

**Summary of
Future Directions for NSF Advanced
Computing Infrastructure to Support
U.S. Science and Engineering in
2017-2020
Interim Report**

Emily Carter, Princeton
Juan C. Meza, UC Merced

November 4, 2014, MPSAC meeting

**COMMITTEE ON FUTURE DIRECTIONS FOR NSF ADVANCED COMPUTING
INFRASTRUCTURE TO SUPPORT U.S. SCIENCE IN 2017-2020**

WILLIAM D. GROPP, University of Illinois, Urbana-Champaign, *Co-Chair*

ROBERT HARRISON, Stony Brook University, *Co-Chair*

MARK R. ABBOTT, Oregon State University

DAVID ARNETT, University of Arizona

ROBERT L. GROSSMAN, University of Chicago

PETER M. KOGGE, University of Notre Dame

PADMA RAGHAVAN, Pennsylvania State University

DANIEL A. REED, University of Iowa

VALERIE TAYLOR, Texas A&M University

KATHERINE A. YELICK, University of California, Berkeley

JON EISENBERG, Director, Computer Science and Telecommunications Board, and Study Director

SHENAE BRADLEY, Senior Program Assistant

Statement of Task

“Examine anticipated priorities and associated trade-offs for advanced computing in support of NSF-sponsored science and engineering research.”

Five major areas

1. Contribution of high-end computing to U.S. leadership and competitiveness in S&E
2. Expected future national-scale computing needs for high-end modeling and simulation, and data analysis
3. Complementarities and trade-offs in investments in hardware, software, data, communications
4. Range of operational models and role of NSF
5. Expected technical challenges to delivering capabilities affordably

Observations on task

- Considering priorities/tradeoffs for both compute-intensive and data-intensive facilities
- Considering full range of needs for infrastructure
 - Including high end and tail
 - Including emergence of demand from SBE
- Emphasis on needs of users
 - But also attention to technical challenges and solutions
- No funding or organizational recommendations

Building Advanced Computing to Support Integrated Discovery

- Advanced computing refers to both compute and data-intensive research
- Common element is that they are expensive enough that computing is shared among multiple researchers, institutions, and applications
- Specific comments sought on:
 - *How to create advanced computing infrastructure that enables integrated discovery involving experiments, observations, analysis, theory, and simulation*

Technology Challenges

- General uncertainty in technical direction
- Unfavorable trends in power consumption and inter-chip communications
- Variability in storage hardware performance and failure rates
- New algorithms and software approaches needed
- Consider comp+ data colocation and cloud
- Specific comments sought on:
 - *Technical challenges to building future, more capable advanced computing systems and how NSF might best respond to them*

Responding to Growing Demand

Committee seeks comments on:

- The match between resources and demand for the full spectrum of systems and impacts on the research community if NSF can no longer provide state-of-the-art computing
- The role that private industry and other federal agencies can play in providing advanced computing infrastructure
- The challenges facing researchers in obtaining allocations of computing resources and suggestions for improving the allocation and review processes

Possible NSF Responses

- Whether wider collection and more frequent updating of requirements for advanced computing could be used to inform strategic planning, priority setting, and resource allocation; how these requirements might be used; and how they might best be developed, collected, aggregated, and analyzed
- The tension between the benefits of competition and the need for continuity as well as alternative models that might more clearly delineate the distinction between performance review and accountability and organizational continuity and service capabilities
- How NSF might best coordinate and set overall strategy for advanced computing-related activities and investments as well as the relative merits of both formal, top-down coordination and enhanced, bottom-up process.

General Observations

- Biggest uncertainty is in estimating demand or requirements for advanced computing
- Historically been a mismatch between resources and demand, but difficult to quantify
- Forward looking requirements would be especially useful, but strategic planning hasn't always been quantitative regarding computing needs
- Examples of science drivers particularly useful, especially if matched with a point of contact
- Recruiting challenge of next generation

Input most useful if received by January 2015

To download/read interim report:

http://www.nap.edu/catalog.php?record_id=18972

To follow project activities:

<http://www8.nationalacademies.org/cp/projectview.aspx?key=49628>

To submit comments:

<http://tiny.cc/inputacistudy>

or email study director Jon Eisenberg,
jeisenbe@nas.edu