U.S. DEPARTMENT OF ENERGY | Office of Science

ASCR Requirements Reviews and Related Activities

Barbara Helland
Facilities Division Director
SC/Advanced Scientific Computing Research

Talk Outline

- **ASCR Exascale Requirements Reviews**
- **Requests for Information on Exascale Applications**
  - DOE
  - Joint with NIH and NSF
- **HPC Operational Review of Best Practices for Scientific Software Architecture for Portability and Performance**

ASCR's Facilities

Providing the Facilities **– 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 6,000 users and 800 projects

- **Leadership Computing Centers at Argonne National Laboratory (ALCF) and Oak Ridge National Laboratory (OLCF)**
  - Delivers highest computational capability
    - Open to national and international researchers, including industry –
    - Not constrained by existing DOE or Office of Science funding or topic areas
  - Approximately 1,000 users and 50-60 projects at each center, each year

Linking it all together **– Energy Sciences Network (ESnet)**

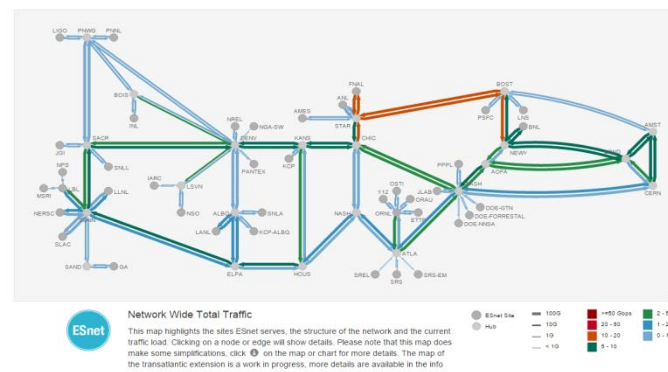

Path to the Future **– Research & Evaluation Prototypes**

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®

Objectives of New "Exascale" Requirements Review (RR)

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

Identify key computational science objectives from DOE SC that push exascale and describe the HPC ecosystem –HPC machine and related resources- needed to successfully accomplish 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.

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

Strengthen and inform interactions between HPC facility experts and scientists as well as ASCR and SC domain office.

