

Research Computing at Notre Dame

A Proposal to Establish a Center for
Research Computing Support

June 10, 2005



The success of the University of Notre Dame's strategy to increase the size of the research program depends heavily upon the ability to provide effective infrastructure and support, especially in disciplines that rely upon information technology resources to accomplish that research. This set of disciplines is growing, with many programs in the Arts and Humanities joining those of Science and Engineering, the traditional source of requests for IT support. The University's support for IT used in research has been limited to relatively modest investments by the Office of Information Technologies (OIT), and limited amounts of support provided by individual Colleges and Departments, often funded from research grants.

The growth of the research agenda of the University has come with a renewed call for increased IT infrastructure support. A number of previous reports and recommendations have been generated by the Colleges detailing the shortcomings of extant support offerings, and suggesting ways to improve the level of support. Those reports have been joined recently by recommendations from an Ad Hoc Committee of Research Computing Infrastructure co-sponsored by the Vice President for Research and the Chief Information Officer, and a concurrent Academic Support Assessment conducted by the Office of Information Technologies.

Virtually all of the previous recommendations converged on the establishment of a 'center' for research computing. While a center concept can take many forms, this report attempts to draw upon those previous reports to generate a 'consensus' proposal outlining the recommendations contained in those earlier proposals, and borrows heavily from work that has been done by faculty such as Peter Kogge from Engineering, Olaf Wiest from Science, and others. While the names vary, the overall recommendation is the same – to establish a new multidisciplinary center for research computing at Notre Dame.

This center (CRC, the Center for Research Computing) will have five major goals. First, and primarily, it will support and enhance the research agenda of the University. Second, it will ensure efficient usage of common resources through professional systems administration, user training and exploitation of economies of scale. Third, it will provide a forum for the exchange of ideas and experience from different areas of the campus research community. Fourth, it will coordinate outreach activities in the area of the application of computational methods to real-life problems. Finally, it will leverage University investment by seeking additional outside funds through facilitating the grant writing of individuals or groups of researchers.

Brief Statement of Need

The recent Ad Hoc Committee drew the following conclusions relating to the support of research computing at Notre Dame:

- 1) Research computing, as currently provided, will not support increased research activity.
- 2) Current strategies for the provision of research computing (originating from the 1980s) have to be revisited.

Moreover, these specific areas of concern were highlighted by the committee's report:

- 1) A support “gap” exists – a lack of coordinated support from staff who have blended expertise in information technology and the various fields of research.
- 2) “Do it yourself” has become the norm. This is a potentially dangerous model for success, oftentimes leading to inefficiencies, security risks, and poor use of University resources.
- 3) Major gaps exist in elements of University technology infrastructure needed by the research community, such as enterprise storage.
- 4) Regional & National assets are poorly leveraged. Notre Dame has not (yet) established a coordinated approach to opening up access to research technology assets from existing and potential collaborative research partnerships.
- 5) As research activity increases, specific growth demands will soon be upon us, including such areas as visualization, network bandwidth, and synchronous research and collaboration.

National Context

Successful research at Notre Dame requires an understanding of the larger context of research worldwide. Notre Dame is part of a larger community of researchers, and we must understand where that community is heading, what changes are coming, and identify the opportunities for Notre Dame in the coming years. Here are some fundamental changes in the (inter)national research scene:

- The modern scientific model is increasingly moving from a "theory to experiment" paradigm to a "theory to computational model to experimental verification" paradigm.
- Research computing, especially at the high end, is no longer simply more and more "number crunching," but an integrated ensemble of high bandwidth communications between labs and research facilities, large on-line storage for increasingly complex data sets, increasingly sophisticated data mining and non-numeric information searching, and highly interactive and real-time visualization.
- A ‘federated model’ is becoming the required norm. In most cases, no longer will a single campus or researcher be able to fully identify a challenge and create the solution. Notre Dame researchers cannot and will not work completely on their own. They must (and will) work together with peers at other locations to collaborate on solutions. World-scale collaborative teams are becoming common place.
- Increased emphasis on an international cyberinfrastructure (CI) for research. This technical infrastructure describes a robust, secure, and dedicated collaborative platform for conducting research worldwide. CI offers new options for *what* is done, *how* it is done, and *who* participates.

