

**Using Information Technology to
Achieve the Strategic Goals of
Rutgers, The State University of New Jersey**

Appendix 2: Report of the Research Subcommittee

Executive Summary

Information Technology (IT) is an essential element of research in virtually every discipline. The Research Subcommittee evaluated the status of IT for support of research at the University in the context of seven categories: Database and Database Servers, Experimental IT for Research, Grant Administration, High Performance Computing, Libraries, Network and Videoconferencing and Streaming. The Subcommittee prepared a list of recommendations which cut across all seven categories. These recommendations are listed in Table 1.

The recommendations are organized into five major categories. *Education* covers a broad range of training requirements. Examples include courses on Message Passing Interface (MPI) for parallel computing in a cluster environment and procedures and requirements for submitting proposals to the NSF-funded National Supercomputer Centers for computer resources. In addition, a Research IT Portal (extension of myRutgers portal) is needed to aggregate many research resource streams and build a common environment including tools that can enable faculty to efficiently and effectively identify potential research colleagues at the University. *Grant Administration and Intellectual Property* includes three major recommendations including development of a web-based proposal creation and submission system for seamless interface to the electronic submission systems of federal agencies (e.g., NSF Fastlane). *Network* comprises five recommendations. Upgrading legacy networks to (at least) RUnet 2000 standards is an essential requirement. Additionally, the University must plan and implement upgrades to inter and intradepartmental networks, and university-wide networks to the next generation hardware and software in a planned manner. Many universities are providing network-based access to expensive, one-of-a-kind experimental equipment as part of multi-institutional research collaborations. Rutgers participation in federated resources using advanced authentication processes for access is essential to enable Rutgers faculty and students to use these extra University facilities. *Systems Support* encompasses eight recommendations. University full funding for core research hardware and software systems acquisitions and upgrades is essential to maintain Rutgers position as a major public research university. Herein, "core" implies those hardware and software systems which federal, state and private agencies are unwilling to fund. In addition, University matching (at least 50%) for acquisition and upgrades of department and laboratory-level clusters ("variable" hardware and software systems) is critical to maintain Rutgers research preeminence. Rutgers University has developed a strongly *decentralized* model for research computing, and thus Rutgers central administration and faculty must work cooperatively to provide the needed research computing infrastructure. Additional needs include adequate common computer datacenter, a *small* (16 to 32processor) central cluster for training and testing (already in acquisition), development of strategic university-wide discounts on hardware and software, and OIT authentication of user of library resources by persons located outside the University. *Video* technology has developed at a rapid pace and

ITSP Appendix 2: Research Subcommittee
Public Draft 1: February 15, 2006

videoconferencing is becoming critically important in the current multi-institutional/multidisciplinary research environment at Rutgers and other universities. The recommendation is the establishment of a Video Technology and Support Group within OIT to assist faculty in the application of video technology in their research.

ITSP Appendix 2: Research Subcommittee
Public Draft 1: February 15, 2006

Table 1. Recommendations								
<i>Category</i>	<i>Recommendation</i>	<i>DB</i>	<i>EX</i>	<i>GA</i>	<i>HP</i>	<i>LI</i>	<i>NE</i>	<i>VT</i>
Education	Continuing education in IT (<i>e.g.</i> , Research IT portal, training courses on High Performance Computing, etc)	•	•		•	•	•	•
Grant Admin/ Intellectual Property	Develop web-based proposal creation and management system for seamless electronic submission to federal agencies			•				
	Maintain significant intellectual property created by Rutgers faculty and students in collaboration with OCLTT	•		•		•		
	Acquire and/or license needed information resources for access through Rutgers Libraries					•		
Network	Upgrade legacy networks to RUNet 2000 standards	•		•	•	•	•	•
	Upgrade intra and interdepartment networks to 1 Gbps as required	•			•	•	•	•
	Add access to extra-university network, National Lambda Rail, at 10 Gbps.	•			•	•	•	•
	Develop policy for funding upgrades to laboratory, intra and interdepartmental and university-wide network hardware and software	•	•	•	•	•	•	•
	Participate in federated resources using advanced authentication mechanisms to access experimental measurement systems and data on the network.	•	•		•	•	•	
Systems Support	University funding (100%) for core hardware (<i>e.g.</i> , 100 TByte data storage facility) and software requirements of research systems	•			•			
	University matching funding (at least 50%) for variable hardware and software requirements of research systems including department and laboratory clusters	•			•			
	Datacenter for servers including adequate space, power, air conditioning, UPS, security and access	•			•			
	Establish small (16 to 32processor) centrally available 32bit and 64bit clusters for testing and training				•			
	Professionally staff IT workshop providing staff and resources to design special purpose systems		•					
	Establish common systems management for secure access, archival and disaster recovery	•		•	•	•	•	
	OIT authentication of the use of library resources and other networked information					•		
	Develop strategic discount arrangements with hardware and software vendors on a university-wide basis	•	•	•	•	•	•	•
Video	Establish Video Technology and Support Group within Office of Information Technology				•	•	•	•

