting minimal computing needs. Of equal importance is providing the support services so that these tools can be used effectively and efficiently.

The goals of the computer support plan are:

* To assure within a reasonable time framework an effective level of computer competence for all interested faculty.

* To provide readily accessible consulting and maintenance services to all faculty.

* To create a responsive and effective computer support infrastructure.

Elements of support:

Some of the more obvious projected support needs include:

* Advising faculty and staff in choosing appropriate hardware and software.

* Training faculty and staff users to operate new equipment and installing software upgrades as they become available.

* Dealing on a daily basis with computer-related problems.

* Providing technical help in installing and maintaining computing and network equipment.

* Supplying accessible documentation for the various software and hardware packages.

* Keeping faculty and staff informed of developments in computing which could be particularly helpful in carrying out their specific research and teaching tasks.

The intelligent use of computers in the educational process will require a different level of support than do other uses of computation. The development of effective computer courseware is a creative process, every bit as taxing on the imagination and intellect as any other research project that goes on in the university. Some of the support personnel should be sophisticated enough to be partners of the faculty in this creative activity.

The currently available staff (about 10 FTE's) is not sufficient to support the existing microcomputers/workstations. An extensive expansion of support staff will be necessary to provide even minimal support in the areas outlined above. A reasonable estimate of the direct support staff that will be needed is about ten additional FTE's.

A staged plan for the enhancement of support staff should be in place as quickly as possible. Because of the difficulties involved in starting up any computing enterprise, it is important that adequate staff should be available before the equipment is installed.

Because the different disciplines have very different computing needs, we also recommend that most of the additional support personnel be familiar with the computational needs of specific fields and that staff be made available within departments, divisions, or perhaps within specific buildings. An optimal organizational model would have most of the support personnel officed with or near the people they serve. There should continue to be some form of central control, however, to coordinate policy, set standards across the university, and arrange for university-wide site licenses.

Replacement

A commitment to provide the Rice faculty with computational tools for carrying out their contractual obligation to teach, do research and perform university service will not end with an initial purchase of equipment and support for those machines and software packages. The useful life of a computer will vary depending on the application. A computer used primarily for word processing and data retrieval may have a life of seven years; a computer used primarily for number crunching might need replacement after only three years. Accordingly, the university should plan to turn over equipment on a revolving basis, replacing 20% of the equipment in a given year. Monies for replacements of hardware and software must be a regular line-item included in the university budget. These considerations, however, also lead to the conclusion that a computer outdated in one situation might still be viable in another. By judicious redistribution of equipment, the university should be able to achieve significant savings.

Cost

In order to estimate the cost of the equipment and software needed to implement the proposed plan, we have conducted a survey of the departments. We conclude that the initial cost will be about $1.5 million and that the annual cost of maintenance, support, and replacement is expected to be $900,000. These costs are detailed in Table 3 below. The costs of local area networks are not included in these figures; they are estimated separately in the section on Networks later in this report.

Implementation

The subcommittee suggests that the proposed plan be implemented over a three-year period. The distribution of equipment should be done through the university's standard administrative channels. The academic Deans are in the best position to judge the microcomputing needs of the faculty in their divisions. In addition, they should be able to assess what other resources are available to their division and thereby optimize the use of the university's own resources. Every effort should be made to seek out external sources of support for the purchase and maintenance of computer equipment. Frequently, however, external support will not be available, and the university will have use its own resources in order to pursue its agenda.

Costs of Microcomputers and Workstations

Capital cost of equipment and software $1,500,000		
Annual costs		
Maintenance (micros) 220 computers x $200/computer/year	$44,000	
Maintenance (workstations) 40 computers x $500/computer/year	20,000	
Software upgrade (micros) 220 computers x $300/computer/year	66,000	
Software upgrade (workstations) 40 computers x $500/computer/year	20,000	
Total annual maintenance	$150,000	
Equipment replacement cost 20% of capital cost/year	300,000	
Annual support cost	400,000	
Total annual costs		$850,000

Table 3.

Mid-range computing requirements in research

Virtually all faculty members need access to personal or workstation-level computing to support their research. For a smaller, but significant number of faculty across the campus, computing is an essential part of their research. For some, supercomputing is needed; their requirements are addressed in a later section. A much larger number do not need super-computing power, but require facilities more powerful than workstations. For example, ionospheric physicists want to simulate the characteristics of solar radiation on plasmas; atmospheric chemists want to model the Antarctic ozone hole; economists want to assimilate large amounts of data to predict economic trends.

