U.S. DEPARTMENT OF ENERGY

DOE Office of Advanced Scientific Computing Research

Presented to the
Fusion Energy Sciences Advisory Committee

by

Steve Binkley
Associate Director

April 10, 2014

Advanced Scientific Computing Research

Delivering world leading computational and networking capabilities to extend the frontiers of science and technology

The Scientific Challenges:

- Deliver next-generation scientific and energy applications on multi-petaflop computers.
- Discover, develop and deploy exascale computing and networking capabilities.
- Partner with U.S. industry to develop the next generation computing hardware and tools for science.
- Discover new applied mathematics, computer science, and networking tools for the ultra-low power, multicore-computing future and data-intensive science.
- Provide technological innovations for U.S. leadership in Information Technology to advance competitiveness.

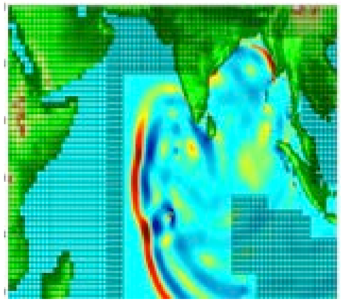

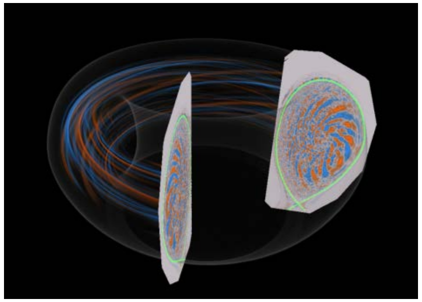

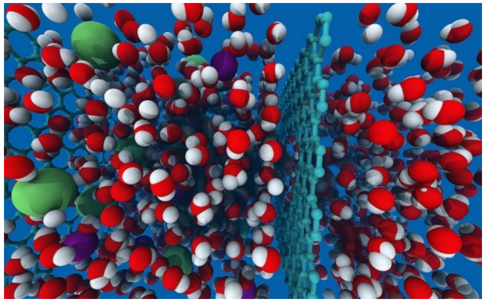

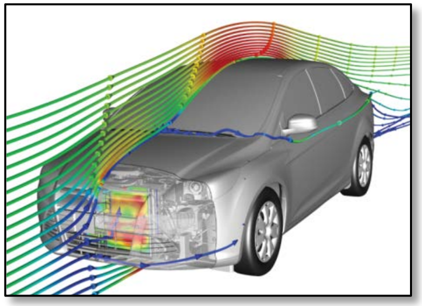

ASCR Investment Priorities

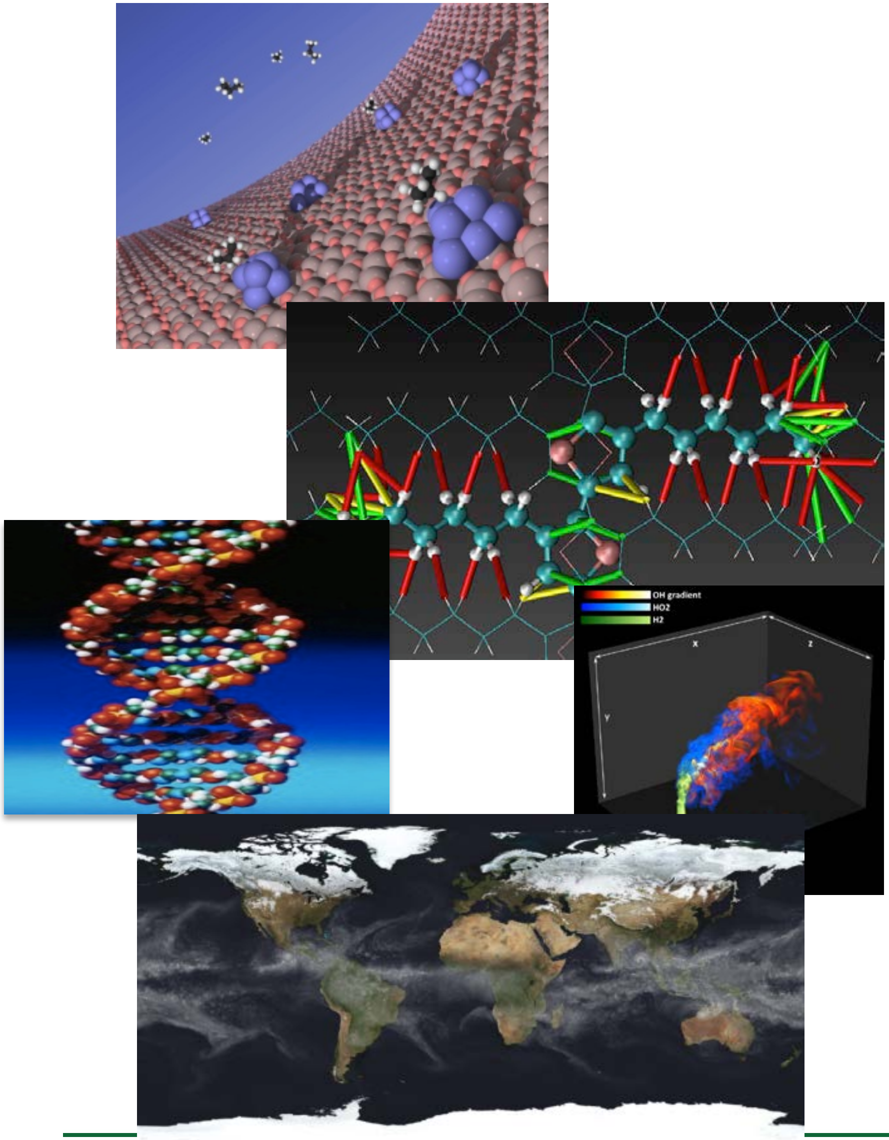

- **Investment Priorities**
  - **Exascale –** conduct research and development, and design efforts in hardware software, and mathematical technologies that will produce exascale systems in 2022.
  - **Large Scientific Data –** prepare today's scientific and data-intensive computing applications to migrate to and take full advantage of emerging technologies from research, development and design efforts.
  - **Facilities –** acquire and operate more capable computing systems, from multi-petaflop through exascale computing systems that incorporate technologies emerging from research investments.