DOE Exascale Drivers

Example: Combustion Models Coupling Biofuels with Turbulence

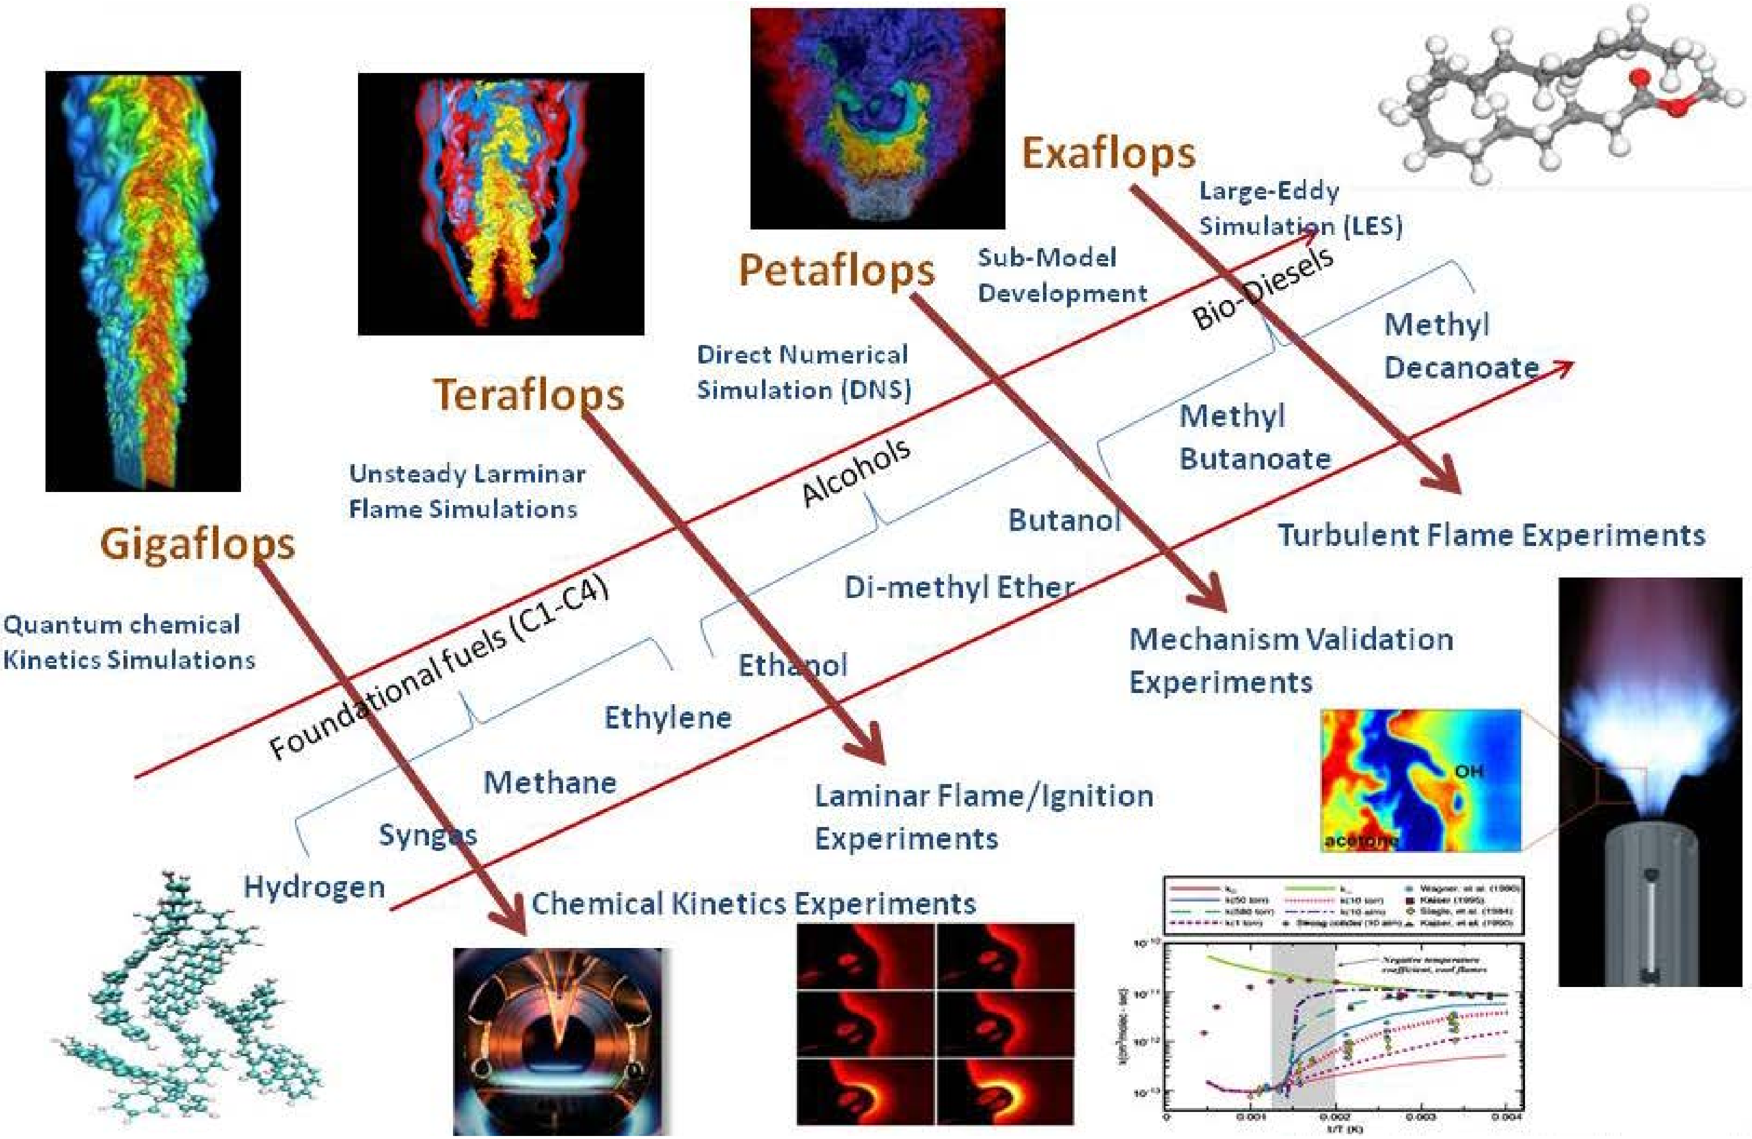

Geologic Carbon Sequestration

Single Injection Site
- Moderate-resolution flow + heat transport
- Low-resolution flow + heat transport + geochemistry + geomechanics

Basin-Scale Injection
- Moderate-resolution flow + heat transport
- Low-resolution flow + heat transport + geochemistry + geomechanics

Basin-Scale Injection
- High-resolution flow + heat transport
- Moderate-resolution flow + heat transport + geochemistry + geomechanics

Multi-Scale Models
- Coupled pore-scale and continuum-scale modeling
- Coupled subsurface and atmospheric modeling