Legend

DB	Databases and Database Servers	EX	Experimental IT for Research
GA	Grant Administration	HP	High Performance Computing
LI	Libraries	NE	Network
VT	Videoconferencing/Streaming Technologies		

1. Planning Process

The Research Subcommittee held nine meetings (see Table 2). At the first meeting, the Subcommittee defined eight categories of IT in research at the University. A draft report on the status and requirements for IT in each category was prepared by a member of the Subcommittee and distributed to the entire Subcommittee. Each report was discussed at a meeting of the Subcommittee and a final version was prepared¹. The final meeting of the Subcommittee focused on defining the priorities for implementation. The Subcommittee members concluded that the recommendations of each report represented essential needs, and therefore ranked all recommendations as equally important². The draft final report was prepared by the Subcommittee chairperson, distributed to all members of the Subcommittee for comments. The final report was prepared incorporating the comments of the Subcommittee members.

<i>Date</i>	<i>Report Discussed</i>	<i>Prepared by</i>
29 November 2004	Report Outline	Subcommittee
27 January 2005	Libraries	Lynn Mullins
24 February 2005	Grant Administration	Keith Osterhage
24 March 2005	Experimental IT for Research	Dinesh Pai
21 April 2005	High Performance Computing	Georgiy Stenchikov
19 May 2005	Database Servers and Databases	Helen Berman
22 September 2005	Network	Doyle Knight
13 October 2005	Centralized IT Support	Diane Davis
	Videoconferencing and Streaming	Tom Grzelak
20 October	Recommendations for Final Report	Subcommittee

2. Description of Current and Future State

The current and future state of the seven areas of IT in research are described in the following reports presented in alphabetical order.

¹ The reports on Database Servers and Databases and Centralized IT Support were consolidated into a single report entitled Databases and Database Servers.

² The exception is the recommendations for provision of a small centrally available 32bit and 64bit cluster, which is already in acquisition, and upgrade of the external network to 10 Gbps (National Lambda Rail), which was seen to be too far in the future to be recommended in this near term strategic planning process.

3a. Databases and Database Servers

To support diverse IT computing needs of the academic and research units within the University a range of services are required. These services may vary from the delivery of turnkey database and analysis servers to providing consulting on leading edge technologies. It is essential that the University adopt a flexible, creative, and collaborative support model capable of supporting these diverse requirements.

Common to all IT projects are requirements for physical, network, collaboration, and procurement infrastructures. In addition to these common requirements are the diverse IT project requirements for server hardware, software, and system management.

Core IT Requirements

The following requirements are largely common to all IT projects.

Physical Infrastructure

A basic requirement of any IT project is the physical space to store server, peripheral, and networking hardware. This space must provide reliable environmental and power conditioning with backup systems to ensure continuous 7x24 service. Secondary space should be available to support servers for applications with critical uptime or availability requirements.

Network Infrastructure

Network reliability, performance and security are critical to all IT projects. The former issues require the University to maintain multiple points of presence on high performance and commodity networks. Security issues common to many IT projects require the partitioning public and private network activities and involve the construction of firewall environments and virtual private networks. Similarly, network load balancing is typically required to support web access to a distributed IT computing architecture. Solutions to all of these problems should be readily accessible to any IT project within the University.

Collaboration Infrastructure

Virtually all IT project involve collaboration among geographically distant groups. The University should maintain expertise and provide facile solutions for constructing network collaborative environments. There is also a need for authentication protocols for collaborative users.

Procurement

The University should aggressively pursue strategic discount arrangements with hardware and software vendors to minimize the cost of acquisition and maintenance of key/common hardware and software products for IT projects.

Variable IT Requirements

The following issues are key to all IT projects; however, these issues may be addressed in different ways depending on the local infrastructure of a particular IT project.

Server Hardware

IT projects may require the identification, acquisition, installation and configuration of server hardware. The specific requirements for server hardware may result from compatibility considerations with existing equipment or with budgetary considerations. In either case the University should be prepared to assist in the selection, acquisition, installation and ongoing maintenance of IT hardware.

Software

Although most commercial software vendors provide some special consideration for academic sites this support is typically limited to instructional use. Applications deployed on the Internet for public use are often excluded from academic coverage and require expensive commercial licenses. Coordinated University licensing for common operating system and application software can achieve significant economies. Some database and analysis software applications require significant support and continuing maintenance. Support should be available for these ongoing services.

System Management and Administration

Every IT project will have specific requirements for system management and administration. Some projects will require a comprehensive system for managing secure access, configuration, monitoring, archiving and routine software updates. Other projects with local infrastructure and staff may require substantially less support. Solutions for common system management tasks such as secure access (e.g. one time password access), archiving and disaster recovery should also be provided by the University.

