ARRA Project Briefing
August 11, 2009

Daniel Hitchcock, Barbara Helland, Walt Polansky and Vince Dattoria

Advanced Scientific Computing Research
Office of Science, U.S. Department of Energy

U.S. DEPARTMENT OF ENERGY

American Recovery and Reinvestment Act (Recovery Act)

- Recovery Act ground rules
  - Shovel ready,
  - Enhancing research infrastructure and supporting high-priority R&D, and
  - No out year mortgages
- ASCR's Recovery Act Projects
  - Magellan
  - Leadership Computing Facility Upgrade
  - Advanced Computer Architectures
  - SciDAC-e
  - Advanced Networking Initiative

U.S. DEPARTMENT OF ENERGY

Magellan - $33M

- What it is:
  - Magellan is a research and development effort to establish a nationwide scientific mid-range distributed computing and data analysis testbed. It will have two sites (NERSC and ALCF) with multiple 10's of teraflops and multiple petabytes of storage, as well as appropriate cloud software tuned for moderate concurrency.
- Why ASCR is Funding it:
  - ASCR recently hosted a workshop (www.sc.doe.gov/ascr/ProgramDocuments/ProgDocs.html) to assess the role of mid-range computing in the Office of Science and revealed that this computation continues to play an increasingly important role in enabling the Office of Science. Although it is not part of ASCR's mission, midrange computing, and the associated data management play a vital and growing role in advancing science in disciplines where capacity is as important as capability. Magellan addresses the current challenges to provisioning sufficient resources in this critical area.
- Expected Outcome
  - Understand how clouds contribute to: 1) wall clock turn around time for results (i.e. does the Magellan cloud produce science results as fast or faster than a local clusters); and 2) ease of use (i.e. does the interface allow the job to be directed to different resources or split between resources without requiring additional input from the user). Both will be measured as compared to existing midrange resources in the Office of Science labs. These results will generate data for a cost-benefit analysis of various mid-range computing options for the Office of Science.
  - Promote open interface specifications for clouds.

OLCF Upgrade - $19.9M

- What it is:
  - Replace 37,376 quad-core Barcelona processors with six-core Istanbul processors.
    - 70% performance boost for less than 20% of the original purchase price of the Cray XT5. Increases the clock frequency from 2.3 GHz to 2.6 GHz, resulting in a node increasing from 74 gigaflops to 125 gigaflops, a 70% increase in node peak performance and a 21 to 27% increase in node memory bandwidth.
    - Increase L3 cache from 2MB to 6MB.
    - Increases the overall peak performance of the XT5 portion of Jaguar to 2.06 petaflops and of the combined XT4 and XT5 complex to 2.3 petaflops.
    - Allows the memory controllers to run independently or in "unganged" mode to use memory more effectively, which increases the memory bandwidth while still allowing "chip kill" error correction to take place. This will further increase the effective memory bandwidth by as much as 7%.
    - Use of HyperTransport-3 more than doubles the intranode bandwidth from 16 to 38.4 GB/s.
    - Improved power management features, including enhanced energy efficiency through a new processor sleep mode that allows the L2 data of a core to be accessed without waking the core.
- Why ASCR is Funding it:
  - In 2009, through DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, the Leadership Computing Facility (LCF) at Oak Ridge allocated 470 million CPU-hours to 38 grand challenge projects from a pool of over 200 projects that had requested over 2.03 billion CPU-hours.
- Expected Outcome:
  - Increase performance of many applications by 1.5X per node. For example:
    - Materials/nano science: DCA++ will be 60% faster, LSMS will be 50% faster
    - Turbulence and Combustion: DNS will be 50% faster, S3D will be 40% faster
    - Weather & Climate: WRF will be 50% faster, HOMME (atmosphere) & POP (ocean) will be 50% & 10% faster, respectively
    - Chemistry: MADNESS and NWChem will be at least 30% and 10% faster, respectively
    - Fusion: GTC will be 33% faster and AORSA will be at least 10% faster

Advanced Architectures - $5.2M

- What it is:
  - New effort to provide early access to DOE researchers of technologies emerging from IBM PERCS effort.
  - Enhancement of UCB RAMP effort to provide focused research on flexible simulations of performance of scientific applications on next generation microprocessors.
  - Both proposals were in hand and were peer reviewed.
- Why we are Funding it:
  - To provide ASCR with a multiple-path approach to ensure that DOE's computational needs are supported in the next generation of advanced computer architectures, which can provide the increases in computational capability necessary for addressing critical national problems (e.g., alternative energy strategies, climate change, and national security) will require fundamental breakthroughs.
- Expected Outcome
  - Research results that will help to inform ASCR initiatives in computer architecture in FY 2010 and beyond.

U.S. DEPARTMENT OF ENERGY

SciDAC-e - $30M

- What it is:
  - One time stimulus of applied mathematics & computer science research efforts to establish computational foundation to advance the DOE mission across a wide range, including developing renewable energy sources and developing smart grids.
- Why ASCR is Funding it:
  - The three (3) components of SciDAC-e are wholly within the ASCR mission:
    - applied mathematics research to enable bigger, better and smarter electrical grids;
    - supplemental awards to SciDAC Centers and Institutes to support collaborative research with BES-supported Energy Frontier Research Centers (EFRCs) to develop a high-performance computing capability relevant to the goals of the EFRCs; and
    - enhanced user support at NERSC, ALCF & OLCF for SciDAC-e and energy users awarded allocations through ASCR Leadership Computing Challenge allocation process; SciDAC-e summer school for training and experience with leadership computing resources.
- Expected Outcome:
  - Algorithms to simulate performance of electrical grids over a full range of operating conditions
  - Provide software environment and intellectual resources to EFRCs to meet computational goals

Advanced Networking Initiative - $66.8M

- What it is:
  - A research and development effort to establish a nation-wide demonstration prototype with 100gbps throughput capability.
  - Connecting four geographically dispersed sites including NERSC, ALCF, OLCF, and a peering point in NYC.
  - A distributed testbed and associated research supporting advanced network research topics.
  - All proposals were in hand and were peer reviewed.
- Why we are Funding it:
  - Accelerate the commercial availability of the next generation of high-speed optical backbone.
  - Explore the challenges and solutions for installing and operating a 100gbps optical backbone.
  - Explore the challenges in leveraging the capabilities of a 100gbps infrastructure for end-to-end bulk file transfers at or near available bandwidth for addressing critical national problems (e.g., LHC full-capacity, climate change modeling for IPCC AR5, etc.).
- Expected Outcome
  - Research results that will help to inform ASCR initiatives in networking tools and middleware in FY 2010 and beyond.
  - Demonstrate the capabilities of 100gbps and map a strategy for an ESnet implementation.