These mid-range computers are typically accessed over a network, with a workstation or personal computer serving as the user's window to the larger machine. They have significantly greater computing power than a workstation and support software systems capable of dealing efficiently with large amounts of data. Rarely do such machines need to be solely owned; they are typically shared among a number of researchers. However the extent to which they are shared differs. In one typical model, a local computer is shared by a small group of users which has specific computational needs that can best be met by a computer dedicated to that purpose. In another typical model, computers (such as the current campus mainframe - an AS/9000) are shared by users from across the campus. Both of these environments will be discussed in this document.

Another important consideration is the type of computational problem and the correspondingly appropriate computer architecture. For example, there are some problems which are data intensive and require a computer which has a large memory space and high speed I/O. IBM mainframe computers (and clones thereof) have these features, but many computers do not. Other computational problems, which might be labelled computationally intensive, do not require large data files, and for these the best computers are those which feature speed of computation, perhaps enhanced by vectorization and parallelism.

Local mid-range computers

Since the installation of the AS/9000 in 1982, there has been a significant expansion in the number of local mid-range computers available on campus, whose combined power is now several times that of the AS/9000. These machines (including VAX's, large workstations manufactured by Sun and other companies, one Sequent Symmetry, and many others) serve groups of researchers with similar computational needs. As it happens, the users are usually members of a single department although, because of networking, physical proximity is not required.

Local computers allow their users to define their own computing environment and give them easy access to the power of that environment. These advantages are sufficiently important to offset the power limitations inherent in local mid-range computing. For many researchers on campus, especially for those whose problems are computationally intense, local computing is the environment of choice.

The local mid-range machines currently on campus have all been purchased by funds from research grants, with some matching funds being provided by the university. This method has worked to the advantage of the university and its faculty and should be continued in the future.

The university should continue to encourage the purchase of local mid-range computers by providing matching funds to supplement research grant funds.

More critical than obtaining a computer is maintaining it. Most modern computers have few failures of either the central processor or the peripherals. Failures can be so infrequent that researchers forget how capricious a computer can be. Yet, they do occur and then render long-running calculations impossible or destroy valuable data. Because these are shared facilities, a support staff is essential for routine computer maintenance and system support. Staff costs can sometimes be borne by grants and contracts, but these monies are usually difficult to find. Furthermore better qualified personnel can be found if the position is not based on "soft-money".

Economies of scale can occur when diverse research groups share support staff for similar computers. Furthermore, researchers considering writing equipment proposals should have a central source for information on machine capabilities, machine costs (purchase and maintenance), and the university's system for approving matching funds.

For all of these reasons, the subcommittee makes the following recommendation.

> **The responsibility of the Computer Resource Center for mid-range computing should be increased and its staff enlarged appropriately.**

Campus wide mid-range computers

The other mode of mid-range computing is that done on large computers available to users across the campus. Presently, this category includes the campus mainframe, the AS/9000, and an IBM 3081 that is available for limited use. In addition, Rice offers special opportunities for high-performance computing in the Center for Research on Parallel Computation, the result of a recent Science and Technology Center award from the National Science Foundation.

Computers for campus-wide use will always be necessary in a distributed computing system for the following reasons:

* Data intensive problems require a computer with large memory space and high speed I/O, features not available on smaller machines.

* There are users without access to an adequate local mid-range computer who need to solve problems beyond the capability of a micro. A central campus computer is essential to meet their needs.

* There are problems which are simply too large for the capabilities of a local mid-range computer.

The first reason indicates that the university must have a sufficiently powerful computer with an architecture capable of coping efficiently with data intensive problems. The current campus mainframe, the AS/9000, has the correct architecture, as do most IBM compatible mainframes. The AS/9000 was installed in 1982: it is already outdated and will become more so over the next five years.

The third reason listed above applies to all fields. A survey of users on the Rice campus indicates that there is substantial unmet demand for campus-wide mid-range computing. Much of this demand involves computationally intensive problems.

It is clear that, in the five year period that is covered by this plan, there should be a substantial increase in campus wide computer capability. No attempt was made to assess more quantitatively the need for campus wide mid-range computers. Consequently, the subcommittee will make a general recommendation, and will present possible solutions.