- With increased frequency, the research data will not exist locally. It will oftentimes be scattered throughout the world. Even more intriguing, with the exploding size of datasets being captured electronically, the actual raw data will never be seen by human eyes!
- Significant growth of research activities in the Arts. There are some stark differences in this activity. For example, in the field of humanistic studies, the data are not developed or collected by the scholarly community or using scientific instruments, rather they are oftentimes publicly available content, such as films, books, and music. These differences bring a set of new challenges, such as source data ownership and archival stewardship.
- Next generation technology is being created (the possible future of the Internet) that will open up exciting new opportunities, potentially changing how we conduct research.

At the national level, the goals of successful research activity closely follow our very own for successful campus research. They include the following:

- More applications, capabilities, efficiency
- Reuse and multiple-use of designs; capture of commonality
- Spread of best practice
- Achieving interoperability
- Provision of tools and services
- Shared facilities
- Assistance and expertise

In order to appropriately make recommendations for our own campus, we must take into account and study this larger national context.

The Proposed Solution

This proposal describes a consensus ‘Center’ approach derived from previous reports. Outlined in this report are several important aspects of the CRC, including the center’s scope, management & governance structure, funding, and space requirements, with suggested solutions.

In scoping the mission of the center, it is conceivable that this center and its leadership could help to shape the direction of research at Notre Dame and especially assist the Colleges in their understanding of how to shape their specific research agenda. However, the initial areas of focus of the center will be around managing IT assets dedicated to the support of research, augmenting existing support models, promoting innovative approaches to solving research challenges with information technology, developing assistance programs, and participating in efforts to strengthen the linkages between research activity and the teaching and learning environment at Notre Dame.

Fundamentally, the center is proposed to be created around five ‘pillars’, which serve as the foundation for its mission. At the highest levels, the five pillars of the new CRC can be explained as the following:

- *Enhance University research agenda.* Enhance advanced research agendas across campus through the availability and use of state of the art high end computing and communication infrastructure and associated resources, especially skilled technical computing staff
- *Enable collaborations across disciplines.* High-performance computing has a major impact in many areas of science, engineering and the arts. There are a wide variety of important problems to which these computational hardware, infrastructure, expertise and software resources can be usefully applied. Efficient sharing of these resources maximizes the benefit to the whole University while increasing the likelihood that any particular field will have available a mélange of expertise that can attack problems in creative and potentially unique ways. The CRC will be a platform for an exchange of expertise, ideas, and resources from many different disciplines.
- *Provide system administration, user support and training.* In order to make efficient use of the investment in hardware, shorten response times, and help researchers explore new applications of high-performance computing in the respective areas, it is important that the user support through highly qualified system administrators is close to the research groups. Previous experience also shows that there is insufficient access to training for potential users of specialized high-end computing resources; thus researchers that could benefit are unaware of the possibilities. The CRC will provide access to training for high performance computing tools and in software design to efficiently use the available hardware and infrastructure.
- *Supply resources for hardware, software, and data storage.* The combination of resources across many academic units on campus exploits economy of scale and allows quick responses to the rapidly changing environment of high performance computing and emerging research areas. The CRC will be the primary resource for these high-end computing needs on campus.
- *Leverage University investment.* Ad-hoc groups of researchers that need high performance computing resources have been very successful in attracting external funding, leading to the creation of the TOP500 computing cluster BoB in 2001 and its upgrade in 2004. The CRC would have an organizational structure designed to aggressively pursue external funding such as the NSF-MRI or DOE initiatives, and for new research initiatives that result from the collaborations within the center. It is also reasonable to expect that the resources, expertise and outreach programs of the CRC will lead to an increase in the success rate of individual researchers that depend on high-performance computing. These resources will complement the University’s investment.

