Site Visit Team (SVT)

Members of the SVT

- David C. Bader, Oak Ridge National Laboratory
- Michael E. Papka, Argonne National Laboratory
- Ralph Roskies, (Chair) Pittsburgh Supercomputing Center
- Karl W. Schulz, Texas Advanced Computing Center
- Jeroen Tromp, Princeton University
- Francesca Verdier, NERSC/Lawrence Berkeley National Lab
- Fuqing Zhang, Pennsylvania State University

NSF personnel in Attendance:

- Sarah Ruth, Anjuli Bamzai, Shaun Young, (all from AGS)

The Review Charge explicitly excluded the NCAR –Wyoming Supercomputer Center as that had been reviewed by a different external panel in Feb 2011 – this review was also available to the SVT

Major Findings

- CISL's accomplishments & plans exhibit a high quality research program that advances knowledge and the overall mission of NCAR;
- CISL team is well matched with both CISL's goal and those of NCAR;
- CISL actively and effectively engages in a wide range of activities designed to broaden the lab's impact through outreach, collaboration and community engagement and through capacity building at regional, national and international scales;
- CISL's exemplary accomplishment: Broad range of data services that benefit many users in the nation and all around the world;
- CISL provides essential domain-specific supercomputing services to the broad atmospheric science community;
- CISL's development of highly scalable algorithms DART for DA and HOMME dynamical core framework for CESM have broad impact in the geophysical community;
- Applauded CISL's direction to add a research component that includes applied mathematics, algorithm research, statistical sciences, and data sciences in IMAGe;
- Impressed with the quality and relevance of the research conducted within CISL;
- CISL management has done an excellent job managing a productive group of individuals who have a great deal of responsibility for how NCAR is viewed by the outside world.

Findings versus Charge

(The overall charge for the review was to assess the accomplishments of CISL since the inception of the current cooperative agreement, and their plans for the remainder of the cooperative agreement)

- Build, operate and maintain shared computational infrastructure and databases, and make them accessible to the community.
- Serve as stewards of high quality scientific data on behalf of the community, through maintenance, enhancement and curation.

Findings:

- *CISL has done both of these well:*
  - *Successfully deployed Bluefire, IBM's first large POWER6 system;*
  - *Successful transition from their MSS archival system to HPSS;*
  - *Computational capabilities have 15,000 users annually,*
  - *Provide excellent user support by people who have deep domain knowledge as well as a deep knowledge of algorithms, numerical and statistical methods, and optimization techniques for compute systems with large core counts;*
  - *CISL's data resources are unique and greatly enhance NCAR's value as a national center;*
  - *Especially impressive are the data archives maintained for the community. CISL has been prescient in this regard, having established its first data repository, the Research Data Archives (RDA), in 1960, and for moving it to the web starting in 2001.*

Findings versus Charge (cont'd)

(The overall charge for the review was to assess the accomplishments of CISL since the inception of the current cooperative agreement, and their plans for the remainder of the cooperative agreement)

- Facilitate the definition, construction and use of numerical models, in particular where these support community-wide collaborations;
- Address challenging scientific problems that require long term focus and integration across global, regional and local scales;
- Apply scientific understanding to recognized and emerging problems of interest to society.

Finding:

- *CISL strongly contributes to all three of these missions through its research by providing the tools and computational infrastructure that both modeling and observational data analysis require;*
- *The SVT was impressed with the quality and relevance of the research conducted within CISL. Applaud CISL's direction to add a research component that includes applied mathematics, algorithm research, statistical sciences, and data sciences in IMAGe;*
- *This group contributes to the intellectual vitality of CISL and enables collaborations with other NCAR groups and extramural partners. The publication record is excellent and the engagement of the staff within their discipline communities is high.*

Findings versus Charge (cont'd)

(The overall charge for the review was to assess the accomplishments of CISL since the inception of the current cooperative agreement, and their plans for the remainder of the cooperative agreement)

- Enable the education and development of diverse human capital for the atmospheric and related sciences, in particular, those groups which are underrepresented in the US science and engineering workforce;
- Attract, support and advance a high quality, diverse scientific and technical staff and provide them with continued training and career advancement as befits a national center and to enhance the community;
- Integrate research and education in an increasingly multi- and interdisciplinary arena.

Findings:

- *CISL engages in a variety of education and outreach activities, which are focused on advanced scientific computing and cyberinfrastructure.*

- *Highlighted by:*
  - *Summer Internships in Parallel Computational Science (SIParCS), which is a 10-week internship program designed to address the shortage of trained scientists and engineers in the computational sciences, particularly applied mathematics, computer science, and software engineering.*
  - *IMAGe's Theme of the Year program (TOY), a cross-cutting series of activities intended to explore the interface between applied mathematics and the geosciences. TOY is designed to advance research and education between the mathematical and geoscience communities*

Recommendations & Responses

1. Establish a coherent vision and strategy for IMAGe to define priorities and document the relationships between IMAGe, the other CISL groups, and rest of NCAR;

   Response: Will review our strategic planning during 2012 with a view to a potential update of the CISL Strategic plan;

2. Provide increased parallelization and code optimization support, especially to the CISL HPC user community;

   Response: Bearing in mind that CISL does not own the scientific applications and cannot unilaterally undertake to modify the codes, CISL has an underlying strategy of extensive collaborations in the area;

3. Pursue partnerships with TeraGrid/XD to provide additional training courses on scientific programming languages targeted to the atmospheric sciences (Fortran, Python, etc.)

   Response: CISL's role in the TeraGrid XD/XSEDE project is through the advanced user support service (AUSS) component. It seems sensible and within project scope to federate our HPC training activities with those of our XSEDE partners;

Recommendations & Responses

4. Benchmark OSD HPC staffing levels against current NSF centers;

   Response: Evaluating staffing levels is an important and continual effort within any organization. As part of this effort, we periodically ask other organizations about their staffing levels in various elements of their HPC organizations;

5. Explore opportunities for coordinating, streamlining, standardizing data collection, storage and distribution practices with other NSF-supported efforts;

   Response: CISL strongly agrees with this recommendation, local developments must be considered in the context of other community systems;

6. Establish and document goals for the NWSC and associated HPC deployment along with specific metrics for success;

   Response: CISL recognizes the need for a successful transition from its Mesa Lab facilities and resources to the NWSC facilities and resources. To this end, CISL is identifying specific metrics and goals through which we can document and achieve success;

7. Increase efforts to investigate code performance on emerging architectures.

   Response: CISL will re-evaluate the balance of staff allocated to evaluating emerging many-core architectures relative to those involved in the other aspects of our acceleration strategy. Will also consider expanding our partnerships in this area.

Questions?