U.S. DEPARTMENT OF ENERGY

DOE Office of Advanced Scientific Computing Research

Presented to the
Advanced Scientific Computing Advisory Committee

by

Steve Binkley
Associate Director

July 27, 2015

Advanced Scientific Computing Research

Computational and networking capabilities to extend the frontiers of science and technology

- **Mathematics research** to address challenges of increasing complexity within DOE's mission areas from a mathematical perspective. This requires integrated, iterative processes across multiple mathematical disciplines.
- **Computer science research** to increase the productivity and integrity of HPC systems and simulations, and support data management, analysis, and visualization techniques.
- **SciDAC partnerships** to dramatically accelerate progress in scientific computing that delivers breakthrough scientific results.
- **Exascale computing** research and development of capable exascale hardware architectures and system software, including the deployment of programming environments for energy-efficient, data-intensive applications, and engagement with HPC vendors to deliver systems that address the exascale challenges.
- **Facilities** operate with at least 90% availability while continuing planned upgrades – begin deployment of 10-40 petaflop upgrade at NERSC and continue preparations for 75-200 petaflop upgrades at each LCF.
- Continue a postdoctoral program at the ASCR facilities and provide funding for the Computational Science Graduate Fellowship to address DOE workforce needs.

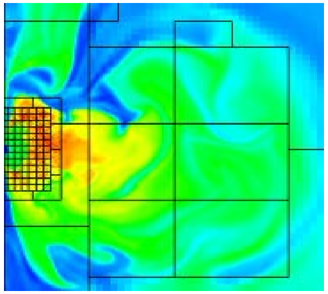
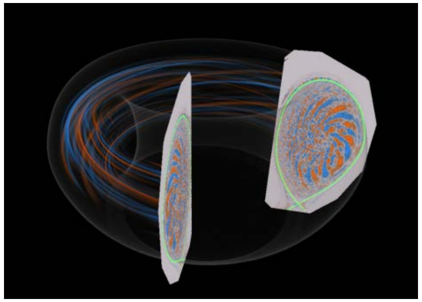
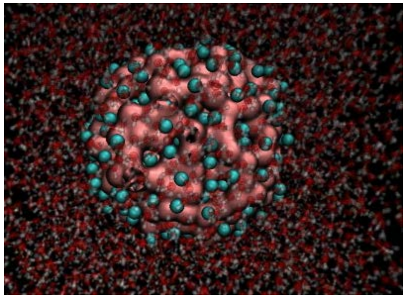
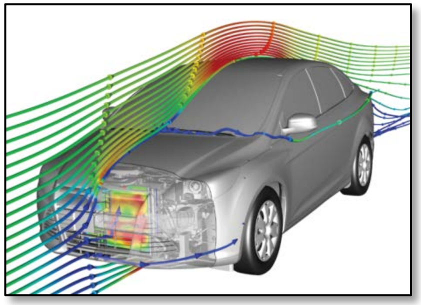

ASCR Investment Priorities

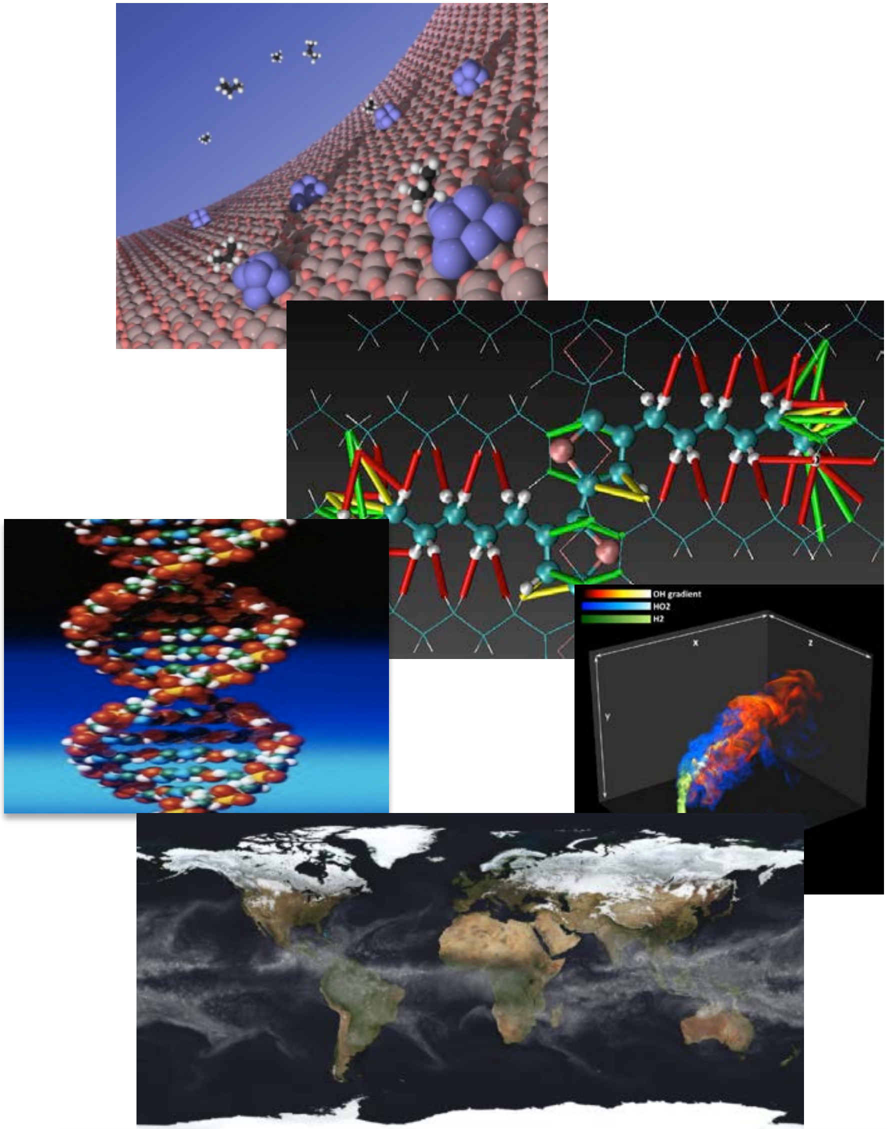

Investment Priorities:

- **Exascale –** conduct research and development, and design efforts in hardware software, and mathematical technologies that will produce exascale systems for science applications.
- **Facilities –** acquire and operate more capable computing systems, from multi-petaflop through exascale computing systems that incorporate technologies emerging from research investments.
- **Large Scientific Data –** prepare today's scientific and data-intensive computing applications to migrate to and take full advantage of emerging technologies from research, development and design efforts.

Office of Science FY 2016 Budget Status

Office of Science

FY 2016 Congressional Action

(B/A in thousands)

	FY 2015			FY 2016						
	President's Request	Enacted Approp.	Current Approp.	President's Request	House Mark	FY16 House Mark vs. FY15 Enacted		Senate Mark	FY16 Senate Mark vs. FY15 Enacted	
SCIENCE										
ASCR……………………	541,000	541,000	523,411	620,994	537,539	-3,461	-0.6%	620,994	+79,994	+14.8%
BES………………………	1,806,500	1,733,200	1,682,924	1,849,300	1,770,306	+37,106	+2.1%	1,844,300	+111,100	+6.4%
BER………………………	628,000	592,000	572,618	612,400	538,000	-54,000	-9.1%	610,000	+18,000	+3.0%
FES………………………	416,000	467,500	457,366	420,000	467,600	+100	+0.0%	270,168	-197,332	-42.2%
HEP………………………	744,000	766,000	745,232	788,000	776,000	+10,000	+1.3%	788,100	+22,100	+2.9%
NP…………………………	593,573	595,500	580,744	624,600	616,165	+20,665	+3.5%	591,500	-4,000	-0.7%
WDTS……………………	19,500	19,500	19,500	20,500	20,500	+1,000	+5.1%	19,500	……	……
SLI…………………………	79,189	79,600	79,600	113,600	89,890	+10,290	+12.9%	113,600	+34,000	+42.7%
S&S………………………	94,000	93,000	93,000	103,000	103,000	+10,000	+10.8%	100,715	+7,715	+8.3%
PD…………………………	189,393	183,700	183,700	187,400	181,000	-2,700	-1.5%	185,000	+1,300	+0.7%
SBIR/STTR (SC)…………	……	……	132,905	……	……	……	……	……	……	……
Subtotal, Science………	5,111,155	5,071,000	5,071,000	5,339,794	5,100,000	+29,000	+0.6%	5,143,877	+72,877	+1.4%
SBIR/STTR (DOE)………	……	……	65,075	……	……	……	……	……	……	……
Subtotal, Science………	5,111,155	5,071,000	5,136,075	5,339,794	5,100,000	+29,000	+0.6%	5,143,877	+72,877	+1.4%
Rescission of PY Bal……	……	-3,262	……	……	-4,717	-1,455	+44.6%	-4,717	-1,455	+44.6%
Total, Science Approp.…	5,111,155	5,067,738	5,136,075	5,339,794	5,095,283	+27,545	+0.5%	5,139,160	+71,422	+1.4%
(SBIR/STTR total)………	(……)	(……)	(197,980)	(……)	(……)	(……)	(……)	(……)	(……)	(……)

ASCR Budget Overview

		FY 2014 Current Appropriation (w/o SBIR/STTR)	FY 2015 Enacted Approp.	FY 2016 President's Request	*FY16 vs. FY15*
	Advanced Scientific Computing Research				
Exascale	Applied Mathematics	47,081	49,155	49,229	+74
Exascale	Computer Science	55,835	55,767	56,842	+1,075
Exascale	Computational Partnerships (SciDAC)	46,261	46,918	47,918	+1,000
	Next Generation Networking for Science	17,852	19,000	19,000	+0
	SBIR/STTR	0	5,830	6,181	+351
	Total, Mathematical, Computational, and Computer Sciences Research	*167,029*	*176,670*	*179,170*	*+2,500*
	High Performance Production Computing (NERSC)	67,105	75,605	76,000	+395
	Leadership Computing Facilities	160,000	184,637	171,000	-13,637
Exascale	Research and Evaluation Prototypes	36,284	57,329	141,788	+84,459
	High Performance Network Facilities and Testbeds (ESnet)	33,054	35,000	38,000	+3,000
	SBIR/STTR	0	11,759	15,036	+3,277
	Total, High Performance Computing and Network Facilities	*296,443*	*364,330*	*441,824*	*+77,494*
	Total, Advanced Scientific Computing Research	**463,472**	**541,000**	**620,994**	**+79,994**

Updates

- Appointees
- CSGF – Annual meeting is July 28-30
- ASCR Personnel Updates & Awards
- Early Career Research Program
- Exascale update
- Requirements Gathering Process

Computational Science Graduate Fellowship (CSGF)

Partnerships to Deliver Future Leaders

- **Started in 1991 to broadly train advanced computational scientists**
- **Funded by both DOE-SC/ASCR and NNSA/ASC**
- **Requires that fellows**
  - plan and follow a plan of study that transcends the bounds of traditional academic disciplines
  - participate in 12-week research experience at DOE lab
- **Benefits**
  - Up to four years of support, including full tuition and required fees paid
  - Yearly stipend of $36,000
  - Academic allowance
- **CSGF Longitudinal Study (2012)**
  - "…The generation of a DOE CSGF community of scholars and the building of collaborative networks – notably, often across generational and disciplinary lines – have been critical outcomes of the Fellowship experience." (page 18)
  - "…By linking individual elements with institutional and external realities and needs, the DOE CSGF program itself has operated to identify and involve individuals who might serve not only the field and their own professional goals, but also the national agenda and society more generally, both directly and indirectly." (page 75)

Lawrence Livermore National Laboratory

LOCKHEED MARTIN

Duke University

MIT Massachusetts Institute of Technology

Berkeley Lab

Procter&Gamble

IBM

The Subcommittee believes that the CSGF is unique in its focus on Computational Science. It provides features that other Graduate research Fellowships do not, such as the Plan of Study, the Practicum, the Annual CSGF Conference and efforts to keep alumni engaged. In this regard, the CSGF is an exceptional program that produces interdisciplinary scientists uniquely qualified to address current and future computational science challenges.