Funding Models

There are a variety of funding models to support the required infrastructure. To be competitive the University must provide the core requirements at no cost to the researcher. At least 50% of the cost of the variable requirements should be borne by the University and a charging mechanism should be put in place for the remaining 50%. The details of how this would work should be part of an implementation plan. In making this plan there must be consultation with the stakeholders.

Examples

There are several projects at Rutgers that have requirements for the types of infrastructure described here. Two examples are given.

The Protein Data Bank

The PDB is the international repository for the three dimensional structures of biological macromolecules. Currently there are more than 33,000 structures in the database

comprising more than 45 GBytes of data. More than 100 new structures are deposited per week. More than 10,000 individual users per day access the site. The PDB is funded by 8 government agencies with an annual budget of \$6,000,000. The PDB is an essential resource for academic and commercial scientists worldwide who have come to depend on it as a reliable source of data for their research. Loss of the PDB or any compromise on service would have a profound impact on the conduct of research worldwide, as this is the single archive for these data.

In the first round of this grant no funds were provided for equipment and it was the expectation of the funding agencies that the University would supply this. Rutgers provided some funds for this. The rest came from gifts. In the second round we did secure some funds for equipment. However, in the current funding environment, it is unlikely that this will continue.

The Large Data Set (LDS) Analysis Platform

The LDS platform is a centrally located resource used by several dozen faculty from more than 18 departments, Centers, Bureaus and Institutes through out the University. The Large Dataset Users of the Social Science Research Community represent a rich source of expertise, talent and energy in conducting major research at the University. The increasing University/state collaboration among faculty and public sector administrators and researchers about policy issues requires the use of cost, services and administrative data with millions of case records that must be merged for analyses. The critical issues for this important research are extreme I/O needs and extreme security needs due to the sensitive nature of the much of the information. In addition it is key that we maintain the ability to read data from the wide range of (sometimes antiquated) media provided by the source agencies. The central computing infrastructure must be capable of supporting the requirements of researchers merging large sensitive datasets in efforts to examine integrated patterns of behavior, policies and costs. Their ability to acquire state and federal grants and contracts is noteworthy, but their ability to continue to attract major funding is dependent on assurances that the computing environment is capable of handling the processing, securing and analyses of large sensitive databases.

The LDS platform was created in 1997 by cobbling Institutional Strategic Resource and Opportunity Analysis (SROA) and CBI (Center Bureau and Institute) dollars together after the decommission of the IBM 3081 research mainframe in 1995 resulted in an inability to meet grant and contractual commitments. Another one time allocation initiative (Reinvest in Rutgers) was used to replace aged disk storage in 2001. It is critical that the University embrace its role in ensuring maintenance and upgrading of this resource as part of the infrastructure of a research University.

3b. Experimental IT for Research

The dramatic rate of change in Information Technology (IT) poses two significant challenges for the utilization of IT in support of research.

(1) Research support for IT must plan for changing priorities. It must be sufficiently flexible to accommodate evolving trends in IT and IT research.

For example, newer areas of IT such as embedded computing, visualization, human interfaces, peer-to-peer networks, mobile computing, and sensor networks could have as much of an impact on research productivity as the more mature IT areas such as high performance computing. For a strategic plan in IT to be effective, the university needs to anticipate future needs, in addition to addressing current ones.

(2) The university must help researchers quickly adopt new information technologies. Researchers are often unaware of the possibilities created by new developments in IT. Even when they are aware of the technologies, lack of background expertise creates barriers to utilizing them. This is unfortunate, since rapid adoption of new technologies could provide a significant competitive edge in research.

For example, wireless sensor networks may allow social scientists and ecologists observe large groups of animals in a natural setting (*e.g.*, ZebraNet³). Human interfaces and visualization techniques may help neuroscientists and molecular biologists to better understand the enormous quantities of data they generate from sensors and simulations.

The university can take several steps to address these emerging challenges in IT. There are no clear role models at peer institutions that we can emulate; Rutgers will have to lead, by creating new organizational entities to support novel uses of IT in research. Some possibilities are listed below.

Continuing Education

The university should create effective mechanisms for disseminating to the community developments in IT relevant to research. Such mechanisms should balance the need to inform, with the need to respect the very busy schedules of Rutgers researchers.

Education efforts could include:

- Brief email newsletters that alert readers to novel uses of IT in research, and provide contact information about people at Rutgers who could assist (see below).
- A "Research IT" web portal that maintains links to IT relevant to research. The portal could include discussion groups and wikis. It could include a "classified ads" section that could be used to solicit IT-related work (*e.g.*, for student programmers).
- Short training courses in research IT. For example, many experimental sciences use complex software, for instance Matlab, to process data. A half-day course, advertised on the Research IT portal, could benefit graduate students in several

³ <http://www.princeton.edu/~mrm/zebranet.html>

disciplines, especially if it introduced new techniques (for example, how to use computer vision techniques in Matlab for counting microorganisms using a video microscope).