Program Updates

- **Path toward Exascale**
  - Achieving new, higher levels of concurrency
  - Affordable power consumption
  - Programmability
  - Resiliency
- **ASCR Facilities**
  - Leadership Computing
  - National Energy Research Supercomputing Center (NERSC)
  - High-performance networking
- **Large Scientific Data**
- **Applied math, computer science, SciDAC**
  - Tools, libraries, software, partnerships to maximally utilize future HPC systems

Mission: Extreme Scale Science
Next Generation of Scientific Innovation

- **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
  - Mission-focused basic science in energy
  - Provide state-of-the-art scientific tools
  - Plan, implement, and operate user facilities
- **The next generation of advancements will require Extreme Scale Computing**
  - 1,000X capabilities of today’s computers with a similar physical size and power footprint
- **Extreme Scale Computing, cannot be achieved by a “business-as-usual,” evolutionary approach**

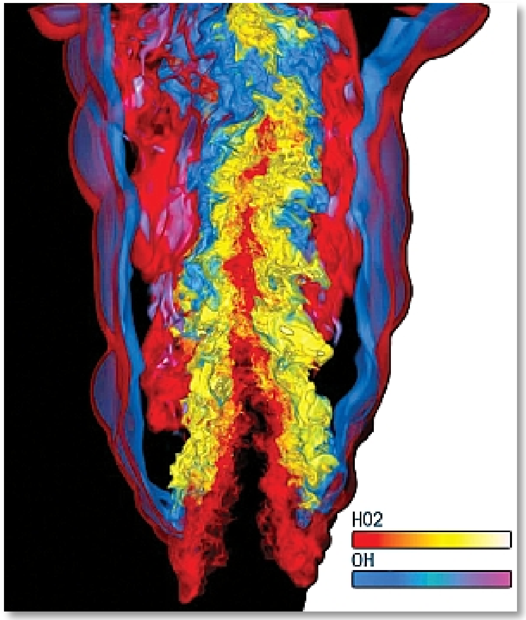

Exascale Computing
The Vision

- **Exascale computing**
  - Achieve order 10^{18} operations per second and order 10^{18} bytes of storage
  - Address the next generation of scientific, engineering, and large-data workflows
  - Enable extreme scale computing: 1,000X capabilities of today's computers with a similar size and power footprint
  - Barriers: billion-way concurrency, energy consumption
- **Productive system**
  - Usable by a wide variety of scientists and engineers
  - "Easier" to develop software & management of the system
- **Based on marketable technology**
  - Not a "one off" system
  - Scalable, sustainable technology, exploiting economies of scale and trickle-bounce effect
- **Deployed in early 2020s**