Philosophically, the CRC will take a paradigm-centric approach. The three paradigms include the current (current practice and heavily commoditized); emerging (commercial

support just becoming available); and next generation (currently at research and evaluation stage). In other words, the CRC will:

- a) Help existing researchers optimize existing applications and resources.
- b) Help non-traditional Notre Dame scholars understand and apply existing information technologies to new areas of scholarship.
- c) Prepare the way for infusion of new information technologies into the Academy.
- d) Provide the stage for integrating multiple technologies.

Research Support

Support for campus researchers is currently provided primarily by the University departments. This localized support is critical for successful research activity, and needs to continue. The CRC will spend time selecting areas to augment this support, such as providing training, promoting best practices, creating collaborative opportunities (support teams from across campus can share ideas and solutions), as well as finding ways to expand the enterprise IT resources available to these support teams.

Localized delivery of support does not mean, however, that those localized support assets cannot be centrally managed. Strong consideration will be given to establishing a support model in which research assistance to the field (departments and centers) is managed by a central support element tied to the Center. In this way, the center director can ensure the localized support elements maintain currency in technology applications specific to the research community and maximize the likelihood of cross-discipline transfer of application knowledge.

CRC staff will likely complement University researchers on a variety of disciplines. Not only would this include the traditional computational sciences, but also some of the emerging research in the social sciences. Depending on the expertise of CRC staff, and the ensuing priorities driven through the CRC governance structure, these areas of focus could include visualization, statistical analysis and methodologies, and possibly extend to the use of various data gathering tools and techniques, such as survey instruments.

The CRC leadership will also focus on policy and procedures. Examples include policies that set IT resource access priorities and procedures for protecting and handling research data (thus minimizing University risk of data exposure or loss of critical research data).

Governance

Governance and oversight of the center will be established through lines of reporting for the Director through the CIO and the VP for Research, and through the participation of a Faculty Advisory Committee and a Strategic Advisory Board.

The faculty advisory committee will consist of appointed and elected members of the faculty and will represent the Colleges, the research computing community, and the OIT. This committee will advise the Director in making decisions about operational policy, architectural standards, cost recovery rates, requisite applications expertise, areas of support emphasis, and other items germane to the operation and development of the CRC.

The Director will be a member of the University Council on Academic Technologies (UCAT). The faculty advisory committee will serve as a subcommittee of UCAT for the purposes of advising UCAT on matters relating to research computing on campus.

The Director of the CRC should be a researcher of distinction in a field such as computational science and engineering. The Director should become a member of the faculty of the appropriate department(s) within the respective college of the Director's field of expertise.

The CRC will also establish a Strategic Advisory Board, to aid the Director and Executive Committee in the formation of long-range strategies and policies, including off-campus partnerships, collaborations and funding-raising opportunities. The Strategic Advisory Board will be comprised of Deans, leaders from industry, and noted outside experts in research computing.

Grant Support

The Center staff would lead the efforts to seek federal, corporate, and foundation grants directly tied to increasing the University infrastructure and services supporting the IT needs of the campus research community. This work would be in collaboration with the OIT, the Office of Graduate Studies, various colleges, and specific researchers on campus.

For any of these requests that are successful, the Center would execute on the requirements of the specific grant awarded – acting as the lead on behalf of the University.

For example, the recent grant awarded to Notre Dame (along with two Purdue University campuses) referred to as the “**Northwest Indiana Computational Grid**” (NICG), would be managed by the Center's Director, under the advice of both the CIO and Vice President of Graduate Studies. In support of this grant, the CRC would lead the project to update the enterprise data storage infrastructure dedicated to research activity, create and manage an access grid node and develop strategies to support any additional requirements or opportunities that arise from the NICG grant.

Campus Research Center Administration

There may be additional existing University research centers and institutes that will fall under the operational management structure of the Center for Research Computing. The identification of existing center candidates will come from academic leadership. Once decided, the CRC Director will work with the current administrative oversight body (for

example, a specific College Dean) to transfer management and funding for that specific Center to the CRC.