2011 ASCAC Review of DOE Computational Science Graduate Fellowship Program

Computational Mathematician at Sandia receives DOE's Lawrence Award

Bochev wins E.O. Lawrence Award – July 23, 2015

- **Pavel Bochev, a computational mathematician at Sandia National Laboratories, has received an Ernest Orlando Lawrence Award for his pioneering theoretical and practical advances in numerical methods for partial differential equations.**
- **Citation in Category of Computer, Information, and Knowledge Sciences: "… invention of new algorithms, the rigorous mathematical analysis of algorithms and the mathematical models to which they apply, the implementation of computational algorithms for solving application problems … to DOE".**

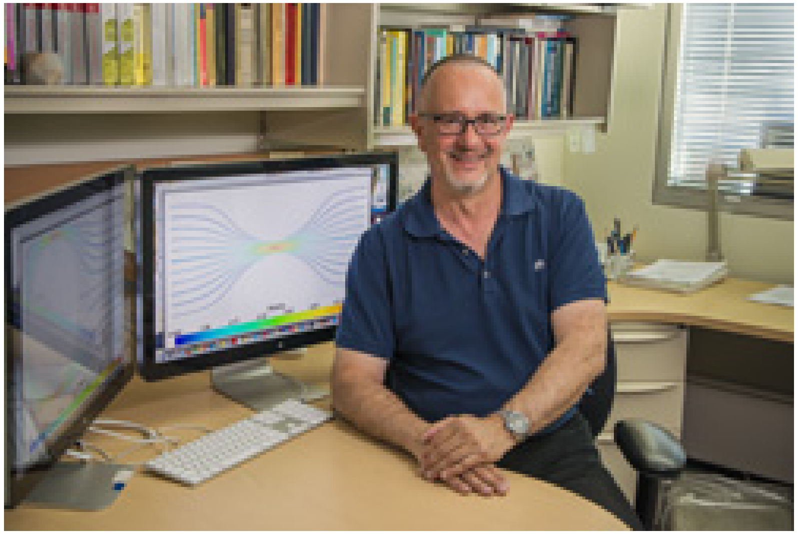

Pavel Bochev

The Lawrence Award honors the memory of E. O. Lawrence, who invented the cyclotron and received a 1939 Nobel Prize in physics for that achievement. Lawrence later played a leading role in establishing the U.S. system of national laboratories

DOE SC Early Career Research Program

- **University Awards**
  - **Minimum award size is $750,00 for 5 years; $150,000/year**
    - PI must be an untenured Assistant Professor on the tenure track or an untenured Associate Professor on the tenure track at a U.S. academic institution as of the deadline for the application
- **National Laboratory Awards**
  - **Minimum award size is $2,500,000 for 5 years; $500,000/year**
    - PI must be a full-time, permanent, non-postdoctoral national laboratory employee as of the deadline for the proposal
- **Eligibility**
  - No more than ten (10) years can have passed between the year the Principal Investigator's Ph.D. was awarded and the year of the deadline for the proposal
  - There can be no co-Principal Investigators

2015 ASCR Early Career Research Program Awards

Area	Description	
Applied Math	***Cory Hauck – Oak Ridge National Laboratory*** 2006 PhD in Applied Mathematics, University of Maryland Title: Hybrid Methods for Complex Particle Systems	
Applied Math	***Jonathan Weare – University of Chicago*** 2007 PhD in Mathematics, University of California, Berkeley Title: Ensemble Simulation Techniques and Fast Randomized Algorithms	
Computer Science	***Henry Hoffmann – University of Chicago*** 2013 PhD in Electrical Eng and Comp Sci, Massachusetts Inst of Technology Title: CALORIE: A Constraint Language and Optimizing Runtime for Exascale Power Management	
Computer Science	***Christian Engelmann – Oak Ridge National Laboratory*** 2008 Ph.D. in Computer Science, University of Reading, UK Title: Resilience Design Patterns: A Structured Approach to Resilience	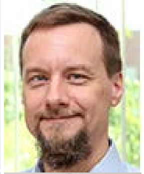
Computer Science	***Daniela Ushizima – Lawrence Berkeley National Laboratory*** PhD: 2004 Ph.D. in Computational Physics, University of Sao Paulo, Brazil Title: Scaling Analytics for Image-Based Experimental Data	

U.S. DEPARTMENT OF ENERGY

FY12-FY14 Early Career Awards

PI	Institution	Awarded	Program
Miranda Holmes-Cerfon	New York University	FY14	Applied Math
Emil Constantinescu	Argonne National Lab	FY14	Applied Math
Guglielmo Scovazzi	Duke University	FY14	Applied Math
Todd Gamblin	Lawrence Livermore National Lab	FY14	Computer Science
Florin Rusu	University of California at Merced	FY14	Computer Science
Milind Kulkarni	Purdue University	FY13	Computer Science
Sriram Krishnamoorthy	Pacific Northwest National Lab	FY13	Computer Science
Krzysztof Fidkowski	University of Michigan	FY13	Applied Math
Aydin Buluc	Lawrence Berkeley National Lab	FY13	Applied Math
Frank Curtis	Lehigh University	FY13	Applied Math
Pavan Balaji	Argonne National Lab	FY12	Computer Science
Hank Childs	Lawrence Berkeley National Lab	FY12	Computer Science
Victor Zavala	Argonne National Lab	FY12	Applied Math
Aleksandar Donev	New York University	FY12	Applied Math
Mattan Erez	University of Texas at Austin	FY12	Computer Science
Haim Waisman	Columbia University	FY12	Applied Math