Gigascale — Terascale — Petascale — Exascale

Increasing grid resolution and coupled process complexity

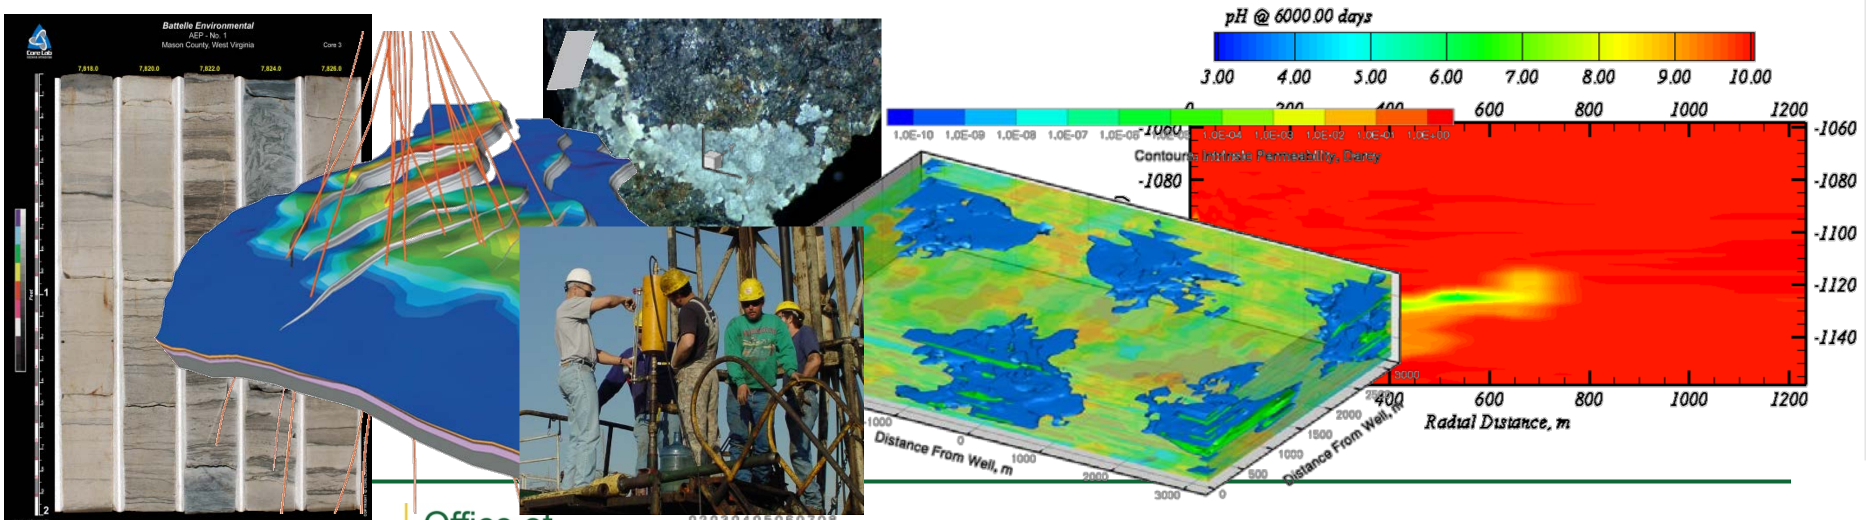

Implementation of Exascale Requirements Review (RR)

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

Proposed Schedule

June 10-12,2015	HEP
November 3-5 2015	BES
January 27-29, 2016	FES
March 29-31, 2016	BER
Target: June 2016	NP
Target: Sept. 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 data
  - White Papers: Broad coverage of Science area
  - Case Studies: Individual Examples
- **Output**: Summary workshop report written for each workshop.

Case Study Template

Case Study Title:

Lead Author(s):

1. **Description of Research**

    1.1 Overview and Context:

    1.2 Research Objectives for the Next Decade

2. **Computational and Data Strategies**

    2.1 Approach

    2.2 Codes and Algorithms

3. **Current and Future HPC Needs**

    3.1 Computational Hours

    3.2 Parallelism

    3.3 Memory

    3.4 Scratch Data and I/O

    3.5 Long-term and Shared Online Data

    3.6 Archival Data Storage

    3.7 Workflows

    3.8 Many-Core and/or GPU Readiness

Answers to the following questions are not required as these issues will be dealt with more globally in the white papers. However, if you have a unique situation that you feel will need to be addressed, please do so.

