U.S. DEPARTMENT OF ENERGY | Office of Science

HEP Requirements Review

Barbara Helland, Facilities Division Director
Advanced Scientific Computing Research

June 10-12, 2015

ASCR Facilities Division

- Providing the Facility – **High-End and Leadership Computing**
  - **National Energy Research Scientific Computing Center (NERSC)** at Lawrence Berkeley National Laboratory
    - Delivers high-end capacity computing to entire DOE SC research community
    - Over 5000 users and 400 projects
  - **Leadership Computing Centers at Argonne National Laboratory (ALCF) and Oak Ridge National Laboratory (OLCF)**
    - Established in response to P.L. 108-423, Department of Energy High-End Revitalization Act of 2004
    - Delivers highest computational capability through **Innovative and Novel Computational Impact on Theory and Computation (INCITE)** program
      - » Open to national and international researchers, including industry
      - » No requirement of DOE or Office of Science funding or topic area
      - » Peer and computational reviews
    - Approximately 300 users and 25-30 projects at each center
- Linking it all together – **Energy Sciences Network (ESnet)**

Previous Requirements Gathering Efforts: "Lead with the Science"

Value of Approach

- **Review meetings establish consensus on requirements, capabilities, services**
- **Scientists, programs offices, and facilities have the same conversation**
- **Provides a solid, fact-based foundation for service and capability investments**
- **Addresses DOE mission goals by ensuring DOE science is effectively supported**

Mission Need for LCF 2017-2018 Upgrades

Science challenges that can be tackled with proposed upgrades:

- ***Energy Storage:*** Develop multiscale, atoms-to-devices, science-based predictive simulations of cell performance characteristics, safety, cost, and lifetime for various energy storage solutions along with design optimizations at all hierarchies of battery (battery materials, cell, pack, etc.).
- ***Nuclear Energy**:* Develop integrated performance and safety codes with improved uncertainty quantification and bridging of time and length scales. Implement next-generation multiphysics, multiscale models. Perform accurate full reactor core calculations with 40,000 fuel pins and 100 axial regions.
- ***Combustion:*** Develop fuel -efficient engines through 3D simulations of high-pressure, low-temperature, turbulent lifted diesel jet flames with biodiesel or rate controlled compression ignition with fuel blending of alternative C1-C2 fuels and n-heptane. Continue to explore the limits of high-pressure, turbulent combustion with increasing Reynolds number.
- ***Fusion:*** Perform integrated first-principles simulation including all the important multiscale physical processes to study fusion-reacting plasmas in realistic magnetic confinement geometries.
- ***Electric Grid:*** Optimize the stabilizing of the energy grid while introducing renewable energy sources; incorporate more realistic decisions based on available energy sources.
- ***Accelerator Design**:* Simulate ultra-high gradient laser wakefield and plasma wakefield accelerator structures.
- ***Catalysis Design**:* Enable end-to-end, system-level descriptions of multifunctional catalysis including uncertainty quantification and data-integration approaches to enable inverse problems for catalytic materials design.
- ***Biomass to Biofuels:*** Simulate the interface and interaction between 100-million-atom microbial systems and cellulosic biomass, understanding the dynamics of enzymatic reactions on biomass. Design of superior enzymes for conversion of biomass.
- ***High resolution climate modeling:*** Simulate high resolution events by incorporating scale aware physics that extends from hydrostatic to nonhydrostatic dynamics. Incorporate cloud resolving simulation codes that couple with a dynamically responding surface.
- ***Rapid climate and earth system change:*** Adequately simulate physical and biogeochemical processes that drive nonlinear responses in the climate system, e.g., rapid increases of carbon transformations and cycling in thawing permafrost; ice sheet grounding line dynamics with ocean coupling that lead to rapid sea level rise; dynamics of teleconnections and system feedbacks within e.g. the (meridional) ocean circulation that alter global temperature and precipitation patterns.

