

Midrange Computing Workshop

Building 66 Auditorium, Tuesday, March 26, 2002

Executive Summary

At the end of a daylong series of presentations and discussions, a group of about 40 scientists and staff came to series of preliminary conclusions regarding efforts to advance midrange computing at LBNL. The concepts of centralized consulting for the purchase of clusters, providing a central room to house clusters, and a “matchmaking” service to link potential partners in the purchase and use of clusters, generated interest from the participants. As currently estimated by ITSD’s UNIX Support Group, the cost of cluster system administration seems to be too steep for most Lab groups. The group is also interested in seeing if another model exists, and benchmarking Lab services and options with those provided by similar institutions. The workshop concluded with organizers agreeing to study the outstanding issues and come up with potential paths forward. The group also seemed interested in identifying the Laboratory-wide requirements for Mid Range Computing and whether the case could or should be made for such a facility or service.

Organizers:

The MRC Working Group:

Paul D. Adams	Physical Biosciences
Jon Bashor	Computing Sciences
Ali Belkacem	Chemical Sciences
Alessandra Ciocio (Chair)	Physics
Kenneth H. Downing	Life Sciences
Gary Jung	Information Technologies and Services
James F. Leighton	Information Technologies and Services
Alexander "Sandy" Merola	Information Technologies and Services
Douglas L. Olson	Nuclear Science
John W. Staples	Accelerator and Fusion Research
Shaheen Tonse	Environmental Energy Technologies
Michel A. Van Hove	Materials Sciences
Tammy S. Welcome	NERSC

With:

Meg Treleaven (ITSD)
Yeen Mankin (CSD)

List of workshop attendees:

http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/Attendance_List.xls

Workshop agenda:

<http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/Agenda.pdf>

Morning Session

Presentations

Welcome, Workshop Purpose

Alexander Merola, Director, Information Technologies and Services Division

ITSD Director Sandy Merola opened the presentations by presenting some background information to put current MRC efforts into perspective. Since the 1970s, computing at Berkeley Lab evolved from a centralized facility (where users couldn't even submit their own data cards) to a departmental function in the 1980s to decentralized, desktop-centric computing in the 1990s. With the arrival of desktops, however, centralized computer support, at the Lab bottomed out, until there began a resurgence several years ago. Now researchers are expecting more support, asking for more affordable backup services, support for clusters and web servers.

The goals of the midrange computing workshop are to:

- Exchange information, and
- Identify what the computer support community and/or researchers can do to advance MRC.

If there is interest in getting institutional help for an MRC resource, the potential user group will first have to get institutional support. A more logical approach is to make it a "grass roots" effort and grow it to a level that attracts institutional interest. There is already a shared MRC model at LBNL -- the PDSF system operated by NERSC to support researchers in Physics and Nuclear Science.

Among the options available are:

- Keep providing the current level of support.
- Provide increased system support (but not scientific support).
- Create a centralized system, administered and shared, using capital and operating funds.
- Establish an institutional resource, whether as a recharge or overhead program, by:
 - Building on PDSF, or
 - Using the existing Alvarez cluster
 - Possible agreement with NERSC

Scientific Computing at the Berkeley Lab

C.William McCurdy. Associate Laboratory Director for Computing Sciences

Bill McCurdy gave an informal overview of midrange computing in the DOE complex. He suggested, in what became a recurring theme, that the group seek a solution that solves part of the MRC problem for part of the Lab (i.e. one size doesn't fit all). Bill noted that midrange computing is important, but is rarely anyone's top priority. To gain institutional support, the group needs to convince a couple of division directors to say MRC is important.

In the ensuing discussion, a suggestion was made to investigate purchasing extra nodes for NERSC's IBM SP, and to dedicate them to LBNL users. One desirable aspect would be fast turnaround for on-the-spot computing linked to experimental facilities, such as the ALS.