    3.9 Software Applications, Libraries, and Tools

    3.10 HPC Services

    3.11 Additional Needs

HEP Requirements Review

Stats:

- Held at Bethesda Hyatt on June 10-12, 2015,
- Headquarters POCs: Lali Chatterjee, HEP and Carolyn Lauzon, ASCR
- Conference Co-Chairs : Salman Habib, ANL and Rob Roser, FNAL
- 71 total attendees: 2 co-chairs, 34 domain/application scientists, 35 observers, including staff from ASCR facilities and ASCR Research and Facilities Program Managers
- 9 white papers and 9 case studies in draft report

Breakout Sessions

- Compute-Intensive Modeling and Simulation
- Data-focused analysis and workflows

BES Requirements Review

Stats:

- Held at Rockville Hilton on November 3-5, 2015,
- Headquarters POCs: Jim Davenport, Mark Pederson, Eliane Lessner, BES and Carolyn Lauzon, ASCR
- Conference Co-Chairs Michael Banda, LBNL; Tom Devereaux, Stanford/SLAC; Theresa Windus, ISU/Ames Lab
- 99 total attendees: 3 co-chairs, 60 domain/application scientists, 36 observers, including staff from ASCR facilities and ASCR Research and Facilities Program Managers
- 50 of 60 attendees submitted white papers prior to the review; collecting case studies

Breakout Sessions

- Quantum Materials, Core Challenges in Heavy Element Chemistry, Exotic States, Emergence
- Catalysis, Photosynthesis, Light Harvesting, Combustion
- Materials & Chemical Discovery
- Computing & Data Challenges @ BES Facilities
- Next Generation Programming
- Advances in Quantum Algorithms
- Math & Computer Science
- Soft Matter, Biochemistry, Bioinspired Materials

Talk Outline

- ASCR Exascale Requirements Reviews
- **Requests for Information on Exascale Applications**
  - DOE Labs
  - Joint with NIH and NSF
- HPC Operational Review of Best Practices for Scientific Software Architecture for Portability and Performance

Related Application Needs Gathering 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. They will also
  - 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 --- **133 responses received**
- **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.
  - Released September 15, 2015 and due November 13, 2015 --- **115 responses received**

Response to RFI Spans Sectors

A vision and need for HPC exascale systems was identified across sectors including National Lab, University, Industry, and Non-University Non Profits.

	DOE RFI	Joint RFI	Total
Lab	129	5	134
University	4	93	97
Industry	0	9	9
Non Profit	0	4	4
Government	0	2	2
Unknown	0	2	2

