

## HIGH PERFORMANCE COMPUTING INITIATIVE

### What Is It?

High Performance Computing (HPC) applies to a variety of computer configurations that allow multiple processors to be linked in order to perform parallel activities at very high speeds. By using common, off-the-shelf equipment linked together into a single system it is possible to build a relatively inexpensive system capable of performing well over a trillion floating-point operations per second. This is especially useful for scientific research and calculating engineering variables in virtual prototypes. HPC has recently been applied to business uses such as data warehouses and transaction processing.<sup>1</sup>

### Technical Issues

HPC systems take up space, use a lot of power, and generate high amounts of heat.<sup>2</sup> Rooms equipped to handle HPC should have redundant environmental controls to avoid data loss, dedicated power sources (with appropriate emergency backup power), and should be linked to high speed networks. The types of research suited to these systems are highly collaborative and highly computational efforts. In order to ensure facile exchange of large data files at speeds necessary to take advantage of HPC systems, network connections should exceed 100 megabits per second. Co-location of HPC systems would eliminate maintenance variables in providing environmental controls. This would also reduce the total cost of ownership for systems supported by individual grants. Information that follows is from institutions currently operating HPC centers.

### Penn State University

PSU operates what is called a Co-Location Center (<http://css.its.psu.edu/es/colo/>). Effectively, the center sells volume, conditioning and good networking on a raised floor. There is authorized access only and someone can choose to buy administration services as well. Center staff do not do modifications to users facilities, but will either let them buy into the central cluster or operate a small handful of machines in the co-lo center. [Operating funds] come out of research funds or departmental funds. The service is cost-recovered.

The ratio of research cycles to teaching cycles is exceedingly high but access to the resource for teaching is always provided. Some central IT dollars were used in acquiring the resources. That investment is amplified with dollars from various research groups and faculty around the university. Clusters have made this possible because acquisitions can happen in a fairly fine-grained manner. A recent 176 node cluster has funding from 12 different investigators.

---

<sup>1</sup> Wikipedia, the free encyclopedia, April 11, 2007.

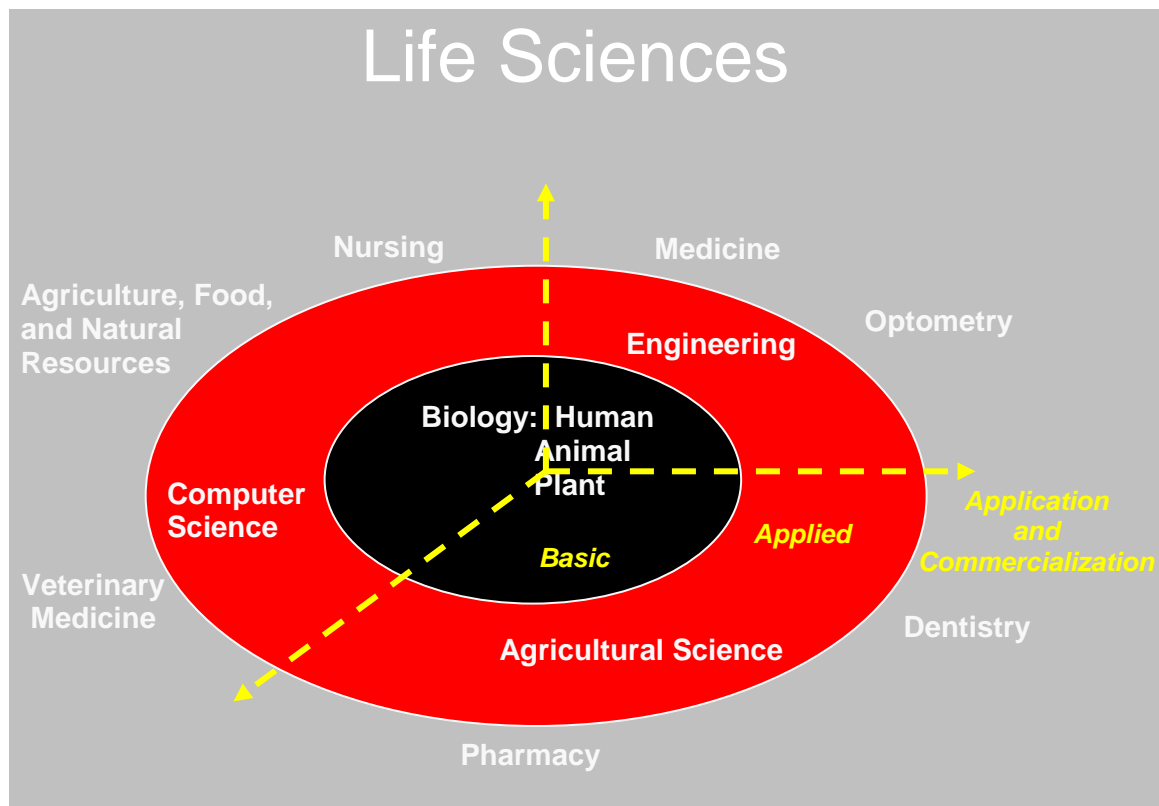
<sup>2</sup> It is possible today to design an HPPC cluster of 36 nodes plus network switches in one seven-foot-high rack. If each node draws 1.5 amps, then the heat to be dissipated would be 6.27 kw. The rack and surrounding space might take up 15 sq. ft. of floor space, thus giving a load of 418 watts/sq. ft. Few computer rooms are designed for this high a load density. MIT Final.IntensiveReport,3-2003, MIT Website