> **The campus wide mid-range computing capability should be substantially upgraded. This upgrade should be done in such a way as to provide for the disparate needs of data intensive and computationally intensive projects.**

At the present time, the only computers with the features needed for data intensive computing seem to be IBM compatible mainframes. Accordingly, to meet this demand, the current mainframe computer should be replaced with another of more modern design, having an architecture that will enable it to cope efficiently with data intensive problems. Furthermore, choosing the IBM architecture would ensure the continued availability of special software which is currently in use on such machines, and the similarity with the current mainframe would provide continuity and ease the transfer to the new computer.

Such a machine might turn out to be suitable for the administrative needs of the university as well, and this combination of needs might make it possible to have on campus a more powerful research computer.

The needs for enhanced capability for dealing with computationally intense projects can then be met in at least two ways. First, the mainframe computer needed for data intensive computing could be enhanced by adding vectorization capability, a UNIX operating environment under the IBM operating system, and possibly a second CPU and more memory. Second, a separate computer more suited to computationally-intense problems could be provided (such as a Convex or an Alliant). The second solution would definitely be preferable to some users, but the relative costs of these two options should be examined.

Networking

Typically a user will use a mid-range computer by logging on remotely from a workstation, even if the computer is located next door. The benefits, in terms of time saved and speed of turn-around, provided by the interactive computing made possible by this procedure cannot be overestimated. In some cases, especially those involving graphics, visualization, or the transfer of large data sets, higher-speed networks will be required. For the most data-intensive problems, the university backbone might be a bottleneck. Thus, a high-speed network connecting the computers on campus is essential.

Supercomputing Requirements in Research

Introduction

The research specialties of an increasing number of Rice faculty involve frontier applications of computing. For these researchers, the complexity of the mathematical problems to be solved and/or the sheer volume of empirical data to be modelled dictate that advances in research are limited primarily by the boundaries of computing resources available. Thus, state-of-the-art computing capability, often called "high performance computing" or "supercomputing," is an absolute requirement for the performance of competitive research by these faculty and correspondingly for the education and training of our students in these areas.

The faculty's diverse and significant use of supercomputing resources is documented in Table 4 which provides a nearly complete list of current Rice faculty, staff, and students with off-campus supercomputing accounts. That the need for these resources is expanding is documented by the recent external funding proposal prepared by the faculty for the establishment of a High Performance Computing Center managed and utilized by the Rice community. Table 5 provides a list of potential users from this proposal.

A "supercomputing facility" provides computing capabilities on a par with the fastest, largest, and most advanced machines in existence. As such, the characteristics of a supercomputer must be constantly redefined. Currently, the defining characteristics include:

* A nominal computing rate of at least one billion arithmetic operations per second (more technically, 1 "Gigaflop") or, alternatively, an actual benchmark speed no less than a Cray XMP on one or more of the standard tests. These rates can only be achieved through efficient vector processing and/or parallel processing.

* A direct access memory of at least 16 million words (128 Mb), indirect rapid access memory of at least 128 million words (1 Gb) and 20 Gb or more disk storage.

* A mid-range "front end" computer to provide efficient access among interactive users on individual workstations and batch processing on the supercomputer.

* A purchase price on the order of 10 to 20 million dollars.

Additionally, a supercomputing facility includes as vital components:

* A high-speed network for the efficient transfer of very large data sets, often involving billions of words.

* A support staff knowledgeable in networking, vector and parallel programming, and advanced operating systems.

Presently, Rice faculty are largely independently responsible for finding off-campus supercomputing time. The listing in Table 4 is testimony to their success. As is clear, virtually all of Rice's supercomputing time is afforded by grants from three sources: the National Science Foundation's Supercomputing Centers, the Cray Corporation, and the Houston Area Research Center. Such unrestricted dependence on external grants of free time is risky, however. First, the NSF sites are rapidly becoming saturated with requests for CPU time from universities throughout the country; easily available time will soon be a thing of the past. Second, the latter two grantors of time are essentially advertising their hardware through their grant programs. While such "free" opportunities are excellent and must be pursued, continued reliance on such programs without any commitment of revenues from Rice could lead to a sudden and uncontrolled loss of the available resources. Third, Rice students and faculty will find it increasingly difficult to compete with their colleagues at most other major universities, where, a survey reveals, faculty research groups are provided large amounts of supercomputing time at no cost. We conclude that Rice must expect to become a cost-sharing partner in the acquisition of grants of CPU time.