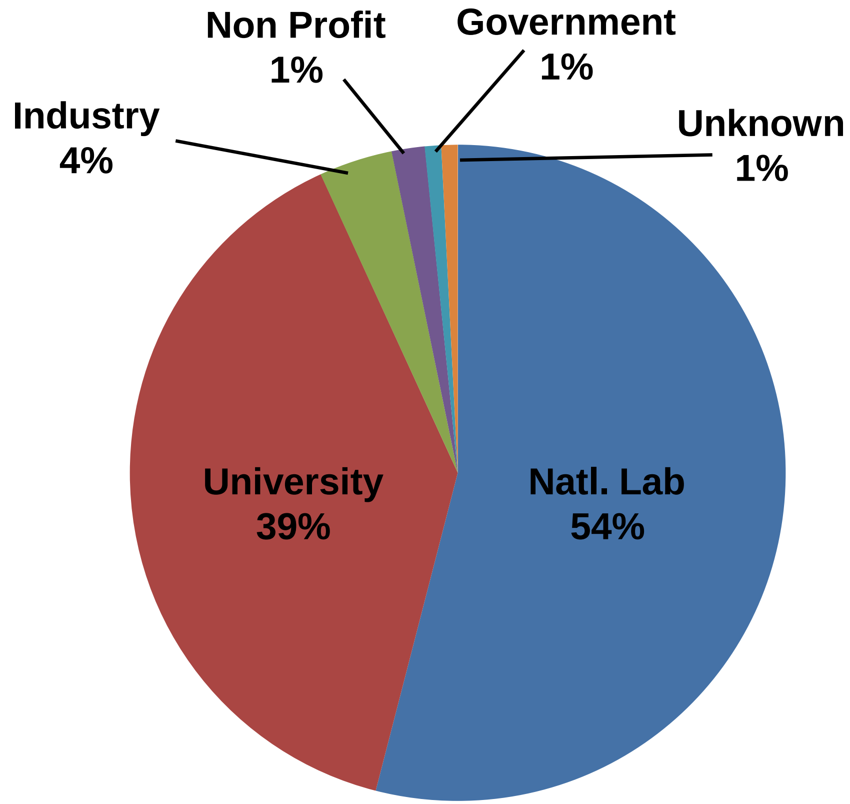

Total Submission

Responses Identify Range of Science Areas

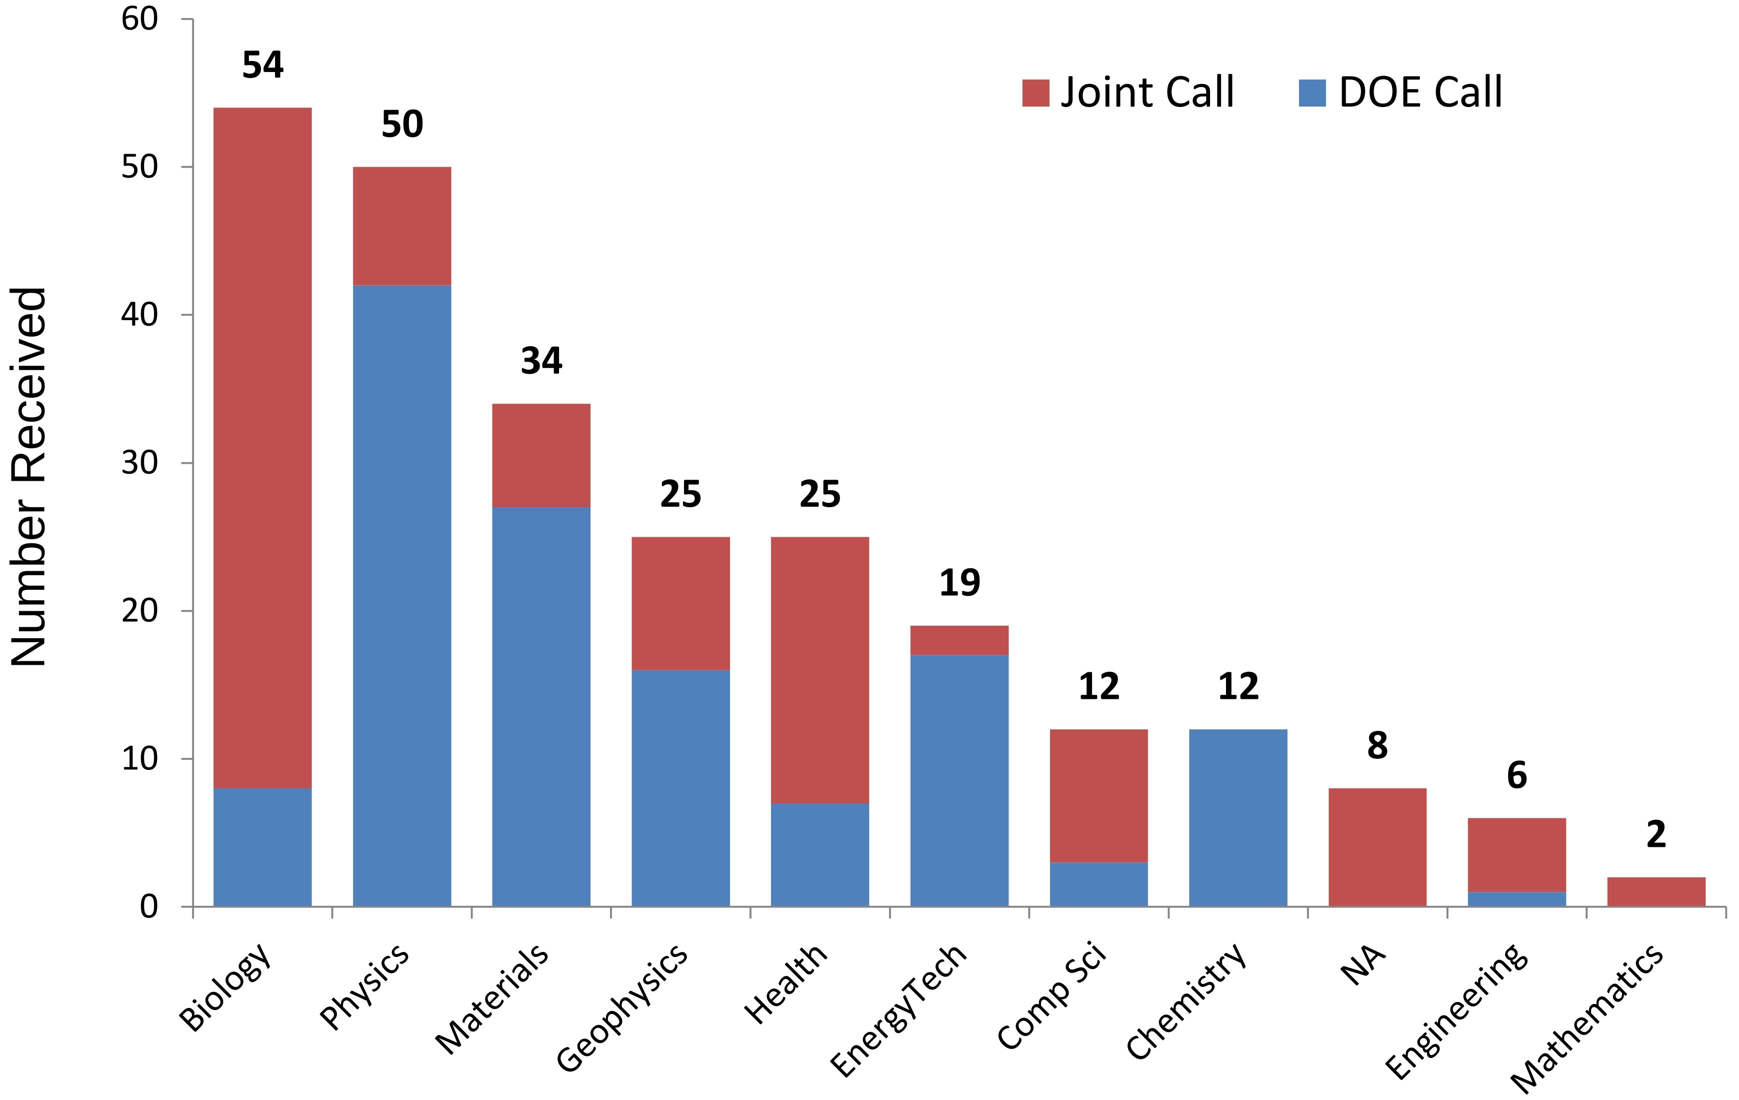

DOE Lab Call for White Papers -- Template

Overview Description and Impact: specific scientific and research challenges that need the projected 100-fold increase in application performance.

System Requirements: Identify the specific limitations of existing HPC systems that must be overcome to perform the planned studies in this area. Describe any specific hardware and software requirements for the system.

Code and Tools: Describe the code and/or tool suite that you have now or will need to address these research objectives.

Models and Algorithms: Describe any new mathematical models and algorithms that will be needed to reach your scientific objectives.

End-to-End Requirements: Describe any capabilities needed by the end-to-end system,

Related Research: Identify any related research areas in your domain of expertise that you foresee that would benefit from this level of augmented computational capability. Identify any barriers in addition to computational capability

10-Year Problem Target: Describe the computational and technical parameters of example problems in this area as you expect them to be in ten years (2025). Indicate key performance parameters specific to this area (e.g. simulated years per day, number of particles, etc.).

Other Considerations/Issues

NSCI Exascale Request for Information (RFI, NOT-GM-15-122)

- The specific scientific and research challenges that would need the projected 100-fold increase in performance
- The potential impact of the research to the scientific community, national economy, and society.
- The specific limitations/barriers of existing HPC systems must overcome
- Any related research areas you foresee that would benefit from this level of augmented computational capability.
- Important computational and technical parameters of the problem as you expect them to be in 10 years (2025)
