LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Institutional Computing Executive Group Review of Multi-programmatic & Institutional Computing, Fiscal Year 2005 and 2006

S. Langer, D. Rotman, E. Schwegler, P. Folta, R.
Gee, D. White

December 20, 2006

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Intra-laboratory memorandum

Mail Stop L-103

Ext: 2-7746

Fax: 3-4908

December 18, 2006

To: Rokaya Al-Ayat, Cherry Murray, Dona Crawford, Michel McCoy
Cc: Laboratory Associate Directors
From: Steve Langer, Doug Rotman, Eric Schwegler, Peg Folta, Richard Gee, Dan White
Subject: ICEG review of Multi-programmatic & Institutional Computing

The Institutional Computing Executive Group (ICEG) review of FY05-06 Multiprogrammatic and Institutional Computing (M&IC) activities is presented in the attached report. In summary, we find that the M&IC staff does an outstanding job of acquiring and supporting a wide range of institutional computing resources to meet the programmatic and scientific goals of LLNL. The responsiveness and high quality of support given to users and the programs investing in M&IC reflects the dedication and skill of the M&IC staff. M&IC has successfully managed serial capacity, parallel capacity, and capability computing resources. Serial capacity computing supports a wide range of scientific projects which require access to a few high performance processors within a shared memory computer. Parallel capacity computing supports scientific projects that require a moderate number of processors (up to roughly 1000) on a parallel computer. Capability computing supports parallel jobs that push the limits of simulation science.

M&IC has worked closely with Stockpile Stewardship, and together they have made LLNL a premier institution for computational and simulation science. Such a standing is vital to the continued success of laboratory science programs and to the recruitment and retention of top scientists. This report provides recommendations to build on M&IC's accomplishments and improve simulation capabilities at LLNL. We recommend that institution fully fund (1) operation of the atlas cluster purchased in FY06 to support a few large projects; (2) operation of the thunder and zeus clusters to enable "mid-range" parallel capacity simulations during normal operation and a limited number of large simulations during dedicated application time; (3) operation of the new yana cluster to support a wide range of serial capacity simulations; (4) improvements to the reliability and performance of the Lustre parallel file system; (5) support for the new GDO petabyte-class storage facility on the green network for use in data intensive external collaborations; and (6) continued support for visualization and other methods for analyzing large simulations. We also recommend that M&IC begin planning in FY07 for the next upgrade of its parallel clusters.

LLNL investments in M&IC have resulted in a world-class simulation capability leading to innovative science. We thank the LLNL management for its continued support and thank the M&IC staff for its vision and dedicated efforts to make it all happen.

FY05-06 Review of Multiprogrammatic and Institutional Computing

*Steve Langer, Doug Rotman, Eric Schwegler, Peg Folta, Richard Gee, Dan White
(for the Institutional Computing Executive Group)*

December 2006

Auspices

The Institutional Computing Executive Group (ICEG), whose members are selected by their respective organizations, reviews M&IC activities every year or two. This report covers both FY05 and FY06 activities. The ICEG also provides guidance on future institutional computing needs. This report summarizes the assessments, conclusions, and recommendations of the ICEG. These findings derive from an analysis of written and oral documentation provided by LC and M&IC staff, discussion at the annual meeting (held June 15, 2006), and informal polls of "seasoned" users of M&IC resources.

Current ICEG membership:

Jim Rathkopf – Defense and Nuclear Technologies, AX Division
Steve Langer – Defense and Nuclear Technologies, AX Division
Eric Schwegler – Physics and Advanced Technology
Peg Folta – Biosciences
Felice Lightstone – Biosciences
Joe Morris – Energy and Environment
Doug Rotman – Energy and Environment
Richard Gee – Chemistry and Materials Sciences
Christian Mailhot – Chemistry and Materials Sciences
Dan White – Engineering
Don Dossa – Computation

Brian Carnes - Chair
Michel McCoy – Technical Advisor
Mark Seager – Technical Advisor
Greg Tomaschke – Technical Advisor
Doug East – Technical Advisor

Background

The mission of Multi-programmatic and Institutional Computing (M&IC) is to provide leading edge computing resources to *all* LLNL programs. Since its creation in FY97, M&IC has used institutional and programmatic investments to enable Livermore Computing (LC) to deploy a high-performance computing infrastructure that supports large-scale simulations across a broad range of scientific and engineering applications.

M&IC has helped to realize the Laboratory Director's desire (1) to extend the use and sophistication of simulation as a core competency for LLNL; and (2) to enfranchise all LLNL programs in this endeavor.

M&IC resources include 3 classes of systems – serial capacity, parallel capacity, and capability. Serial capacity resources provide for the many simulations that can be run on a single shared memory computer. Parallel capacity clusters enable simulations that require a large number (hundreds) of processors tied together by a fast interconnect and are potentially available to the whole M&IC user community. M&IC runs its largest system as a capability cluster that supports a few grand challenge applications. Procurements carried out at the end of FY06 will cause major changes in all three classes of systems.

M&IC constantly works at balancing investments to meet the widespread demand for capacity computing with providing cost-effective capability platforms. M&IC leverages the computing and storage resources of the Stockpile Stewardship Program (SSP) for the mutual benefit of the SSP, other programs, and the Institution.