FY10-FY11 Early Career Awards

PI	Institution	Awarded	Program
Kalyanaraman Anantharaman	Washington State University	FY11	Computer Science / Computational Biology
Peter Lindstrom	Lawrence Livermore National Lab	FY11	Computer Science
Alexandre Tartakovsky	Pacific Northwest National Lab	FY11	Applied Mathematics
Alireza Doostan	University of Colorado	FY11	Applied Mathematics
Xipeng Shen	College of William & Mary	FY11	Computer Science
Christiane Jablonowski	University of Michigan	FY10	Computational Applications/ Climate
Youssef Marzouk	Massachusetts Institute of Technology	FY10	Applied Mathematics
Patrick Chiang	Oregon State University	FY10	Computer Science
Michelle Strout	Colorado State University	FY10	Computer Science
Grigory Bronevetsky	Lawrence Livermore National Lab	FY10	Computer Science
Kalyan Perumalla	Oak Ridge National Lab	FY10	Applied Mathematics

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®

U.S. DEPARTMENT OF ENERGY

ESnet Goes Global: Extension to Europe

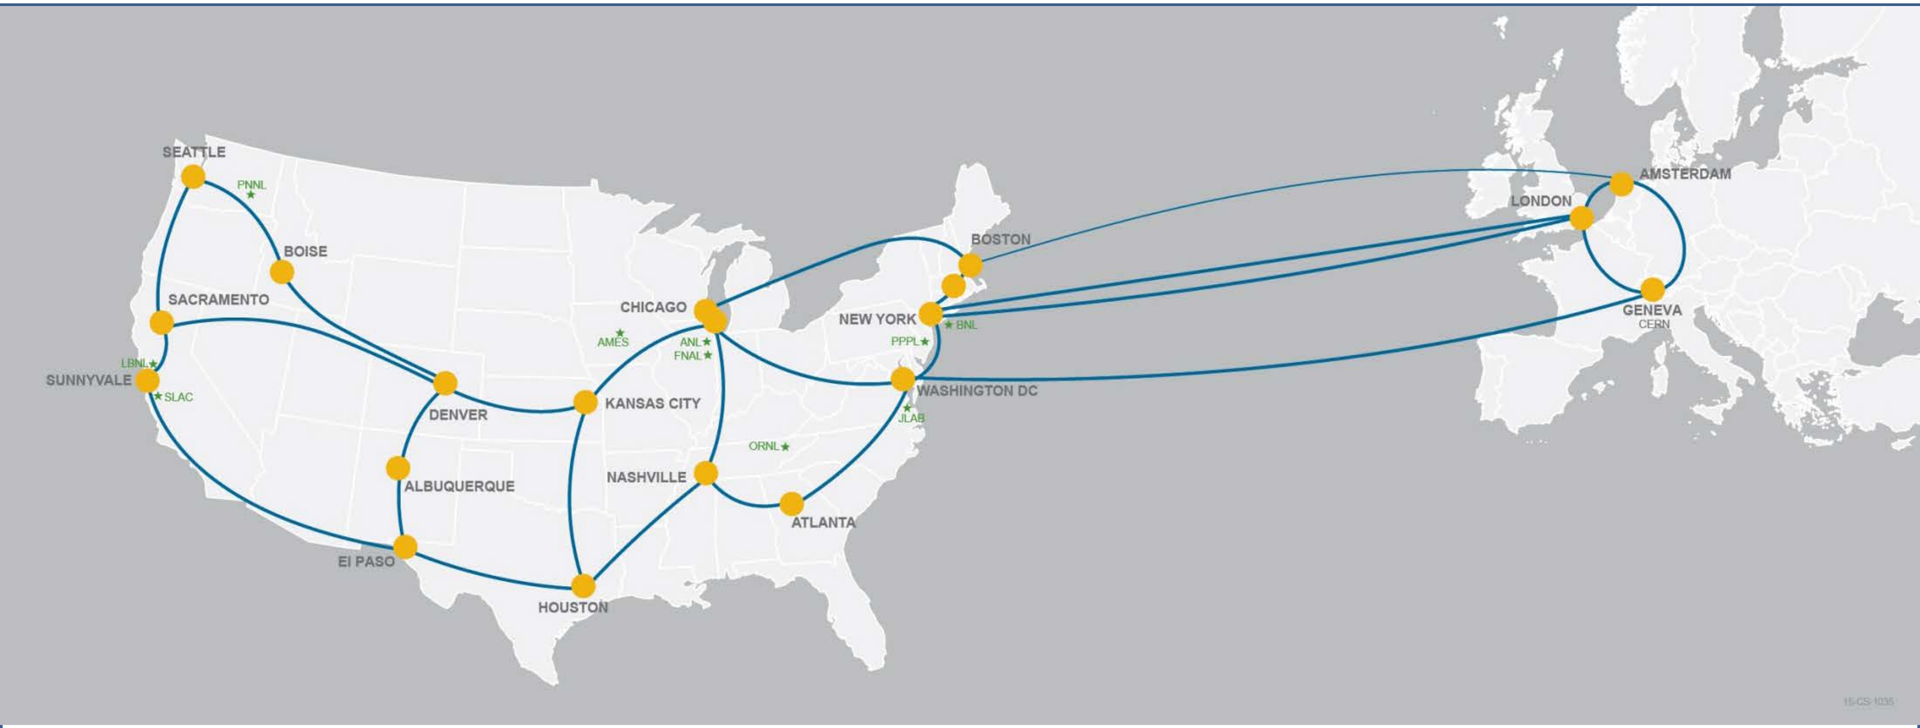

- **LHC experiments (Run 2) resumed June 3, 2015**
- **ESnet 10x increase in transatlantic traffic from Large Hadron Collider**