- Alternative models of deployment and resource accessibility arising out of exascale computing
- Capabilities needed by the end-to-end system, including data requirements such as data analytics and visualization tools, shared data capabilities, and data services which includes databases, portals and data transfer tools/nodes.
- Foundational issues that need to be addressed such as training, workforce development or collaborative environments.
- Other areas of relevance for the Agencies to consider.

NIGMS_exascale@nigms.nih.gov

Common Taxonomy to Understand Science Drivers

Goal - Create common science taxonomy to understand science drivers and portfolio across ASCR facilities

- Works across science disciplines and also highlights DOE science areas
- Science categories provide weight to high use application areas
- Inclusive of possible future application

Process for Developing Taxonomy

- First draft authored by OLCF, ALCF, NERSC, and HQ
  - Started with existing taxonomies used by centers and HQ
  - Integrated NSF organization chart and NIH mission
  - Collected input from DOE PMs to help fill knowledge gap of the team
- Draft shared with NIH and NSF, some updates will likely occur but no large updates have been proposed
- Test new taxonomy on Joint and DOE RFI

Top Level Categories Provide View Across Basic and Applied Sciences

1. Mathematics
2. Computer and Information Science
3. Physics
4. Chemical Sciences
5. Biological Sciences
6. Geosciences
7. Materials Sciences
8. Engineering
9. Applied Energy Technologies
10. Social, Behavioral, and Economic Sciences
11. Health Sciences

Basic Science Categories (items 1–7)

Applied Science Categories highlighting areas important to DOE, NSF, and NIH (items 8–11)

Sub-categories Provide Informative Topic Areas

Examples:

7. Materials Science and Engineering
   - 7.1 Condensed-matter physics
   - 7.2 Electronic Properties
   - 7.3 Mechanics (structural, failure, nucleation)
   - 7.4 Functional Materials
   - 7.5 Soft materials/polymers