Research IT Workshops

In a research setting, the IT required is often not available for purchase from vendors, or may require significant effort to integrate. Individual faculty and their students usually lack the expertise required to develop such experimental technologies. A professionally staffed IT workshop could serve as a bridge between researchers and the experimental IT they need for their research.

Perhaps an analogy will help: many engineering and science departments have staffed workshops where researchers (esp. graduate students) can get professional help in prototyping research ideas (e.g., to machine a novel lens mount, to fabricate electronics, etc.). Their job is not to maintain the heating in the buildings. IT needs such workshops too, which can help a researcher build prototype software and hardware systems.

For example, psychologists may benefit from a facility that provides access to high level technicians with knowledge of the state-of-the-art in integrating electronics and computers, who can assemble a special purpose data acquisition system. In some IT-related disciplines, the departmental IT staff may be able to meet this need to an extent, but their primary responsibility is usually to support teaching and shared infrastructure.

Since these IT workshops are new types of entities, some experimentation will be needed with their organizational structure. The workshops will have to be carefully staffed, with individuals who enjoy project work and who view the challenge of keeping up with a constantly changing IT landscape as a positive aspect of their jobs. The workshops will probably need to be supported by a combination of internal and external funding. One way to achieve this is to offer contract-based staff services which could be paid from grants.

One possibility is to create centralized units that specialize in specific aspects of experimental IT. For example, there could be a "Embedded Systems Workshop" specializing in building real-time systems which integrate sensors, electronics, and microprocessors. Another could be a "Human Interface Workshop" that could help researchers assemble novel display hardware and interaction software that is tailored to a particular research setting.

An alternative is to have workshops that specialize by discipline; for example, one could have a "Biological Sciences IT Workshop" which would be staffed with professionals familiar with typical requirements of researchers in the discipline.

3c. Grant Administration

Sponsored research grants and contracts administration is properly broken into two distinct but related pieces as described briefly below:

- 1) Pre-award functions that address primarily non-financial issues, such as the grant proposal creation and submission process (including budget preparation), award notice and acceptance, compliance issues having to do with use of human subjects, animals, homeland security and a growing list of other federal requirements, negotiation of contract terms and conditions, interfacing with federal and other granting and contracting agencies, and the creation and maintenance of a comprehensive grants and contracts database,
- 2) Post-award functions that address financial administration of award expenditures, compliance with financial and related contract or grant terms and conditions, creation and maintenance of an expenditure database that is either compatible with or part of the institution's financial accounting system and is cross linked to the pre-awards database, financial audit issues, often in partnership with the pre-award function, and any issues of a financial nature that may affect expenditures under sponsored research grants and contracts.

The computing and software support needs for these two functions are generally quite different, although there are important areas of overlap that offer the opportunity of shared use. In particular, the pre-award software needs are primarily non-financial and typically are not well served by a system that is designed around the core of a financial accounting system, such as the major enterprise systems now being sold to universities that were first designed and created in a commercial setting to meet commercial needs. The post-award expenditure accounting needs may be accommodated relatively more easily in such systems, but the grants management issues that are at the core of the pre-award administrative function are typically not well accommodated. For this reason, the Oracle system recently has adapted a version of the Coeus Grants Management system that was created in a primarily pre-award user environment at MIT.

As a historical note, Rutgers was scheduled to be the first user of the Oracle/Coeus Grants Management module in 2001 until the effort was abandoned almost three years ago in 2002 when Rutgers scaled back its implementation of the Oracle suite. Following this scaleback, the Associate Vice President for Research reevaluated the situation and began to implement the Coeus Grants Management system, which had previously been licensed from MIT (\$500 nominal licensing fee). This decision was reinforced in May of 2003 by the Executive Vice President for Administrative Affairs, who set a go live date of July 1, 2003. That deadline was met and the Office of Research and Sponsored Programs has been using Coeus as its Grants Management tool and database of record since that time. ORSP now passes data to the post award unit (DGCA) through bridge software written for the purpose at Rutgers. A number of Coeus features are designed to help

support the post-award process and Coeus is available to the post-award staff. In addition, a new web-based version of Coeus is due out soon that will support electronic proposal creation and submission to federal and other granting agencies. In short, Rutgers currently has a successful experience with the Coeus system that merits further development for both pre and post award use.