The staff that support these central resources are funded centrally and are not cost recovered. They will from time to time, be written in at reasonable percentages as “in-kind” for things like NSF grants.

Since central dollars help to fund the resource, anyone can get an account on it. There are quite a few faculty who have done "test drives" and made purchasing/technology decisions of their own based on that. Center staff help get them started and provide general consulting for them but do not get involved in day-to-day support. It has been regularly demonstrated over a 10+ year period, the cost advantage of collaborating to acquire and run resources. Still, some researchers prefer to run things for themselves and staff try to help them too. Additionally, the center has some site license contracts that they are able to take advantage of in the running of their clusters.<sup>3</sup>

### University of Missouri

The University seeks to maintain a nationally competitive position in research and scholarship in targeted areas – including the Life Sciences. In order to meet that goal, it is recognized that benefit will be gained by leveraging investment, both financial and intellectual, across all four campuses and with outside partners. Indeed, the synergies of cross-institutional partnerships will very likely result in programs and accomplishments that simply would not be realized if attempted by single campuses working in isolation.



<sup>3</sup> Penn State Official Response to Task Force questionnaire on campus HPC computing configurations, MIT, March, 2003.

We are leveraging current investment through MOREnet's provision of high-bandwidth network connectivity among the four campuses. MOREnet also provides access to the Great Plains GigaPop in Kansas City thereby connecting the System campuses to the Internet2 Abilene network, which provides access to the entire Internet2 community – currently comprised of more than 190 universities working in partnership with industry and government to develop and deploy advanced network applications and technologies.

The campuses contribute by providing support for their individual faculty engaged in bioinformatics or other Life Science research, and local technical support for facilitating access and use of bioinformatics tools provided by the Consortium.