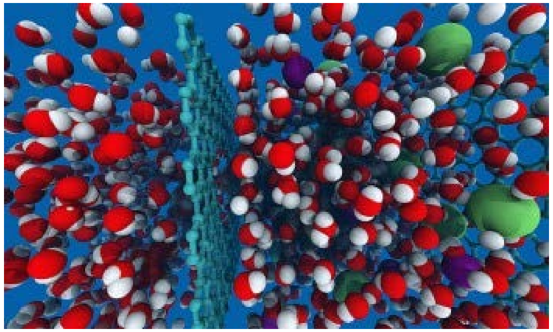

9. Applied Energy Technologies
   - 9.1 Nuclear Energy
   - 9.2 Fossil Energy (fracking, CO2 sequestration, oil and gas exploration, clean coal)
   - 9.3 Renewable Energy (wind, solar, bio, batteries)
   - 9.4 Energy Efficiency (vehicles, buildings)
   - 9.5 Electricity Delivery and Energy Reliability

11. Health Sciences
   - 11.1 Precision Medicine
   - 11.2 Epidemiology
   - 11.3 Population Health
   - 11.4 Imaging

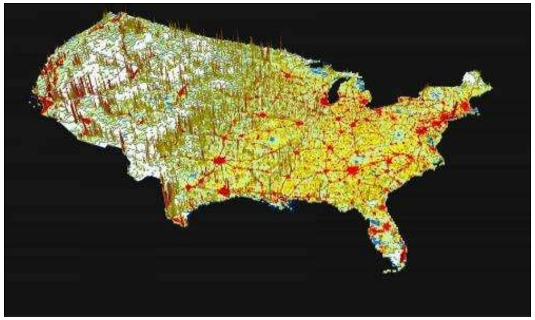

Responses Provide Compelling Science Drivers for Exascale

Climate Challenge Problems: 1km resolution in global simulations, Fully coupled atmosphere-ocean-land-ice models

Impacts:

- Forecast water resources with increased confidence
- Project future changes in severe weather,
- Address food supply change with defensible estimations,
- Inform decision on the resilience of energy and public health infrastructure

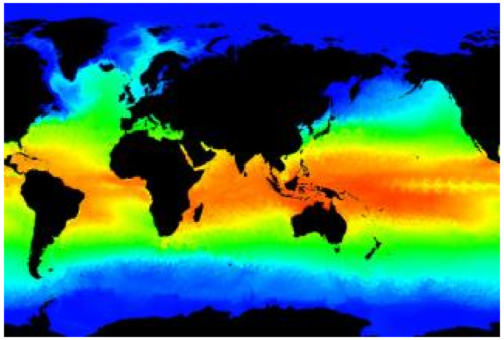

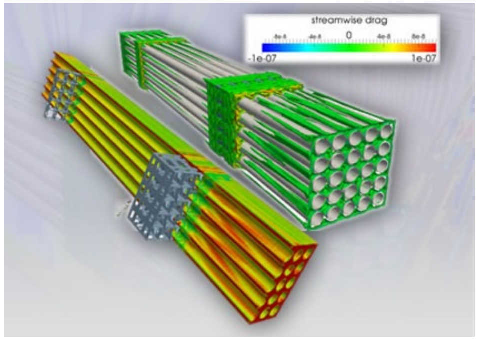

Nuclear Energy Challenge Problems: Simulations over full 18-month fuel cycle, LES of reactor core with 10B elements, Explicitly resolve 51K fuel rods in 3D

Impacts:

- Predictive capabilities for untested new designs
- Inform engineering design in operational nuclear reactors
- Improve radiation therapy treatments for clinical use

Health Science Challenge Problems: Integrated imaging , genomic and medical data; High-resolution large-scale multi-factor simulations of infectious disease

Impacts:

- Transform patient care with early and accurate diagnosis
- Inform decisions on public health policy and response
- Economic benefit of cost-effective public and personal health

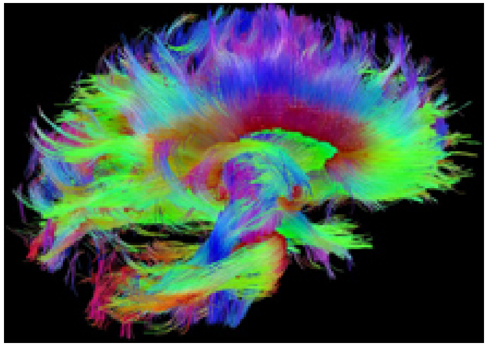

Responses Identified Common Barriers and Challenges

Barriers to Exascale

- **Applications**: scaling, using accelerators, lack of HPC ready codes, lack of software containers
- **Workforce development-** scientific professionals who can use highly parallel systems
- **Data workflows -** memory and I/O footprint on future architectures, data protection, data access from multiple sources, data sharing
- **Ease of resource access**

Exascale Data Challenges

- **52 of the 248** responses emphasized data as important component
- **56% of the 52 in Biology and Health**: mostly "Omics" and large scale clinical data analysis

 # Responses Emphasizing Data

Biology (18 of 54)	Geophysics (3 of 25)
Health Science (11 of 25)	NA (3 of 8)
CIS (8 of 12)	Mathematics (1 of 2)
Physics (3 of 50)	Social Sciences (1 of 1)