Response to Research Subcommittee Questions:

1. What are Current and Near Term IT Infrastructure Needs for Grants Administration?

The most urgent need for ORSP is to implement a web-based electronic proposal creation and management system. This will require a server, dedicated and maintained by OIT, and appropriate support and training. Training will be required for ORSP and DGCA staff, as well as 3040 unit business managers. Implementation of the project will require a pilot phase that will gradually include all units with significant grants and contract activity. Such a system must mesh easily with the financial software system of the university. Other needs include roll out of the web based IRB protocol management system, and future development of a similar web based animal use protocol system.

2. Current/Present Status of IT Infrastructure for Grants Administration

Currently ORSP has an IT manager and small staff who are responsible for acquiring and maintaining the servers for ORSP, recommending, installing and maintaining software for support of the pre-award administrative function, providing email and other RUNet access, creating and maintaining websites and performing other computer and software related support. This small IT group has successfully transitioned the pre-award unit from dependence on an in-house created database system (ORSPLOG) to implementation of the Coeus Grants Management System in July of 2003. At this time Coeus is used primarily as a login and database system, supporting mostly pre-award functions, but also supporting an increasing amount of post-award function. This small IT shop also supports Laboratory Animal Services, the Office of Corporate Liaison and Technology Transfer and the ever expanding federally regulated human subjects and animal use compliance function within ORSP. Currently ORSP/Compliance has worked with a commercial software vendor, using an NIH grant as partial support, to develop a web based IRB management tool, which should be coming on line within the next month or so. It hopes to transfer some of this development effort into a similar project for managing its animal protocols as well.

3. How do we propose to fill the gap between a) and b) above (questions 1 and 2)? What are the priorities for funding?

The major need for an electronic proposal creation and management system will be met by moving to the soon to be released COEUS web based environment. This system has been designed specifically to address the proposal creation and management issues in a user oriented environment at MIT. It also has been created in close cooperation with the NIH and other major federal funding agencies so that it will be fully compatible with the newly emerging federal electronic grants submission and management environment.

The human subjects and animal use protocol management tools are being created in cooperation with a specialty software vendor and with support from an NIH grant. This program is moving quickly toward full roll out and use of the human subjects module, and work has begun on the corresponding animal protocol module.

3d. High Performance Computing

High Performance Computing (HPC) is a general term used to describe supercomputing across various technologies such as cluster computers, massively parallel processing computers, and grid computing. Rutgers University has a varied history of employing these technologies. According to a survey of Rutgers HPC clusters conducted by the Office of Instructional and Research Technology (OIRT), there are currently more than 35 clusters in operation by 17 groups. The efficient use and promotion of HPC at Rutgers University is very important for continuing progress and competitiveness in research and academics.

HPC resources have historically been available for many years at national supercomputing and research centers. Rutgers University operates central HPC facilities within the Center for Advanced Information Processing (CAIP) consisting of Sun Microsystems' Starfire E10000 (2) and Sunfire E12000 (1) units. These units are obsolete and costly to maintain. The largest on-campus use of HPC is through the proliferation of computing clusters using commodity PCs. These efficient yet affordable supercomputing facilities allow them to be purchased and maintained by individual research groups. A sample cost for 12CPU cluster is approximately 15K25K. The relatively low cost use of commodity PCs clusters contributes to their widespread use as a sustainable HPC resource.

The HPC environment (hardware, software, etc.) rapidly changes with a turnover cycle of 34 years. Advanced centralized university planning including development support and consulting is necessary for efficient use and promotion of available resources. The Office of Instructional and Research Technology should play a leading role in this process. HPC activities should be supported either using onsite (clustering) or external facilities at the national supercomputer centers. In both cases, a key question is how to grow and maintain central university expertise in HPC.

There are several issues within HPC that need to be addressed:

- Expertise in parallel programming and efficient designs of HPC clusters systems require significant expertise and knowledge. There should be resources devoted to developing a systematic approach of disseminating this information, consulting and training faculty and students.
- Most commodity PC clusters are built upon 32bit hardware and operating systems which limit processes to 2GB of addressable memory. There are university researchers whose problems require migration to 64bit hardware

which removes the 2GB restriction. Migration from 32bit to 64bit involves issues that require further expertise.

- National supercomputing facilities are available for providing HPC resources to researchers. Their use needs to be facilitated and promoted. In this case it is important to provide assistance to researchers using offsite HPC in testing their software and adjusting to the particular programming environment and hardware at the remote sites. Efficient offsite computing is impossible without well organized system of transport and archiving information on site.

OIRT's survey of HPC activities at Rutgers University notes that 4050 research activities at the university are already using or very soon will be conducting supercomputer computations in a very diverse programming environment. The available systems are mostly in use by the individual departments and research groups. There little evidence of a sharing capacity because the resources are mostly used by the owners of the particular computing systems.

To further develop and improve HPC at Rutgers University, the following measures seem to be desirable at present stage:

Training

Provide training (for faculty and students) on parallel programming including leading programming languages and parallel libraries like (MPI, Open MP). It is important also to have access to the universal debuggers like TotalView and teach how to use this software. Onsite consulting services on all related problems will be very useful.