Moreover, a comparison of the list in Table 4 to that in Table 5 reveals that there are many important supercomputing research projects which could have been performed but have not been due to lack of supercomputer access. The primary barriers to access are: (1) lack of available resources on campus provided by the university; (2) lack of funding opportunities for purchasing computer time for research in these areas; (3) lack of computing support expertise required for connecting to, and operating, the high performance machines offered by Cray and HARC. Because of these problems opportunities to enhance Rice's position as a premier research university have been lost. Prospective and current faculty must have access to supercomputing technology if Rice is to remain competitive in computer-related research areas.

The conclusion of the subcommittee is that:

The university must commit itself to supporting supercomputing research by providing to the Rice students and faculty:

1. **Grants, in the form of CPU time, for access to supercomputing facilities.**

2. **Staff personnel with expertise to support supercomputing research.**

3. **Networking.**

Supercomputing Access

Providing these resources to the Rice community is obviously problematic. In order to remain at the leading edge, however, Rice must be prepared to keep pace with constant advances in technology and software. To this end, we provide a set of alternative recommendations for supercomputer access, based on the underlying recommendation that Rice position itself to seize opportunities to avail itself of any of these alternatives.

In this regard, Rice offers unique opportunities for high-performance computing users. The Center for Research on Parallel Computation will provide state-of-the-art parallel computing capability to researchers at Rice.

The university should establish as a goal the provision of a minimum of 10 hours of CPU time on a supercomputing facility to every faculty research group which requests it and can establish a need for it. A further goal should be to support a small number (perhaps one per division) of larger projects each year at a level of 50 hours of CPU time. Initially, the university should anticipate the need to provide 600 hours of CPU time to meet these two goals.

This CPU time could be acquired in a variety of ways that are not mutually exclusive. The configuration of options ultimately selected will depend, in part, on decisions about university support for mid-range computing. However, it is the subcommittees' feeling that any combination of options must address the needs of the research faculty across all divisions.

Option 1:

The university could expand current agreements with the Cornell University Supercomputing Center and with HARC through which CPU time is made available to researchers at Rice. Rice could also pursue agreements with other NSF-supported supercomputer centers which are now being upgraded. The required 600 hours of CPU time might be obtainable at minimal cost in this way.

Option 2:

Another arrangement would be to purchase a mid-range supercomputer such as an IBM 3090 or to upgrade to vector and multiprocessor capability the mainframe the university decides to acquire. Researchers with constant access to the most recent technology would then continue to direct their high-end use at supercomputing centers off campus through funded research grants. This arrangement might also diminish the need for a mid-range computing center.

Option 3:

A third scenario involves the joint purchase or lease of a supercomputer in a consortium of schools (for example, HARC). The supercomputing center could be at Rice. Presumably, such an arrangement would give us more control over supercomputing resources and would provide direct access to them. Moreover, the administration and staff of our mid-range computing center could also oversee the administration and staffing of the supercomputer center. Although it is impossible to estimate the cost to Rice in such a scenario, since the number of consortium members is unknown, the cost to the consortium as a group would be an initial capital outlay in excess of $20 million and a yearly operating expense of $4 - 6 million.

Support Personnel

Although the <u>nature</u> of technical support varies with the level of computing, the <u>need</u> for such support is independent of the level. Even for faculty researchers whose technical expertise in computing is above the norm, the sophistication of the operating systems and the complexity of the programming methods involved in supercomputing demand that the researcher either be a computer scientist, collaborate with one, or become one. Furthermore, introducing newer faculty with exciting and innovative research ideas to the capabilities of high performance computers requires training that is not provided in any existing facility at Rice. Finally, given that we are discussing state-of-the-art technical capabilities, if the university is to remain competitive in research, it will need staff expertise in this rapidly advancing field.

The university must commit itself to having a minimum of three permanent staff personnel with expertise in the access and use of high performance computers.

Networking

The university's needs, in terms of hardware, software, and personnel, for a supporting electronic network are documented elsewhere in this report.

We emphasize here that reasonable use of supercomputing is not possible without such a network.