Exascale Computing Initiative: Next Generation of Scientific Innovation

- **Departmental Crosscut – In partnership with NNSA**
- **"All-in" approach: hardware, software, applications, large data, underpinning applied math and computer science**
- **DOE's missions push the frontiers of science and technology:**
  - Discovery science – next-generation materials
  - Mission-focused basic science in energy – next-generation climate software
  - Use current Leadership Computing approach for users
- **The next generation of advancements will require Extreme Scale Computing**
  - 100-1,000X capabilities of today's computers with a similar physical size and power footprint
  - Significant challenges are power consumption, high parallelism, reliability
- **Extreme Scale Computing, cannot be achieved by a "business-as-usual," evolutionary approach**
  - Initiate partnerships with U.S. computer vendors to perform the required engineering, research and development for system architectures for capable exascale computing
  - Exascale systems will be based on marketable technology – Not a "one off" system
  - Productive system – Usable by scientists and engineers

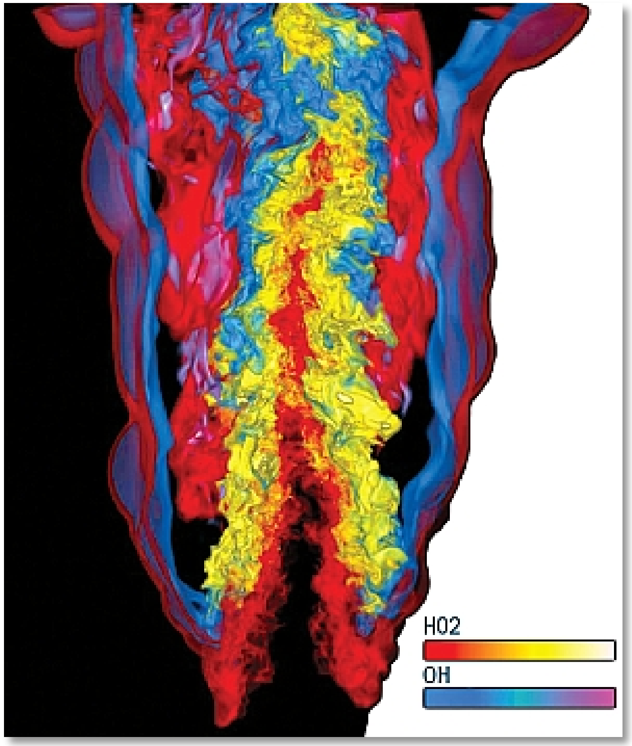

"Projectizing" the Exascale Initiative

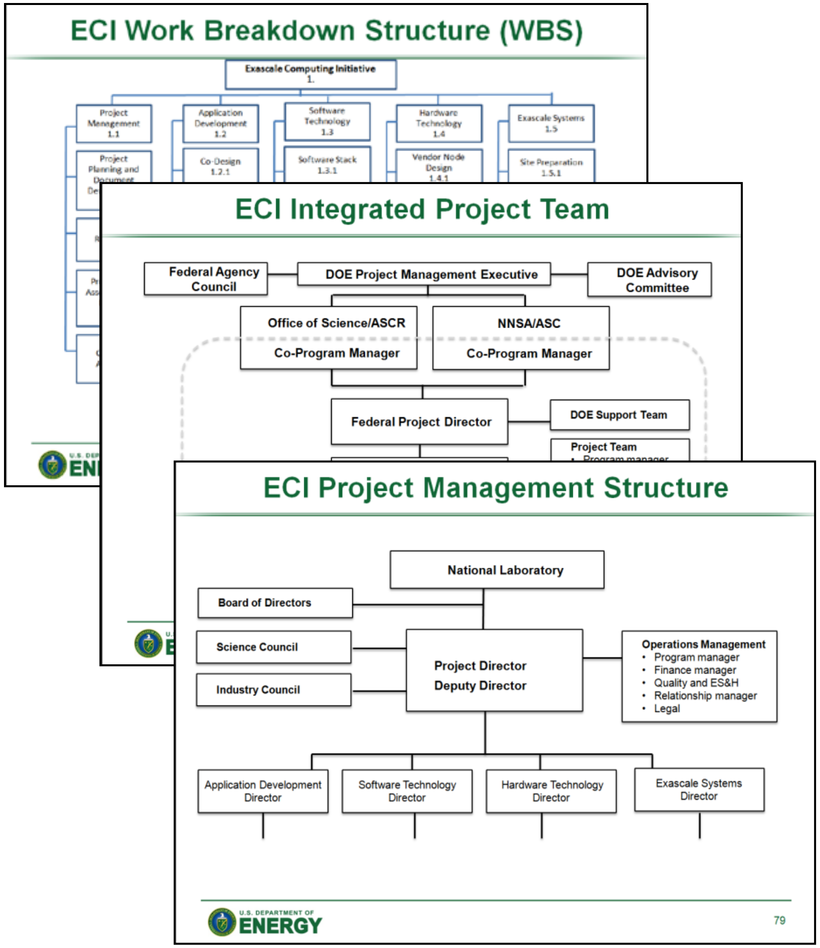

- The exascale initiative will follow established DOE review and decision protocols for its execution
- A project office has been established at ORNL with representation from the major participating laboratories (ANL, LANL, LBNL, LLNL, ORNL, SNL)
- An Integrated Project Team (IPT) has been established, analogous to execution of previous, large, Office of Science projects
- The IPT is refining the work breakdown structure (WBS) and is preparing required project documentation (e.g., critical-decision packages, preliminary project execution plans, etc.)
- A top-level WBS activity has been established to develop and implement exascale applications, based on a labs-wide request for information

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**

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 2016	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.