Onsite HPC

University must provide centralized consulting services on building cluster systems for purposes of individual departments and scientific groups.

A program of 50% University matching funding for new department-based cluster systems is essential to support growth of high performance computing in education and research in the University. Assuming ten new clusters per year, the estimated cost of this program is \$240K per year.⁴

There should be *small* (e.g., 16 to 32 processor) centrally-available 32bit and 64bit clusters for common use to test available software and optimizing parameters of the systems that are planned to be purchased for particular applications. These centralized facilities would be used strictly for testing of hardware and software and not as computational platform.

⁴ The estimate assumes ten new clusters per year. Each cluster is 32 processors at \$1.5K per processor (including network interconnect, rack, etc) for a total of \$48K per cluster. At 50% cost sharing from the University, the total annual cost is $10 \times \$24K = \$240K$ per year.

Offsite HPC

Provide information about available offsite HPC resources and consulting on using those resources. Training on working in different programming environments would be useful.

Centralized Support of HPC

- Develop a massive storage facility (UniTree type) with capacity of about 100 Tb.
- Provide disk space for temporary use for a reasonable fee.
- Provide backup services for the departments and groups for some reasonable fee.
- Consider centralized support of GRID computing.

Centralized technical support

Plan on installing computer facility/data center designed for Large High Performance Computing Clusters that will be equipped with sufficient electric power supply (12 MW) lines and air conditioning. This facility will be administered by OIT and will be available for the University community. This will help to solve already existing problems with maintaining large computing clusters at Chemistry and Physics Departments that do not have sufficient electric power supply in their buildings.

3e. Libraries

The Rutgers University Libraries seek to support and advance the research and scholarly needs of the University and its strategic partners in diverse ways. In terms of information technologies, these include:

- the provision of access to scholarly information in all formats
- the preservation and accessibility of the intellectual content of the Rutgers community in digital format
- the provision of a technical infrastructure and tools that support both user access
- and internal library processes

Access to Scholarly Information

The provision of access to information continues and builds on the Libraries' historic mission of building and preserving collections and expediting the delivery of information. Given the rapid changes in the information environment and the technologies, the fundamental and ongoing need is to continue to acquire/license/and otherwise make available needed information resources, especially in digital format, to researchers throughout the University in support of their work. Given the anticipated offering of more digital video and audio, there is a need for an expanded bandwidth of the network.

There is also a need for OIT authentication and authorization services so that the user community, including research collaborators at other institutions, such as UMDNJ and NJIT, and incubator companies, is fully identified and authorized for access to resources that is in concert with licenses and/or other agreements. The Rutgers Libraries are well aware of the challenges presented by the need for the preservation of digital information on a national and international basis, and, as a research library, will continue to monitor advances in this area.

Rutgers Digital Archive

The maintenance of the significant intellectual property of the Rutgers University community, particularly the unique information created by Rutgers faculty and graduate students and by University publishers, is envisioned as being addressed by the Libraries through the development of a Rutgers digital archive or institutional repository. A somewhat newer role, the Rutgers digital archive is part of an effort of research libraries to capture and preserve the intellectual output of their universities and make this available as part of a federation of systems. The development of a “trusted repository” means that four key tests need to be met. These encompass the following: the safeguarding of intellectual property; fidelity to the source of information; authenticity of information; and the preservation of user privacy. Rutgers depositors and users will be able to trust the new digital archive to provide authentic and useful information for many generations of researchers and students.

The creation of the Rutgers digital archive is part of the changing system of scholarship and scholarly communication. Activities of the archive/repository will include the creation, storage and preservation, description, access, and display of content, and digital rights management. At RUL, work is being done on infrastructure development. It is envisioned that primary clients for the initial phase of the repository may include faculty, with early content to possibly include research materials that would otherwise not be exposed to researchers and students, such as working papers. Services and tools offered in support of the University’s research process will include: consultation on data design and management for research projects; “best practices” recommendations for archiving research products, such as data sets; tools for creating, publishing, archiving, and indexing research products, such as databases, data sets, websites, reports, and multimedia files; and a sustainable platform for maintaining the results of research projects (frequently mandated by granting agencies).

The initial phase of this digital archive build out will include working with Rutgers University publishers, such as graduate schools, through the development of an electronic thesis and dissertation system. It may also possibly focus on working with faculty developing federal grants that require ongoing efforts to maintain and preserve the underlying data for their research. This would allow the investigators to focus on their research and innovation, while the Libraries could contribute expertise to the management of the intellectual products produced. Apart from the continued work on

the infrastructure, work needs to be done in many areas prioritizing content development, planning for long-term digital preservation, collaborative program development, the creation of a business plan (for financing, sustainability, contingency planning, marketing, etc.), and the like.

