

To: Dr. Ray Orbach

From: ASCAC Subcommittee on Future Large Facilities

Preliminary ASCAC Facilities Recommendations

Introduction

To support the leadership position of the Department of Energy in scientific research, it is necessary for the Office of Science to maintain premier advanced scientific computing facilities. The essential need for high-end computers, networking, storage, special application software, algorithms and applied mathematics have been documented in the various department publications. (Do we need to reference them here?)

It is clear from current events that current resources cannot support Office of Science computing needs. For example, SciDAC researchers requested twice as many computer hours for their projects as they were allocated and expect to require more as the projects mature. Furthermore, new initiatives such as the Fusion Energy Simulation Initiative, the Nanoscience Master plan and the Genomes to Life program envision the need for substantial new computational resources that far exceed any currently planned facilities.

As ASCAC has noted in previous communications, satisfying the advanced computing needs of the Office of Science, requires a balanced, thoughtful, well-managed, but also aggressive plan for the growth and development of a comprehensive set of computational resources. Though the rapid pace of development in computing technologies makes it very challenging to predict with certainty the technological details of future systems, it enhances the need for investment (and adventure). Fortunately, the last several years of study, planning, and analysis of computing needs has laid the foundation for a systematic approach to this task. Developing and managing the computing capabilities and capacities required for success of all aspects of the Office of Science's missions is perhaps the most important facilities challenge we face for the next 20 years.

Recommendations

- a) The Office of Science should invest in the creation and maintenance of at least one "Leadership Class" computing facility that will support high capability computing. Leadership class systems, which will enable and sustain computation at a world-class level, must be reinvested and renewed with approximately a 3 year lifecycle. This facility would require advanced computation hardware and software support. This type of facility is required to answer the difficult scientific questions posed by the old and new initiatives. We expect that acquisition, maintenance, and operation costs for such a facility would amount to about \$300M over a 3-year timeframe. We expect at least this amount of funds would also be needed in the subsequent 3-year periods. We consider this facility to be

- “absolutely central” and would depend on the availability of “leading edge hardware”.
- b) The Office of Science needs to invest in capacity computing to satisfy the needs of many researchers who individually do not need the dedicated capabilities of the leadership class systems but whose aggregate need is substantial. It is essential to provide such capacity in order to be able to sustain the effective use of the leadership class systems for those problems they are uniquely suited to address (instead of timesharing them out into infinitesimal fragments to many people). This need can be addressed by an appropriate combination of continued use of leadership class facilities and continuous upgrade of user support software. We estimate this facility at \$200M over a 3-year timeframe. We consider this facility to be “absolutely central” and is “ready to begin”.
 - c) It is essential for the Office of Science to support experimental research in networking, computer hardware, and software for both applications and system support. To support the petaflop environments needed for both (a) and (b) above, it is especially important that we provide experimental systems at a scale that enables researchers to assess how algorithms, protocols, applications, etc scale. A separate facility for this research is advised. We estimate this activity at \$100M over a 3-year timeframe. We consider this facility as “important” and “ready to begin”.

Management

The advanced computing facilities, though they may be physically distributed among the labs and centers, must be managed and coordinated as an Office of Science resource as are other major DOE facilities. In addition, it is quite possible that these facilities could be part of a larger national open computing infrastructure that could be used by researchers across the country. DOE’s ability to provide such facilities is what makes it unique among science agencies. This is nowhere more important than in computing.

From the point of view of users, the facilities should be viewed as a coordinated, seamless entity. A user or group of collaborators should be able to have an integrated view of its computing resources linking components at the leadership class systems, development systems, and distributed systems. Also, all offices under the Office of Science should participate in allocation and planning including any special purpose computers.

The time for leveraged and planned investment in high performance computation is now. Other government agencies are currently investigating activities in this arena. The DOE Office of Science has an important role to play in deploying unique facilities that support leadership of the US and DOE in computational science. It can leverage the work of these other agencies in HPC and provide the first open instances of the scientific methods that will dominate this century.

ASCAC Subcommittee on the Future of Large Facilities for HPC

Roscoe Giles, Boston University

Helene Kulsrud, Center for Communications Research, Princeton, Chair

Michael Levine, Pittsburgh Supercomputer Center

Juan Meza, Berkley National Laboratory

Stephen Wolff, Cisco

Margaret Wright, New York University

Mary Zosel, Lawrence Livermore Laboratory

March 17, 2003