Requirements Reviews 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

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.

Objectives of “Exascale” Requirements Reviews

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.

ASCR at a Glance

Office of Advanced Scientific Computing Research

Associate Director – Steve Binkley
Phone: 301-903-7486
E-mail: Steve.Binkley@science.doe.gov

Research

Division Director – William Harrod
Phone: 301-903-5800
E-mail: William.Harrod@science.doe.gov

Facilities

Division Director – Barbara Helland
Phone: 301-903-9958
E-mail: Barbara.Helland@science.doe.gov

Relevant Websites

ASCR: science.energy.gov/ascr/

ASCR Workshops and Conferences:
science.energy.gov/ascr/news-and-resources/workshops-and-conferences/

SciDAC: www.scidac.gov

INCITE: science.energy.gov/ascr/facilities/incite/

END

BACKUPS

2015 Proposal Submission Demographics

- **61 Total submitted proposals**
- **15 Laboratory proposals**
- **46 University proposals**
- **29 States represented**
- **25% Nat Lab; 75% University**
- **13% Female; 87% Males**

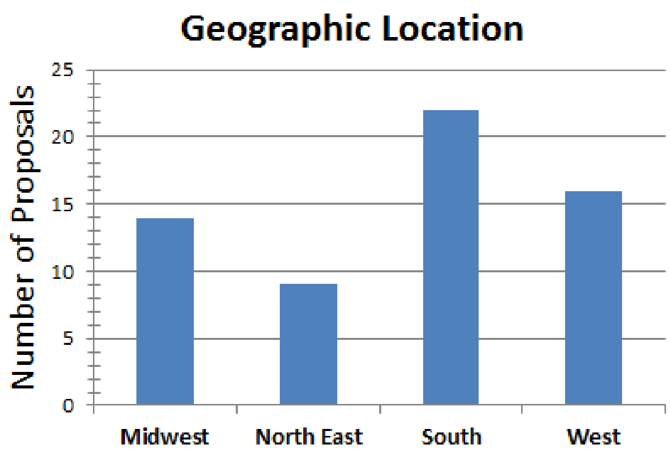

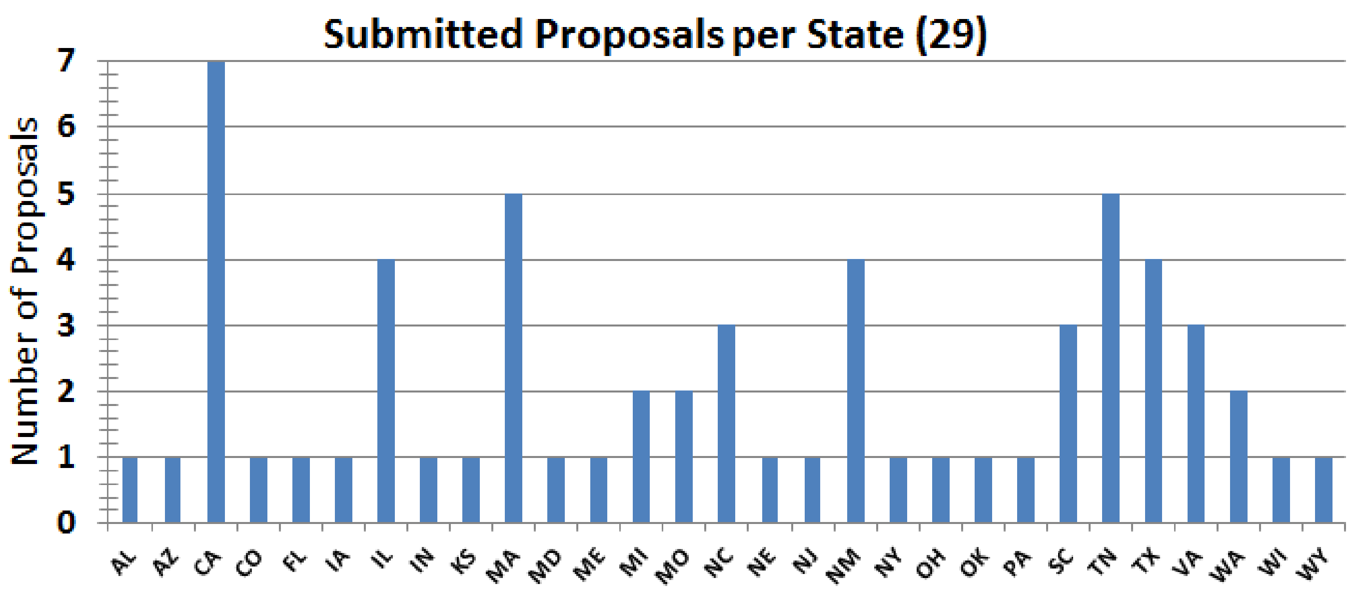

ECI Work Breakdown Structure (WBS)