Rice Supercomputer Users

Name	Department	Title	Source of Support
Kennedy, Kenneth	Computer Science	Professor	NSF
Callahan, David	Computer Science	Research Associate	NSF
Balasundaram, Vasanth	Computer Science	Graduate Student	NSF
Kalem, Marina H.	Computer Science	Graduate Student	NSF
Rosene, Carl	Computer Science	Graduate Student	NSF
Subhlok, Jaspal	Computer Science	Graduate Student	NSF
O'Dell, Charles R.	Space Physics	Professor	NSF
Wolf, Richard A.	Space Physics	Professor	NSF
Hill, Thomas	Space Physics	Research Scientist	NSF
Reiff, Patricia	Space Physics	Research Scientist	NSF
Spiro, Robert	Space Physics	Research Scientist	NSF
Bergman, Rachelle	Space Physics	Research Associate	NSF
Bales, Bryan	Space Physics	Research Associate	NSF
Kinsey, James	Natural Sciences	Dean	NSF
Wheeler, Mary F.	Math Science	Professor	NSF, Cray
Symes, Bill W.	Math Science	Professor	NSF
Lane, Neal F.	Physics	Provost	DOE
Hayes, Edward F.	Chemistry	Vice President	LANL
Chacko, Susan			NSF
Min, Kyoung W.			NSF
Williams, James			NSF
Richardson, J.R.	ICSA	Consultant	Cray
Spanos, P.D.	MEMS	Professor	Cray
Ghanem, Roger	MEMS	Graduate Student	Cray
Dessler, Alex	Space Physics	Professor	Cray
Bergman, Rachelle	Space Physics	Research Associate	Cray
Cheatham, John	Civil Engineering	Professor	Cray
Lin, Yu-Hsu	Civil Engineering	Research Associate	Cray

Hill, Tom	Space Physics	Professor	Cray
Toffoletto, Frank	Space Physics	Research Associate	Cray
Miettenin, Hannu	Physics	Professor	Cray
Hutchinson, J.S.	Chemistry	Assistant Professor	HARC
Sickles, Robin	Economics	Professor	HARC

Table 4

Potential Supercomputer Users

Researcher	Department	Research Interest
Ed Akin	Mechanical Eng.	Finite Element Analysis
Guy T. Almes	Computer Science	Distributed Computation
Athanasios C. Antoulas	Elec. & Computer Eng.	Control Theory
Philip B. Bedient	Environ. Science & Eng.	Hydrology
Robert E. Bixby	Mathematical Sciences	Combinational Optimization
Hans J. Boehm	Computer Science	Programming Languages
Andrew Boyd	Mathematical Sciences	Combinatorial Optimization
Sidney C. Burrus	Elec. & Computer Eng	Signal Processing
Robert S. Cartwright	Computer Science	Specification Languages
Nancy J. Cooke	Psychology	System Interfaces
Keith Cooper	Computer Science	Languages; Compilers
John Dennis	Mathematical Sciences	Nonlinear Optimization
John W. Freeman	Space Phys & Astronomy	Satellite System Modeling
Peter R. Hartley	Economics	Economic Modeling
Robert T. Hood	Computer Science	Programming Environments
William C. Howell	Psychology	System Interfaces
John S. Hutchinson	Chemistry	Quantum Mechanics
Bruce Johnson	Chemistry	Photodissociation
Don H. Johnson	Elec. & Computer Eng.	Stat. Signal Processing
J.R. Jump	Elec. & Computer Eng.	Parallel Architecture
Ken Kennedy	Computer Science	Optimizing Compilers
James L. Kinsey	Chemistry	Chemical Physics
David Lane	Psychology	User Interface Design
Neal Lane	Physics	Theoretical Physics
Tamara Ledley	Space Physics	Climatology
Alan R. Levander	Geology & Geophysics	Seismology
Clarence A. Miller	Chemical Eng.	Fluid Mechanics
J.B. Pearson	Elec. & Computer Eng.	Control Theory
George N. Phillips	Biochemistry	Computational Biochemistry
David Scott	Statistics	Computer Graphics

Gustavio Scuseria	Chemistry	Quantum Mechanics
Robin C. Sickles	Economics	Nonlinear Estimation in Econ.
J.B. Sinclair	Elec. & Computer Eng.	Parallel Architecture
Pol D. Spanos	Civil Eng/Mech Eng	Random Vibrations
William Symes	Mathematical Sciences	Geophysics
Richard Tapia	Mathematical Sciences	Nonlinear Optimization
James R. Thompson	Statistics	Parallel Statistical Analysis
Linda Torczon	Computer Science	Interprocedural Analysis
Virginia Torczon	Mathematical Sciences	Numerical Optimization
Gerd-Hannes Voigt	Center for Space Physics	Magneto Hydrodynamics
Mary F. Wheeler	Mathematical Sciences	Numerical PDEs
Karen Williamson	Mathematical Sciences	Optimization Theory
Richard A. Wolf	Space Phys & Astronomy	Magnetospheric Modeling
Willy Zwaenepoel	Computer Science	Distributed Computation
Kyriacos Zygourakis	Chemical Eng.	Chemical Reactor Design

Table 5.