The GPS and ILX serial capacity clusters will both be shut down by early in CY07 and will be replaced by the Yana cluster. The Yana cluster consists of shared memory multi-processor (SMP) systems using four dual core AMD Opteron processors and running the Linux operating system. Nodes have either 16 or 32 GB of memory. The Yana cluster supports jobs using up to 8 processors when using OpenMP or other threaded approaches to parallelism.

The mcr parallel capacity cluster will be retired by early CY07 and replaced by the thunder cluster (re-purposed from capability computing) and the new zeus cluster. Thunder uses Intel Itanium processors, a Quadrics interconnect, and runs Linux. Zeus uses nodes like those in the Yana cluster and adds an Infiniband interconnect to support distributed memory parallel computing.

The new atlas cluster will become M&IC's capability cluster by early in CY07. The atlas cluster uses the same processor and interconnect as zeus and has roughly twice the theoretical peak performance of the thunder cluster. When the mcr and thunder clusters were first procured they made the top 5 on the Top 500 list of the fastest systems at running the Linpack benchmark. Atlas would have ranked around 15th on the November 2006 list if a Linpack run had been completed in time for an official entry.

Atlas does not have as high a relative ranking as mcr and thunder did when first introduced, but it should be competitive with systems at other unclassified research centers. The ICEG does not believe that the lower relative ranking of atlas is a cause for alarm. The widespread acceptance of affordable commodity based clusters has greatly increased the number of computer centers with large clusters. The gaps between systems in the top 10 are smaller than they were a few years ago. The expected fierce competition for time on thunder and zeus is likely to be a bigger issue for ICEG users.

The new GDO data storage facility on the “green” network is intended to enable data intensive collaborations with external scientists. This service made its debut in FY06 and has already been very helpful in sharing climate model data.

M&IC Usage and Resources

No new clusters reached general availability (GA) status in the period covered by this report, but several new clusters were purchased at the end of FY06. The new clusters will replace all existing M&IC systems except for thunder. Taken together, this represents a major move towards the x86_64 processor. The AMD Opteron used in these clusters and the Intel EM64T are (mostly) software compatible, so M&IC could use Intel processors in the future with only modest impact on users.

Serial capacity resources had not changed significantly since the upgrade of ilx in July of 2003. The retirement of the existing ilx and GPS clusters and their replacement by the Yana cluster is a significant change. Serial capacity users should be familiar with systems using 64 bit addressing (GPS) and systems running Linux (ilx), but some assistance with migration will likely be required to get codes running under 64 bit Linux. We encourage M&IC to pay attention to this issue in addition to helping the users of the new parallel clusters. We also encourage M&IC to investigate creating two batch queues on Yana – one with a 12 hour time limit to provide prompt access to a batch slot and another queue with a longer time limit (48 hours?) for users whose runs can not easily be broken up into 12 hour chunks. Interest in serial nodes with more memory and more processors was expressed even before Yana became available to users. M&IC should monitor this area – it should be possible to meet these needs with a relatively modest investment.

Many highly successful simulations were carried out on the MCR and Thunder clusters at or near full system scale. However, the development of a number of very successful parallel codes led to a demand that far outran the available resources. A white paper describing the success of the MCR and Thunder clusters and the kinds of science that could be performed with additional computing resources was presented to LLNL management in August 2005. Laboratory management approved the procurement of the 40 TFLOP/s (peak) Atlas cluster and the 10 TFLOP/s (peak) Zeus cluster to help meet the needs described in this white paper. The Atlas cluster will be used for “grand challenge” simulations. The selection of the first grand challenge submissions for Atlas had not been completed at the time of this writing.

The Zeus and thunder clusters will support parallel capacity computing. The ICEG recommends that M&IC consider tuning one of these systems to favor small parallel jobs (roughly 16-64 processors) and another to favor large jobs. Experience indicates that it is very difficult to support both classes of jobs from the same batch queue.

The ICEG report on FY02 endorsed M&IC’s strategy for developing Lustre, a new parallel file system. The Atlas, Thunder, and Zeus clusters as well as the Prism visualization cluster use Lustre. Lustre was also chosen as the parallel file system for

ASC's Blue Gene/L. LLNL, in partnership with other NNSA sites and industry, has been responsible for the development of Lustre. Lustre is based on Open Source software and may be used with clusters made by a wide variety of vendors. A single Lustre file systems can be mounted on multiple clusters simultaneously. Lustre has made a great deal of progress since FY02, but additional work needs to be done on performance, reliability, and features.

In this report, "Lustre" means the combination of hardware and software making up a parallel file system. Problems with Lustre performance and reliability were the one common theme from all directorates providing input for this report. Only a few M&IC users run parallel jobs on thunder, so the problems are probably based mostly on experience on mcr. One possible problem in the area of performance is the effect of load imbalance. Load imbalance induced by contention with other jobs might be more severe on a cluster like mcr that runs many jobs simultaneously than on a cluster like thunder that has been operated in capability mode.