Midrange Computing Working Group Process and Goals

Alessandra Ciocio, MRC Working Group Chair, Physics Division

http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/presentations/A_Ciocio.pdf

Alessandra gave an overview of the working group's creation and efforts since early 2001. From the CSAC recommendation for the development of a strategic plan for MRC at LBNL in addition to the LBNL-NERSC-DOE agreement, that created the MRC WG in the fall of 2000, the MRC process went through a first phase that culminated in a document (<http://www.lbl.gov/ITSD/CSAC/MRC/Reports/>) describing the assessment and findings from this group and a letter to ITSD with initial recommendation for a path forward.

The second phase consisted on following the evolving situation of scientific computing at Berkeley Lab and to identify a number of options for institutional support of midrange computing for the Lab's scientific programs. Given the fact that the number of small clusters is increasing, what can be done to make the growing midrange computing at LBNL a consistent success across the whole Laboratory?

A Survey of current and potential midrange computing users within the various divisions was then conducted to help determine the institutional requirements for MRC and/or support services. And with this Workshop, which is bringing together current and potential MRC users, for a discussion of MRC users requirements and needs, options and identified offerings, the MRC WG hopes to help evaluate and determine the best path forward for Scientific Computing at Berkeley Lab.

**Overview of Current Midrange Computing at LBNL:
Computer Clusters
Gary Jung, UNIX Support Group Lead, ITSD**

http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/presentations/G_Jung.pdf

In trying to gauge the demand for MRC at the Lab, Gary Jung conducted an informal survey to learn about existing clusters and systems with more than four processors. McCurdy emphasized that to make a compelling argument about the contributions of and need for MRC, the group needs to compile a complete list of existing MRC resources currently in use at the Lab. This information is also important for computer security purposes, he said.

**Overview of Current Midrange Computing at LBNL:
Scientific Projects
Ali Belkacem, MRC Working Group member, Chemical Sciences Division**

http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/presentations/A_Belkacem.pdf

The goal is to make scientific computing a success for the Laboratory. Groups are spending the money already, so it is an important effort and computing is completely integrated with many research projects. Belkacem said that scientists who use, or want to use, MRC need to find the solution. It has to come from scientists, not management, or ITSD.

**PDSF Model
Craig Tull, MRC Working Group Member, NERSC High Energy and Nuclear Physics Group**

http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/presentations/G_Tull.pdf

PDSF, the Parallel Distributed Systems Facility, is a cluster that grew out of the now-cancelled Superconducting Supercollider. The PDSF is a cluster consisting of almost 400 processors, having been constantly upgraded and expanded with hardware paid for by users of the system.

Tull noted that PDSF is not suited for all jobs and is not currently tuned for massively parallel processing. However, the system could be used as a model for other MRC efforts, as long as the appropriate elements are adopted. It may be possible to expand PDSF to accommodate new MRC users, and while this would give MRC a faster ramp-up, there are some ownership issues to be addressed.

Afternoon Session

Survey Results, Options, and Roundtable Discussion Led by Sandy Merola

<http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/presentations/AXMerola.pdf>

The discussion noted that there is a range of interest and readiness for using an MRC resource. One area of needed support would be in helping scientists parallelize their existing codes to take advantage of parallel computing. In general, centralized support was identified as a real need. New experiments coming on line will also benefit from greater computing capacity.

The Path Forward:

The key issue is how to determine if a shared resource or institutional support is worth pursuing? One possibility is to invite potential users to try out existing systems, such as PDSF, the alvarez cluster and the new CIS cluster. Also determine if the general need is serial or parallel processing.

Final Roundtable Comments:

The workshop concluded with each attendee invited to make a final comment.

The general areas of agreement were:

- Keep investigating and discussing ways to provide an MRC.
- Use PDSF as a model for a shared resource.
- Providing a cluster for users to “test drive” would help assess the overall demand.
- Not one size of cluster will meet everyone’s needs.
- Creating a centralized, shared resource will be difficult without institutional support.
- Pre-purchase consulting and “matchmaking” between partners would be useful.
- The ability to use a centralized computer room is important.
- More financial information about the various options is needed.

Detailed comments and roundtable discussion can be found at:

http://www.lbl.gov/ITSD/CSAC/MRC/Workshop/proceedings/Roundtable_notes.doc