Networking

All modern computers above the level of microcomputers are multi-user machines and are generally accessed remotely, using an electronic network. This is true whether the computer is a small mid-range computer in the next room or a supercomputer thousands of miles away. The rapid transfer of large documents across thousands of miles via an electronic network enables researchers at different locations to collaborate more effectively. Electronic networking allows researchers at Rice to access large data bases located at other institutions. These examples make it clear that modern electronic networking is changing the way research is done. High-speed communication via computer networks has become a necessity for almost all researchers in the sciences, social sciences, and engineering.

The benefits of a network to research at the university are matched by those to education and administration. The subcommittee makes the following recommendation.

A high speed network should be installed to connect all research computers on campus and provide access to regional and national networks.

The core of the network at Rice is the university backbone. This is a high speed fibre-optic cable which is now partially in place. Eventually it should provide easy access to all buildings on campus. Completion of the backbone should have high priority. The backbone should have sufficient capability to provide for the high band width required in those circumstances where large data sets must be transferred.

Every department should be connected in a local area network (LAN) which is connected to the university backbone. This will provide the following services:

1. Electronic mail
2. The capability of sharing software and files
3. High quality shared printing on local laser printers
4. Data back-up facilities
5. Access to the library
6. Access to administrative and academic services
7. Access to fileservers where necessary
8. Remote access to other computers
9. Access to external databases

Although AppleShare networks will suffice in some areas, the university should aim at ethernet networks being the standard within five years, in accordance with developments by the computer manufacturers. The local area networks should have a dedicated file server with a large hard disk to provide enough central storage for networked software and common files. The LAN should be so configured as to be accessible by both IBM-type and Macintosh computers. Where it is physically possible, each department should be on one LAN, but a LAN could be shared by several departments. In some cases it might be appropriate for an entire building to be on one network. Special equipment will be needed when high bandwidth data is transferred on a network.

It is estimated that the initial connection cost will average about $1400 per machine for a LAN with 15 workstations together with the gateway to the university backbone. The total incremental initial cost of LAN's across the campus is estimated to be $360,000. The cost of extending the university backbone to the point where there is convenient access to every building on campus is estimated to be $340,000. The annual costs of maintenance for both the LANs and the backbone is estimated to be $50,000.

A centrally located staff will be required for support of the backbone and to handle major maintenance on the LANs. The distributed support staff will have responsibility for the LANs. However, to ensure the adequate performance, each LAN should have a network coordinator whose principal functions would be routine tasks, such as back-ups, maintaining users lists, and performing relatively routine maintenance. In many cases, such a coordinator might be able to serve up to three or four networks, but it is important that the coordinator be available locally.

Computer Graphics and Visualization

Introduction

Computer graphics and visualization capabilities have evolved rapidly to being an integral tool for all types and all phases of research. Access to state-of-the-art visualization capabilities will greatly enhance individual productivity and allow for better communication of research results to broader audiences. Visualization is an area of computing that allows results to be "seen", and thus expands the use of computing in research. University researchers need these capabilities to remain competitive in attracting research funding.

Visualization has some very important applications in teaching as well. As the importance of visual understanding is becoming recognized, educators increasingly use multi-medium teaching techniques. Many universities are moving rapidly into this new area, and their efforts are based on the concepts and techniques of visualization.

While individual faculty at Rice have long been involved in the use of computer graphics, the origins of any comprehensive efforts in this area can be traced to the May, 1985 Computer Committee Report titled "Toward a Long Range Plan for Computing at Rice University." Paragraph 3.4 identified the "great need for high quality graphics not only in engineering, the sciences, and architecture, but also in the Social Sciences and Humanities." All subsequent computer plans and committee reports have emphasized the need for computer graphics and, more recently, visualization.

Computer graphics comprises the area of tool-making for the display of visual images. Visualization is a newer term that applies to using graphic tools and visual imagery in problem investigation and problem solving. As visualization is a broader term that encompasses computer graphics, we will use it in this report.