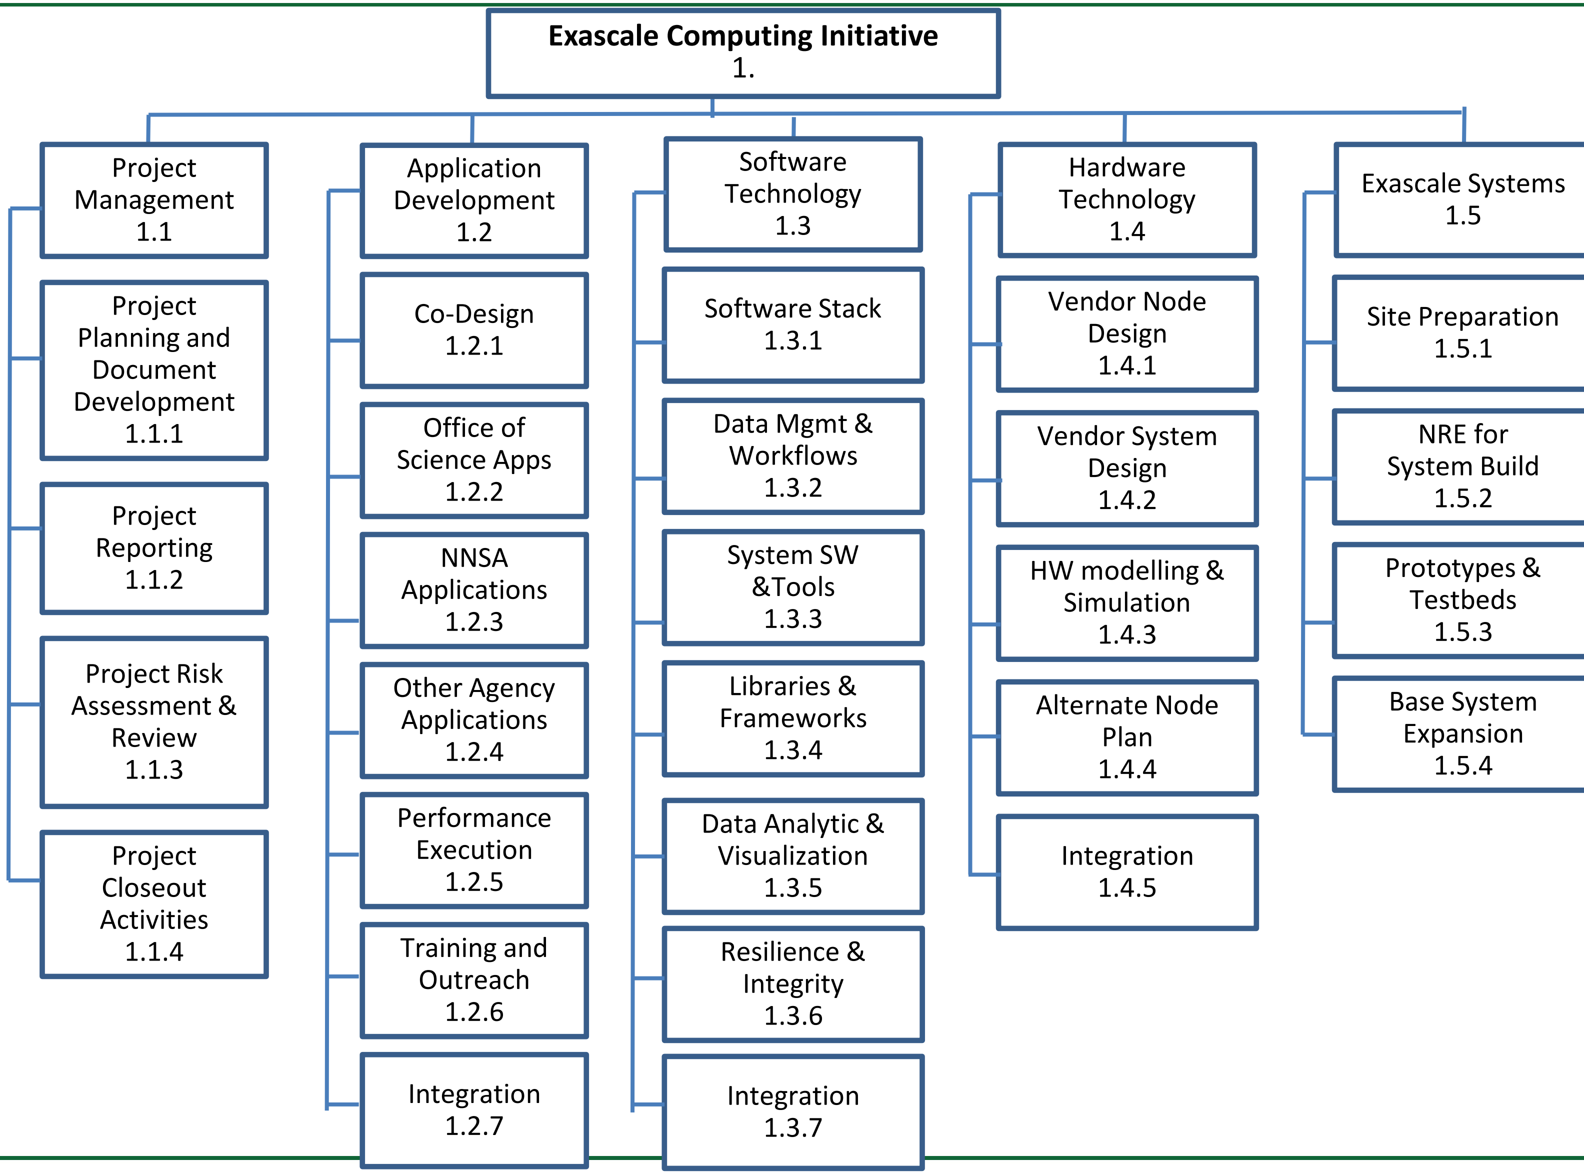

Exascale Applications

- **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. 126 application white papers received. Input will be used by ASCR/ASC in planning for the Exascale Computing Initiative**
  - 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.