High-performance computational systems to analyze massive sets of data have been created centrally. Very large storage devices to house major data collections. High speed networking services to facilitate location-independent access and collaboration among investigators. New technical support staff, such as database managers and networking support personnel to fully realize and leverage the investment in this Center, have been hired. This has provided resources across the UM System in the form of accessible and sharable computational capabilities and databases for collaborating investigators on any campus. This has fostered the development of advanced information technology applications in the Life Sciences in collaboration with peer Internet2 institutions. It has also facilitated training of a new generation of students and scientists in the interdisciplinary areas of Bioinformatics and computational biology. Life sciences research initiatives statewide will benefit from enhanced networking facilities for fast, reliable data exchange and access to data repositories. Many of the proposals for which letters of intent were submitted to this MLSRA initiative are dependent upon networking, computational and storage resources. Through enhancements requested here, MBRN will place these networked resources in reach of Missouri researchers and leverage them for support from funding agencies and industrial partners.<sup>4</sup>

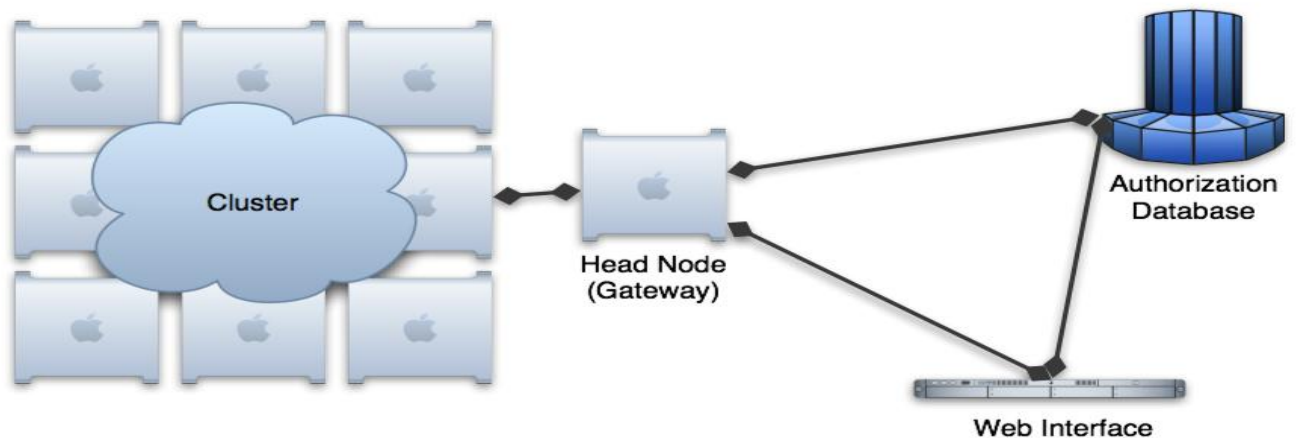
## **MIT**

A student or faculty member connects to our Web Interface with any standard browser. Users are authenticated with their MIT client certificate. Users can select which cluster computer and which application they would like to use. Users can then make requests for computing resources to be allocated for them. The request will be granted based on their level of authorization for that cluster and application.<sup>5</sup> A diagram of this appears on the next page.

---

<sup>4</sup> Copyright Gary K. Allen, 2002. This work is the intellectual property of the author. Permission is granted for this material to be shared for non-commercial, educational purposes, provided that this copyright statement appears on the reproduced materials and notice is given that the copying is by permission of the author. To disseminate otherwise or to republish requires written permission from the author.

<sup>5</sup> Copyright Phillip Long & Patrick Dreher, 2004. This work is the intellectual property of the authors. Permission is granted for this material to be shared for non-commercial, educational purposes, provided that this copyright statement appears on the reproduced materials and notice is given that the copying is by permission of the author. To disseminate otherwise or to republish requires written permission from the authors.



## New Mexico University

### Centralized Computational Center

The Center for High Performance Computing at the University of New Mexico: Our goal is to provide the infrastructure and facilities needed to continue the growth of computing based research at UNM. We strive to foster new, interdisciplinary collaborations, based on computation and to encourage novel applications of computation in research while continuing to grow traditional applications of computing in the science and engineering disciplines.