A recent report, entitled "Visualization in Scientific Computing," published by the ACM SIGGRAPH (American Computing Machinery Association's Special Interest Group in Graphics) presents a comprehensive view of the rapid growth and importance of the area of visualization research support. Excerpts from that report are used here to answer the question, "What is *Visualization* ?"

> *Visualization is a method of computing. It transforms the symbolic into the geometric, enabling researchers to observe their simulations and computations. Visualization offers a method for seeing the unseen. It enriches the process of scientific discovery and fosters profound and unexpected insights. In many fields it is already revolutionizing the way scientists do science.*
>
> *Visualization embraces both image understanding and image synthesis. That is, visualization is a tool both for interpreting image data fed into a computer, and for generating images from complex multidimensional data sets.*

It studies those mechanisms in humans and computers which allow them in concert to perceive, use and communicate visual information. Visualization unifies the largely independent but convergent fields of:

* *Computer graphics*
* *Image processing*
* *Computer vision*
* *Computer-aided design*
* *Signal processing*
* *User interface studies*

One researcher observed many years ago that 'The purpose of computing is insight, not numbers.' The goal of visualization is to leverage existing scientific methods by providing new scientific insight through visual methods.

An estimated 50 percent of the brain's neurons are associated with vision. Visualization in scientific computing aims to put that neurological machinery to work.

Visualization represents a very rapid area of growth in computing and several Advanced Visualization Labs (AVLs) have been started recently in leading universities. Recognition of the importance of visualization has led to setting up the Rice Advanced Visualization Lab (RAVL). RAVL is only one part of the overall university needs in visualization, however. RAVL complements but does not overlap with the distributed computer graphic hardware and software needs detailed in the section of this report dealing with Microcomputer/Workstation requirements. RAVL would provide the centralized facilities and support called for above. Keeping all these needs in mind, this section of the report will concentrate on the role and needs of the central facility, RAVL.

Rice Advanced Visualization Lab (RAVL)

Although the goal is to make RAVL a central resource and service lab for teaching and research at Rice by providing state-of-the-art computational visualization equipment and expertise, it is not intended to be the only source of visualization equipment and capabilities at the university. In addition to RAVL, campus computer labs, faculty, and students must have easily accessible on-site resources available to them as well.

The role of RAVL as a central visualization facility can augment both use and access to visualization by performing at least the following functions:

* Provide staff with state-of-the-art computer graphic and visualization capabilities in this rapidly growing area.

* Provide access to equipment and software.

* Work with commercial vendors to demonstrate current equipment and techniques.

* Be a site for visualization equipment that is either too large or expensive for distributed computing.

* Provide multimedia capabilities to all Rice University departments.

RAVL will be a service facility. Visualization service can be sub-categorized into three levels on the basis of difficulty.

* Low Level: including the well defined areas of desk-top publishing, statistical graphics, CD-ROM and Laser Disk Courseware.

* Mid level: including Computer-Aided Design (CAD), Digital Video Interactive (DVI), and image processing.

* High level: this includes large scale data analysis, computationally intensive algorithms, and other super-computer associated visualization support.

Service must be provided through a variety of means including:

* Expert staff for support, consultation, training, teaching and advice for hardware and software purchases.

* Providing access to hardware and software.

* Preparation of hardcopy, slides and video (VHS).

In order to stay in contact with the university user community, communication channels must be put in place. These will include:

* Newsletters and announcements
* Lectures and presentations

* Demonstrations
* Reports and concept papers

Current Status of RAVL

RAVL, currently in the first phase of development, has already passed several important milestones:

1. A physical space to support computing has been created by remodeling room 218 in Anderson Hall. This type of space, difficult to obtain, is the essential first step to future success in this area.

2. An inventory of visualization hardware and software has been built up in connection with the Rice Architecture Computer Lab (RACL) which provides a start in this highly technical area of computing.

   Current equipment is minimal but consists of a Sun fileserver and one workstation, several PC/XT's that are not capable of supporting real graphics, and some older equipment on loan from a vendor as well as two sets of commercial graphic software for Computer-Aided Design (CAD)

3. An ongoing teaching effort involving RAVL is currently underway. All Architecture Computer-Aided Design (CAD) courses make use of RAVL, as do some short projects from other departments and two Continuing Studies courses.