Lustre file systems had reliability issues throughout the period covered by this report. M&IC implemented a new purging policy in 2006 that should reduce (but not eliminate) the reliability problems that occur when file system components become too full. The ICEG recommends that M&IC continue to monitor the effects of OST's filling up and search for other sources of instability.

A bothersome "feature" of Lustre is that it can return from a read operation with less than the requested number of bytes. This behavior is permitted by the standards, but does not occur on other popular file systems. Lustre forces all users to implement error handling to deal with this behavior in their own codes. It would be more efficient to implement this code once in Lustre. There is no known benefit to users in returning early with incomplete data. The ICEG recommends that Lustre be modified to wait until it can return the requested number of bytes, if those bytes exist in the file.

Linux based visualization clusters have been the mainstay of M&IC visualization for the past several years. The ICEG recommends that M&IC continue to pursue this approach to provide visualization support for new compute clusters. User interviews revealed that one of the problems with the current visualization clusters is the performance and purging policy of the parallel file systems on the production clusters. The ICEG recommends that M&IC investigate the benefits of file systems dedicated to visualization. A dedicated file system could have a longer purge time to help in the analysis of large "science runs" and deliver higher performance either directly or by serving as a cache for the file system on a compute cluster.

In summary, the M&IC staff has dramatically enhanced the LLNL capability computing resources through the procurement of the Atlas cluster. Parallel capacity resources have been significantly enhanced through the procurement of the zeus cluster and the repurposing of the thunder cluster. The Yana serial capacity cluster is significantly more powerful than the ILX and GPS clusters it replaces. The new GDO storage system

supports data intensive collaborations with external scientists. LC responded very well to the recommendation to enhance the capability computing environment in the ICEG's report on FY03-04.

Recommendations

The ICEG recommends the following FY07 M&IC activities in priority order.

1. The ICEG recommends continued support for the operation of the existing M&IC clusters.
2. The ICEG recommends continued work on the Lustre parallel file system with an emphasis on improved stability and enhanced performance
3. The ICEG recommends that M&IC begin planning for the next generation of clusters. We expect the new Atlas and Zeus clusters to become saturated in short order. One possibility is a processor upgrade for the existing clusters in a year or so and the purchase of a new capability cluster in 18-24 months.
4. The ICEG recommends that M&IC evaluate the use of the new petabyte class storage system on the green network and develop an upgrade path if necessary.
5. The ICEG recommends that M&IC evaluate an upgrade to the Yana cluster to provide more processors and memory in a single node. The obvious candidate is a four socket, quad core Opteron node with 64 GB of memory.
6. M&IC should continue to acquire Linux clusters to provide visualization support for compute clusters and should investigate the benefits of separate visualization file systems.

Challenges and Concerns

Through strong and consistent LLNL investments, M&IC has grown into a powerful unclassified computing resource being used across the laboratory to push the limits of computing and its application to simulation science. Through these efforts, LLNL has become a premier laboratory in simulation science resulting in world-class scientific insight and the recruitment and retention of leading physical and computer scientists. All LLNL programs and projects are bolstered through these efforts. To maintain this leadership position, we strongly urge LLNL to continue to invest in the M&IC at a level comparable to FY05 and FY06. The Atlas, Thunder and Zeus clusters provide LLNL scientists with access to world class unclassified systems. Future spending plans should include an ongoing commitment to capability systems sufficient to give LLNL researchers access to systems competitive with those at other unclassified computing centers. Future plans should also devote significant resources to providing resources to

support parallel and serial capacity computing, the key day-to-day requirement of compute intensive research at LLNL.

The number of transistors on a chip continues to follow Moore's law of doubling every 18-24 months. The clock speed of processors is increasing at a much more modest rate. Vendors now increase the performance of chips by increasing the number of processors (cores) on a chip. This will lead to at least two challenges for code developers. First, the number of processors in a cluster will increase so grand challenge codes will need improved scalability. Second, the total computing power on a chip is increasing faster than the memory bandwidth available to the chip. We encourage M&IC to investigate methods to effectively use these new systems and communicate those methods to the M&IC code development community.

The emphasis on open source software and commodity hardware and interconnects has allowed M&IC to deliver world-class systems for a remarkably low price. These systems have been quite productive, but they appear to be less reliable than commercial systems from IBM. Power supplies and some other components appear to be less reliable than those used in systems like white and purple. The ICEG feels that the price/performance benefits of the commodity approach are sufficient to warrant its continued use, but more emphasis on long-term support may be required.

Summary

M&IC's performance continues its tradition of providing outstanding services and access to world-class capacity and capability resources. Through their efforts, LLNL stands at the forefront of unclassified simulation science. We strongly recommend that LLNL invest in M&IC resources to maintain this community standing.

The ICEG recommends that the proposed FY07 activities be fully funded. The ICEG urges M&IC to begin planning for new capacity and capability computing resources with a target of purchasing systems in FY08. The ICEG urges M&IC to investigate methods of improving the reliability of Lustre. The ICEG recommends that M&IC upgrade (as need demands) the petabyte class GDO storage facility on the green network to support external collaborations.

Finally, the ICEG wishes to thank the Institution for its investment in, and dedication to, high-end computational resources and simulation science.