The center was established in 1994 as a sister center to the Maui High Performance Computing Center which was, at the time, run by *Frank Gilfeather*, *John Sobolewski* and *Brian Smith* from the University of New Mexico. In its short history, the UNM Center has enjoyed a number of fairly significant firsts, including: the first Linux cluster available through NSF allocation and the first use of Access Grid technology. For more history, see the [history page](#).

The center serves as the academic unit in charge of the **Computational Science and Engineering (CSE) Certificate** program, a graduate degree certificate. *Marc Ingber* from the Mechanical Engineering Department and *Tom Hagstrom* from the Mathematics department serve as the coordinators for this program. For more information about the CSE program, see the [CSE page](#).

The UNM research community is making extensive use of computing in their research activities. The [ARTS \(Art, Research, Technology, and Science\) Lab](#) represents an exciting confluence of art and technology with an emphasis on computation. ARTS Lab was created in response to New Mexico Governor Bill Richardson's [Media Industries Strategic Plan \(MISP\)](#).<sup>6</sup>

---

<sup>6</sup> <http://www.hpc.unm.edu/>

## Training

As technology changes and as instructors' and researchers' needs for computational resources change, standard training programs in the use of HPC hardware, software, and account management is necessary. Here is an example of regular educational offerings at the Center for High Performance Computing at the University of New Mexico:

## Education<sup>7</sup>

- [Courses](#)  
The center offers training and workshops for UNM faculty, staff, students and others.
- Access Grid classes  
During the year we host a number of conferences, classes, and other activities over the Access Grid. To schedule an access grid session, call our [main number](#).
- [LCI information](#)  
Several members of our staff are involved in planning and teaching in the Linux Clusters Institute.
- [CSE information](#)  
The center serves as the administrative unit for the graduate certificate in Computational Science and Engineering (CSE).
- UNM related classes  
UNM courses that are relevant to the center include
  - CS442 [Introduction to Parallel Systems](#)
- Books
  - [Using MPI : portable parallel programming with the message-passing interface / William Gropp, Ewing Lusk, Anthony Skjellum Cambridge, Mass. : MIT Press, c1999](#)

## Total Cost of Ownership

1. The Gartner Group and other industry watchers indicate that the hardware component of an HPPC project is only about 15 to 20% of the total cost of operations.
2. Maintenance over the total life cycle of the project is the majority cost.

A review of what ancillary costs are involved in running an HPC system shows why they are often surprisingly high.<sup>8</sup>

- Power and cooling requirements: With respect to cooling a sample survey has shown that the costs for cooling can vary from \$261 per kilowatt to \$1,873 per kilowatt where redundant cooling is supplied; and from \$257 to \$2,786 per

---

<sup>7</sup> <http://www.hpc.unm.edu/education>

<sup>8</sup> MIT Final.IntensiveReport,3-2003, MIT Website

kilowatt where there is only a single unit, the loss of which would leave no air conditioning.

- **Building renovations:** Changes to the building can be required at several stages. First, there is the problem (sometimes severe) of finding a space for the initial installation of the equipment. Then, as the project expands, and still more power and cooling are required, further modifications to the physical plant may be needed.
- **Operations support:** This includes system administration, computer security, maintenance of hardware, and miscellaneous operations of the facility (such as hardware replacement, inventory of spare parts). The customized design of both education and research HPPC clusters requires more hands-on involvement because of the unique attributes of each cluster.
- **Turnover:** Those who are most familiar with a given system leave, taking their knowledge and experience with them. These systems, unfortunately, are rarely fully documented, and so expertise needs to be reacquired.

The above items can be quantified. Opportunity costs are far more subjective, yet they are no less real. Projects requiring significant building preparation can cause significant delays in conducting research and expanding academic programs.

### **Recommendation**

Well conditioned room (redundant power and environmental controls) for locating, and eventually co-locating, a clustered HPC environment. This means that operations space must be expandable within the unit. The room should be connected to Internet 2. The HPC Center should be staffed for system management and continuity. Centralized training and support should be offered for all who would use the facility.