Repository interfaces will need to be developed so that the archive can be easily integrated into the workflow of the primary users. There will be needs for mass storage space, for the OIT authentication and authorization services noted above, and for resources for further development and implementation. Collaboration with OIT to support the Sakai learning content management system is also envisioned, and work will need to be done to develop and test the interface between the Fedora repository platform and Sakai so that content from the archive can be mined for use as course documents. Increased dedicated bandwidth, such as NLR or Internet2, will probably be needed over time, as faculty in different departments use the repository web interface to send large amounts of data to the repository simultaneously.

As envisioned above, the digital archive content will include the unique research products resulting from academic research projects, particularly the data sets, lab notes, images, prepublication articles and grant reports that would otherwise not be easily discovered and utilized. In addition to theses and dissertations, additional collection areas will include the Rutgers University Press; the building of digital collections based on collection strength and use; and collections in the Libraries' Special Collections and University Archives and the Institute of Jazz Studies.

The digital archive development work is part of the Libraries' strategic planning initiative that is currently taking place.

Technical Infrastructure

The provision of a technical infrastructure refers to the ongoing need to insure that the Libraries' information systems are state-of-the-art and that there is a systematic replacement and upgrading of equipment and infrastructure throughout the Libraries to support the stewardship of the collections and the management of services. The many tools include the array of web-based user services that are directed at assisting faculty and others with identifying, obtaining, transforming, and using information.

3f. Network

Present Status of Network

The present Rutgers University Network infrastructure serving research computing is primarily based on RUnet 2000, RUWireless and RUWired projects. During the period between the years 2000 and about 2002, the RUnet 2000 network infrastructure was implemented on Rutgers campuses. It provided 100 Mbit/sec connection to desktop computers and servers within departmental buildings, libraries and student apartments through Virtual Local Area Network (VLAN) switches and 1 GBit/sec interconnect

between the switches and the main routers. However, some buildings have not been brought up to RUnet 2000 standards and still rely on legacy network systems (*e.g.*, a single T1 [1.5 Mbit/sec] link into an entire building).

Beginning in 2002, Rutgers initiated RUWireless and RUWired projects to enable students, staff, and faculty with a valid university NetID to connect their laptops and desktops to the Internet through 802.11b (WiFi) wireless network and Ethernet access points. At the present time, 11 Mbit/sec RUWireless and 100 Mbit/sec RUWired cover a limited number of areas in the New Brunswick/Piscataway campuses.

Recommendations

The following recommendations for improving intradepartmental, interdepartmental, and extra-university network connections are presented:

- Intradepartmental networks
The trend in High Performance Computing at the university level has shifted from multi-user environments on multiprocessor machine (*e.g.*, SUN E10000) to laboratory and departmental-based computational clusters and desktops. The computational resources on these distributed systems need to be consolidated through high performance networks at a minimum bandwidth of 1 GBit/sec for file transfer between computational clusters and desktops, shared data storage and archival devices.
- Interdepartmental networks
The recent pronounced shift in Federal funding from single investigator, single discipline research to multi-investigator, multidisciplinary research has led to increased collaboration between university research groups from the different departments (*e.g.*, the School of Engineering and Department of Chemistry). This, in turn, has led to the need for faster network connection between their computers (desktops and servers) for efficient sharing of data and software. A minimum of 1 Gbit/sec bandwidth is needed between departmental networks to achieve the same level of network performance as for interdepartmental networks.
- Extra-university networks
Research Collaboration between Rutgers University departments and the outside institutions (for example UMDNJ) as well as access to the National Supercomputing Centers (such as NCSA, Pittsburgh and San Diego) requires high network bandwidth connection with the Internet for data transfer and database access. The merger of Internet2 and the National Lambda Rail represents the next generation of transcontinental network system providing 10 GBit/sec to universities on dark fibre. Implementation of such high bandwidth network connections requires provision of concomitant intra-university network bandwidth and, in selected areas, restructuring of intra-university network topology.

Videoconferencing is becoming an increasingly important tool for research collaboration. However, videoconferencing software standards are rapidly evolving and individual faculty must invest significant time and effort to achieve and maintain a videoconferencing capability with colleagues elsewhere. A central support service for videoconferencing (*e.g.*, within OIRT) is definitely required.

- Upgrading Buildings with Legacy Network to RUNet 2000
A number of buildings on all three campuses still suffer from low bandwidth network connections due to legacy network connections. These buildings need to be upgraded to RUNet 2000 standards.
- Planning for Converged Communication Technologies
Rutgers OIT needs to develop specific plans for implementation of converged communication technologies on data networks, *i.e.*, VoIP (Voice over Internet Protocol), video (including digital television) and data.
- Experimental Facilities on Extra-university Networks
Experimental measurement systems (*e.g.*, electron microscopes) are typically configured with TCP/IP connections rather than the classical hardwire RS232 connection. This approach enables remote access and sharing of expensive experimental devices. A concomitant requirement is user authentication. Rutgers must participate in the Federated Authentication Process currently under development.
- Evergreening of Network Equipment
A clear policy and procedure for upgrading laboratory, departmental, interdepartmental and university-wide network equipment and software is needed.
- Special Needs
Specific research projects may need higher bandwidth network connections to a local campus site than is afforded by the campus network, even considering the upgrades recommended above.