ASCR Computing Upgrades At a Glance

System attributes	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrades	
Name Planned Installation	**Edison**	**TITAN**	**MIRA**	**Cori** **2016**	**Summit** **2017-2018**	**Theta** **2016**	**Aurora** **2018-2019**
System peak (PF)	2.6	27	10	> 30	150	>8.5	180
Peak Power (MW)	2	9	4.8	< 3.7	10	1.7	13
Total system memory	357 TB	710TB	768TB	~1 PB DDR4 + High Bandwidth Memory (HBM) +1.5PB persistent memory	> 1.74 PB DDR4 + HBM + 2.8 PB persistent memory	>480 TB DDR4 + High Bandwidth Memory (HBM)	> 7 PB High Bandwidth On-Package Memory Local Memory and Persistent Memory
Node performance (TF)	0.460	1.452	0.204	> 3	> 40	> 3	> 17 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron Nvidia Kepler	64-bit PowerPC A2	Intel Knights Landing many core CPUs Intel Haswell CPU in data partition	Multiple IBM Power9 CPUs & multiple Nvidia Voltas GPUS	Intel Knights Landing Xeon Phi many core CPUs	Knights Hill Xeon Phi many core CPUs
System size (nodes)	5,600 nodes	18,688 nodes	49,152	9,300 nodes 1,900 nodes in data partition	~3,500 nodes	>2,500 nodes	>50,000 nodes
System Interconnect	Aries	Gemini	5D Torus	Aries	Dual Rail EDR-IB	Aries	2nd Generation Intel Omni-Path Architecture
File System	7.6 PB 168 GB/s, Lustre®	32 PB 1 TB/s, Lustre®	26 PB 300 GB/s GPFS™	28 PB 744 GB/s Lustre®	120 PB 1 TB/s GPFS™	10PB, 210 GB/s Lustre initial	150 PB 1 TB/s Lustre®

Requirements Reviews Need to Meet Multiple Needs

- **Facilities needs**
  - Develop mission need statements for proposed upgrades (stretch your imaginations!!)
  - Identify emerging hardware and software needs of researchers, including experimentalists at SC or other scientific user facilities or experiments
- **Headquarters needs**
  - Articulate the case for future upgrades to SC and DOE management, OMB and Congress
    - What are the potential impacts from the investments in upgrades
    - How broad is the reach – industry, other user facilities, other agencies
  - Identify emerging hardware and software needs for SC, including research
    - What gaps can we fill
  - Develop strategic roadmap for facilities division based on scientific need
    - Who are our customers
    - What niche are facilities filling
    - What gaps should we fill

Objectives of Current "Exascale" Requirements Review (RR)

Goal: Ensure the ability of ASCR facilities to support SC mission science in the exascale regime (2020-2025 timeframe).

HEP: Identify key computational science objectives from High Energy Physics that push exascale and describe the HPC ecosystem –HPC machine and related resources- needed to successfully accomplish your science goals

- Capture the whole picture:
  - Identify continuum of computing needs for the program office from institution clusters to Leadership computing.
    - » *Note: ASCR focus is on HPC and Leadership computing*.
  - Include modeling and simulation, scientific user facilities and large experiments needs, data needs, and near real time needs.
- Information gathered will inform the requirements for ecosystems for planned upgrades in 2020-2023 including the pre-exascale and exascale systems, network needs, data infrastructure, software tools and environments, and user services.

ASCR: Communicate to DOE SC scientists the known/fixed characteristics of upcoming compute system in the 2020-2025 timeframe and ask the computational scientists for feedback on proposed architectures.

Strengthen and inform interactions between HPC facility experts and scientists as well as ASCR and HEP.

Implementation of Exascale Requirements Review (RR)

Series of workshops, one per SC Office (a hybrid between NERSC requirements reviews and Scientific Grand Challenges)

- **Location:** Washington DC area
- **Program Committee:** Representative community leaders from SC domain program office and ASCR facility staff
- **Attendance**: ~50 attendees including DOE program managers, DOE SC community representatives, ASCR supported applied mathematicians and computer scientists and a small number of Postdocs and senior CSGF fellows

Proposed Schedule

June 10-12,2015	HEP
November 3-5 2015	BES
January 2015	FES
April/March 2016	BER
June 2016	NP
September 2016	ASCR

- **Agenda**: Plenary session and themed breakout sessions determined by program committee
- **Pre-meeting homework**: Templates will be developed and provided to chairs and attendees of breakout session for discussing and documenting case studies
- **Output**: Summary workshop report written for each workshop.

Related Activities

- **Call for input from all 17 DOE Laboratories released on May 31, 2015 to identify potential applications that could deliver new science capabilities on exascale systems. Input will be used by ASCR/ASC**
  - to identify additional key scientific areas for exascale discovery, and specific opportunities for new and existing scientific applications.
  - to provide broad input on the kinds of partnerships and investments required to address technical challenges of exascale applications.
  - Short time frame – lab responses due June 15th
- **NIH-NSF-DOE Request for Information to identify scientific research topics that need High Performance Computing (HPC) capabilities that extend 100 times beyond today's performance on scientific applications.**
  - Information will be used to assist agencies to construct a roadmap, build an exascale ecosystem required to support scientific research, and inform the research, engineering and development process. It is likely that a range of advanced capabilities will need to be developed to respond to the varied computing needs across science disciplines.
  - To be released shortly

Thanks to our leaders

Rob Roser and Salman Habib for HEP

Richard Gerber for taking the lead for ASCR's facilities

Carolyn Lauzon and Lali Chatterjee for the programs

And to all of you who have gathered here today to share your ideas, expertise and opinions.