Software Technologies Cited in the DOE Lab RFI Responses

- **Math Libraries, Components, and Functions**
  - Forward/Inverse Fast Fourier Transforms (FFTs), Fast Multipole Method (FMM)
  - Numerical Linear Algebra: BLAS/cuBLAS, LAPACK/ScaLAPACK, MAGMA, ATLAS
  - Sparse Linear and Nonlinear Solvers: PETSc, Trilinos, HYPRE
  - Eigensolvers: SLEPc
  - Others: ensemble Kalman Filter (EnKF), SUNDIALS
  - Hardware/Software Specific: Intel MKL, NumPy
- **Compilers and Languages**
  - C, C++, Fortran (77/90/95/03), LLVM, Clang
  - Python, Swift, Lua, Matlab, R, Basic, Java, Scala, ArcGIS, Kokkos
- **Accelerator Programming**
  - CUDA, OpenACC, OpenCL
- **Shared and Distributed Memory Parallelism**
  - MPI, Pthreads, OpenMP, PGAS (Global Arrays)
- **Development Environments and Runtime Systems**
  - Legion, HPX, Charm++, Uintah, Meld
  - HADOOP

Software Technologies Cited in the DOE Lab RFI Responses

- **Data Science, Data Management, and Visualization**
  - Input/Output: HDF, netCDF, ADIOS
  - Meshing: SCOREC
  - Visualization: ParaView, VisIt, IDL, VTK
  - Data Assimilation/Calibration/Sensitivity, UQ, Optimization: OpenDA, DAKOTA
  - Environments and Frameworks: Spark
  - Geographic Information: ArcGIS, VGM
- **Domain Specific Frameworks, Environments, Runtime Systems, Libraries, Workflows**
  - Astrophysics: Bellerophon
  - Cosmology: CosmoTools, PDACS
  - BioInformatics/Precision Medicine: BioSAL, Thorium, AutoDock, OpenMM
  - Particle Physics: QUDA, QPhiX, QOPQDP
  - Neutron and X-Ray Scattering: Sassena
  - Adaptive Mesh Refinement (AMR): BoxLib, Chombo
  - Discrete Events: NS3, OMNeT++
  - Plasmas: BOUT++

Talk Outline

- ASCR Exascale Requirements Reviews
- Requests for Information on Exascale Applications
  - DOE
  - Joint with NIH and NSF
- **HPC Operational Review of Best Practices for Scientific Software Architecture for Portability and Performance**

HPCOR: Scientific Software Architecture For Portability And Performance

- **80 participants + observers**
- **Sept 15th- 17th, 2015 Gaithersburg Marriott**
- **Three breakouts**
  - Application Architecture
  - Libraries and Tools
  - Software Engineering
- **Practices, Emerging Ideas, Failures, Opportunities**

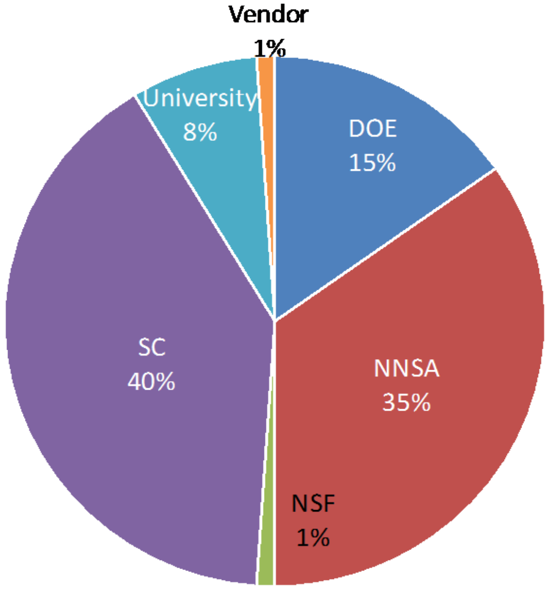

Mission: Identify best practices for scientific software to increase portability and performance

Workshop Findings

Primary Findings

- True performance portability is very hard, and the community still needs to determine what is possible
- Software engineering for applications and libraries is key to taking the application challenges of the future
- Communication paths are needed to share the huge about of work being done today – workshops, forums, etc.

Specific Comments

Application Architecture

- Portable functionality is easier than portable performance
- Best practice today is separate node and core-level parallelism
- Separate concerns, including data structures
- Libraries prohibitive for early functionality and performance in early science periods

Libraries and Tools

- Users and developers need confidence in funding and persistence
- Transparent roadmaps
- Homogenous software stack helps adoption
- Better communication with science/application developers needed

Software Engineering

- Scientific software needs test driven development
- Good practices impact productivity, portability, performance, and people retention
- Not enough training or incentives for scientist to adopt practices