ECI Integrated Project Team

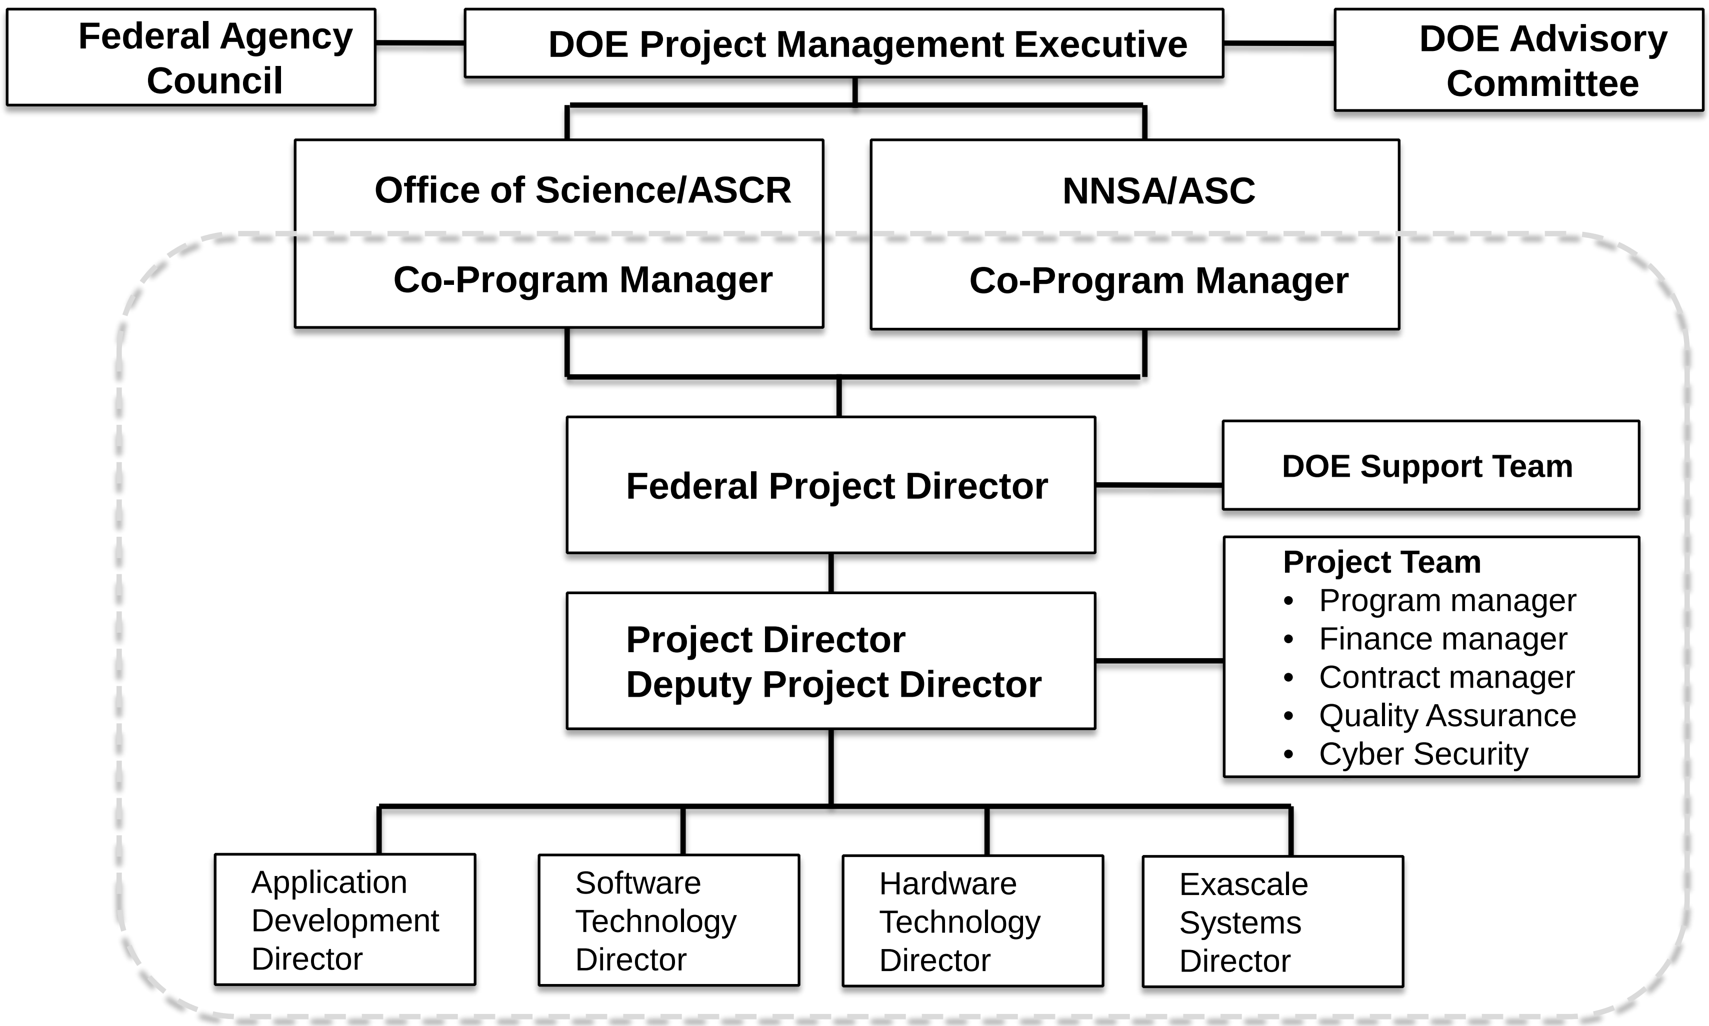

ECI Project Management Structure