3g. Videoconferencing and streaming

The use of video in research and instruction is expanding rapidly in both capability and future promise. It is an important technology consistent with the emerging premise of the convergence of voice, video and data onto a single media/data channel. Rutgers University needs to attain and maintain support for this important method of communication. It is envisioned that desktop videoconferencing will become as ubiquitous as the telephone. The use of video is particularly critical to Rutgers simply by the distributed nature of its three main campuses and its numerous off-campus facilities. The effective employment of video technologies provides a basis for the way the university conducts its research, instructional and administrative business.

Current Video Applications and Directions

Peer universities have been active in promoting the use of video technologies in their research and instruction. Video has been used as an effective collaborative resource in a number of ways. Videoconferencing now scales from conference rooms to the desktop. It can incorporate data sharing, interactive document creation and high-definition video. It can integrate multiple feeds at remote sites into multi-windowed videoconferences (e.g., Access Grids <http://www.accessgrid.org>). Videoconferencing provides a proven mechanism for research and instructional collaboration, remote site training and K12 outreach.

The role of video in instruction is multifaceted. Distance learning courses have traditionally used live (aka streaming) and recorded video as part of their curriculum. There is significant interest in the use of video to archive class content including guest speakers, provide ability for off-campus students to attend a class (e.g. Route 18 project issues), etc. Libraries are actively creating video-on-demand archives for streaming into the classroom (e.g., Northwestern University). It is further evident that traditional analog video (e.g., television) will be supplanted by digital video. Streaming of traditional television video across the data channel has been noted as an emerging technology and is operational at other universities (<http://datn.wisc.edu/>).

Current University Support for Video

Support for video technologies at Rutgers is not broadly available. The Division of Continuous Education (DCEO) provides video support including schedulable videoconferencing rooms, advice on departmental purchases, technical documentation and video streaming services on a limited basis. Other streaming servers are available on various campuses, but not widespread or organized. Purchase programs have promoted the deployment of videoconferencing equipment with little backend support to ensure successful integration and use. There is no organized, staffed effort to obtain, test and implement new and existing video technologies. With many technologies, initial triumph or failure greatly influences future use. Successful employment of video requires the development of technical standards and practices supported by staff that can assist units in deploying and using this equipment.

Video technologies are cross-discipline and involve nearly all aspects of university life in one way or another. Its importance cannot be understated. The Office of Instructional and Research Technology (OIRT) has been reviewing video technologies available within and external to the university for nearly a year. Their consultations with faculty, deans and research units universally discuss video as a key technology for their respective unit. It is clear: support for video is broadly needed.

Development of Support Structure for Video

Video technologies continue to evolve and expand at a rapid pace. University adoption of these technologies has varied greatly unit to unit. Video support should be developed

within a central group whose responsibility is university-wide and whose primary mission is consistent with providing technical assistance (such as OIRT.) The size of the video support group needs to be large enough to address the current and future issues of the university. The group must be a visible and available resource, providing training and expertise in video technologies to all faculty and staff.

A three-member group is proposed in order to support the broad video needs and development of the university. The first member is a video coordinator whose responsibility is to lead the group, assess the video needs of the university, identify external resources and serve as the university-wide coordination point for video. This coordinator would guide the development of technical standards, conduct user and local IT training and perform the required planning to ensure a level playing field with respect to video technologies at the university. The remaining two members would be video technicians. These technicians would be responsible for evaluating available video hardware/ software (e.g., videoconferencing units, ISDN bridges, streaming servers, etc.), developing documentation and training, and conducting onsite assistance and assessments.

A preliminary budget required to develop this resource is on the order of \$320K/year, broken down as follows:

Annual salaries \$170K Video Coordinator: IT grade 6/7 \$70K, Video Technicians:
IT grade 4 – 2@\$50K

Equipment & Hourly Support \$150K includes part-time support staff on all 3
campuses, includes development of shared, high-end facilities.

Please note this budget is an entry point for institutional support for video technologies. Future growth and development will require additional resources.

Conclusion

From collaboration to demonstration, communication is a critical element of successful endeavors. Video technologies are a powerful, multifaceted mechanism for supporting effective communication that is cross-discipline and increasingly intrinsic to instruction and research. The university must recognize the significance of video and its importance to many of the university's goals. Its support must be proactive, purposeful and broadly applied in order for it to become a true university platform for communication.