Funding

Funding needs of the center are generally categorized as capital needs, as stimulated by demand for expansion of existing hardware or the addition of new hardware/software and related expenses, such as networking, and recurring needs, to address expanded or new staffing needs or hardware/software maintenance expenses. While a significant amount of startup capital and recurring funding is being provided by the university administration and participating colleges and departments, it is unlikely that the existing funding sources can support additional demands resulting from the center's acceptance of responsibility for the management of assets contributed by researchers or research programs, absent some additional source of incremental funding to accommodate the growth in the size of the asset base.

There are generally three methods by which additional funding can be obtained. First, a budget request for necessary incremental funding can be presented through the university budget process. This imposes a requirement that the university leadership prioritize the request alongside all other budget requests offered for consideration, and ties the availability of funding to the uncertainties of the university budgeting process.

A second approach is to commit a percentage of research overhead dollars to support of the center, with annual incremental increases tied to the overall growth of the research revenues of the university.

A third approach is to establish a direct cost recovery, or chargeback, model for services and support offered to the research community, and paid directly from grant revenues. The rates charged for services could be offered on an a la carte basis (cycles, storage, bandwidth, or staff hours), or a 'blended' rate could be offered on an annual basis, and executed through a Service Level Agreement.

Contributions of assets (hardware, software, staffing) to the center could offset some of all of the costs to individual researchers or research groups, assuming agreement on appropriate value of the contributions could be achieved.

Rates would need to be established that provide for sufficient funding to maintain support levels across all center operations at adequate levels, and may need to be renegotiated on an annual basis. Establishment of the rate structure could be a responsibility of the faculty advisory committee.

See Exhibit A "*Operating Budget Proposal*" for CRC budget detail and potential funding sources.

Transfer of Assets

To leverage university investment into the CRC, the transfer of assets from grant of individuals or groups of PI's to the center are encouraged. The details of a transfer will

be decided by the center director and the executive board. Two main types of transferable assets are envisioned:

- **Hardware assets:** These are computers or storage facilities that are purchased using funds of individual or groups of PI's and are incorporated in the center for use by all research groups associated with the center. The center assumes responsibility of system administration, but not of repairs or maintenance of the donated equipment. In exchange, the PIs allocation of the total resource is increased in proportion to the contribution for the lifetime of the donated equipment (typically 4 years).
- **Monetary assets:** Individual or groups of PI's can transfer monetary funds to the center to help defray personnel costs. In return, the appropriate portion of the time for one of the center professional staff is allocated to work on a project of interest to the PIs and the center.

Educational/Outreach

In addition to the primary mission of the Center, which is to support the research agenda for the University, the CRC will also strive to enhance the teaching and learning environment at Notre Dame, and to strengthen linkages between research and teaching. This may be done in several ways, including:

- 1) Operation of high-end desktop workstation clusters that will meet curricular needs, and can also be used for research.
- 2) Development of interdisciplinary courses focused on various aspects of research computing. By development of a core curriculum of such courses, development of graduate concentrations or undergraduate certificate programs in computational science and engineering may also be possible.
- 3) Sponsorship of a seminar series focused on topics relevant to research computing.
- 4) Outreach to the overall Notre Dame and local community, perhaps through the use of high-end visualization tools.

Space

There have been many conversations around optimal locations. The three components of necessary space include staff offices, collaborative work space, and data-center quality floor space for the technology assets. Our recommendation about space addresses all three areas:

- 1) **Administrative offices.** Offices for the Director and administrative staff would be accommodated within the first floor of the Information Technology Center. Also included in this space would be a single conference room for working group sessions.
- 2) **Data center space.** Through an existing agreement with Union Station, much of the CRC hardware (other than office equipment) would be located within their

data center. Some of the equipment currently on campus (in colleges and ITC Data Center) would be moved to Union Station.

- 3) Space for technologists would be broken into two components:
 - a. System engineers would be located both at Union Station and the ITC.
 - b. Discipline specific technologists would be located within the colleges, the ITC, and existing research centers.

If Notre Dame is committed to significantly increase research activity, those who have contributed to this report feel strongly that a new multidisciplinary center for research computing is critical to our success. Notre Dame's fundamental aspiration "to be one of the great universities of the world" requires our continued tradition of research innovation such as those carried out by our University founders.