4. An ongoing research effort is underway. Currently, an agreement with IBM is pending to develop an interface to the IBM implementation of the Programmer's Hierarchical Interactive Graphics System (PHIGS) named graPHIGS. When approved, this agreement will result in additional hardware:

   * 8 PS/2 50's
   * 1 PS/2 80

   And software:

   * graPHIGS for the IBM 3081
   * Networking - LAN and Token Ring
   * Speakeasy for the 3081

Beginning from this base, Rice is in good position to build up an excellent visualization facility. This enhancement effort will require expansion of existing personnel, hardware, and software resources.

Support Requirements

The efficient use of modern visualization techniques requires the advice of experts. Support staff is as important here as in any other area of computing. Currently, three part-time staff are struggling to start and maintain this entire area of activity. The immediate need is for three FTE positions in the RAVL, growing to at least four over the next five years.

Equipment and Software Requirements

Major areas of equipment and software needs are:

* Output devices
* Video equipment
* Computer in the "graphic engine" class of computers
* Sufficient number of workstations

Costs of the Visualization Laboratory

To date about $150,000 has been expended or set aside for equipment acquisition. With this the facility is at a minimum visualization capability. Support for RAVL needs to be continued and augmented over the next five years. The estimated costs are in the following table.

Capital cost of hardware and software	$600,000
Annual costs	
Maintenance	$90,000
Equipment replacement cost	120,000
Support staff	170,000

Total Annual Costs
$380,000

Table 6.

Computer-Based Data Acquisition and Analysis

Relevance

Like any university, Rice exists to educate students, produce new knowledge, and provide some services to the outside community.

Computer-based data acquisition and analysis (CBDAA) is becoming an essential feature both in graduate and undergraduate education in the physical sciences and engineering. It has become critically important because most Rice graduates in these disciplines will be employed in environments where they must themselves use CBDAA or where they must manage projects whose success depends on its application.

In almost all experimental work in the physical sciences and engineering, CBDAA is a requirement for an effective and efficient research program. CBDAA is now so widely available, and its enhancement of capabilities so substantial, that hardly anyone can be competitive in his or her experimental work without access to this capability.

At present, Rice has externally-funded research amounting to about $20,000,000 per year, at least 40% of which (>$8M, based on responses to a survey) depends significantly on computer-based data acquisition and/or analysis; most of that research would be severely limited without CBDAA.

Research groups with bona fide computer needs have usually been successful in obtaining and using the hardware and software necessary for their work. This decentralized mode of procurement and installation is, however, often inefficient because:

1. The purchaser has to sort through a lot of information in making the initial acquisition of equipment and software. This information includes the relative merits

of different vendors' equipment and pricing practices, software upgrades, and hardware and software compatibility with existing resources.

2. The purchaser, often not a computer expert, is ill-equipped to make these decisions. One may make a poor choice of equipment and consequently spend three times its price (in labor) in making it work.

3. The individuals usually responsible for installing and operating the data acquisition system (graduate students and postdocs) are usually not familiar with computers at the level required for real-time experiment-interface work. Within reason, setting up a computer-based data acquisition and analysis system is a useful educational exercise, but it is difficult for a student to get a Ph.D. in, say, Chemistry, when most research time is spent doing computer interfacing and programming. This kind of graduate experience does happen occasionally in groups that use computers extensively.

What support should be provided?

A first recommendation would be to provide a group of up to 6 employees. This group should teach courses in CBDAA both to users and to undergraduates, diagnose and fix difficult CBDAA system bugs, write peripheral device drivers for several common operating systems, install off-the shelf data-acquisition peripherals, and provide sound advice on appropriate CBDAA configurations for specific apparatus and on purchases of CBDAA hardware and software. If the group is unable to solve a problem, personnel must be able to procure outside consulting services that can solve it.

In addition, this support group would be responsible for knowing Rice's resources in individual research groups and departments and for making individual groups aware of computer resources on campus that might be useful to them. These resources include personnel with specific expertise as well as computer hardware and software. Further, this group would be responsible for at least examining the issue of development of software that is of general utility in CBDAA. Some simple things can be quite useful, such as the ability to import CBDAA-generated files into commercial programs for graphic manipulation, numerical manipulation, spreadsheets, and word processing. These are relatively easy to do but yet so difficult that few researchers make good use of available commercial software products.

This group should identify and support use of a few standard systems for CBDAA. At present, there are over a dozen different kinds of computer systems in use at Rice for CBDAA. We should be able to reduce this number to three or four. If this activity is effective, we will ultimately increase our research productivity somewhat through standardization and resource sharing. If (and only if) the "standard" systems are truly well-supported internally, people will migrate to their use.