Progress in CMOS CPU Technology

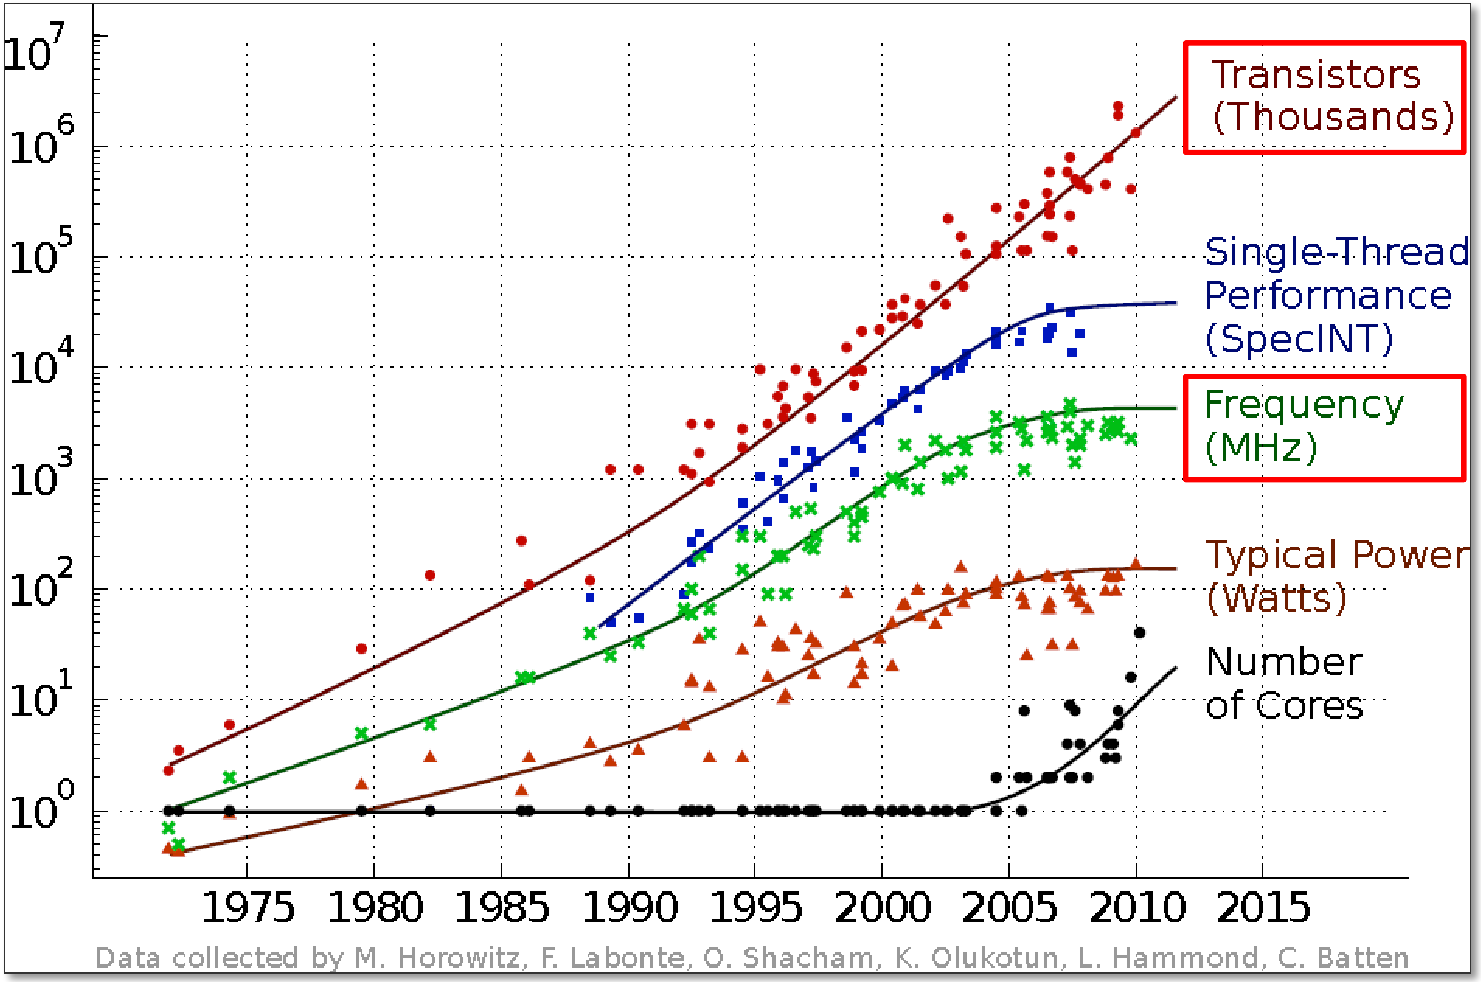

Moore's Law continues

- Transistor count still doubles every 24 months

Dennard scaling stalls

- Voltage
- Clock Speed
- Power
- Performance/clock

Microprocessor Transistor Counts 1971-2011 & Moore's Law

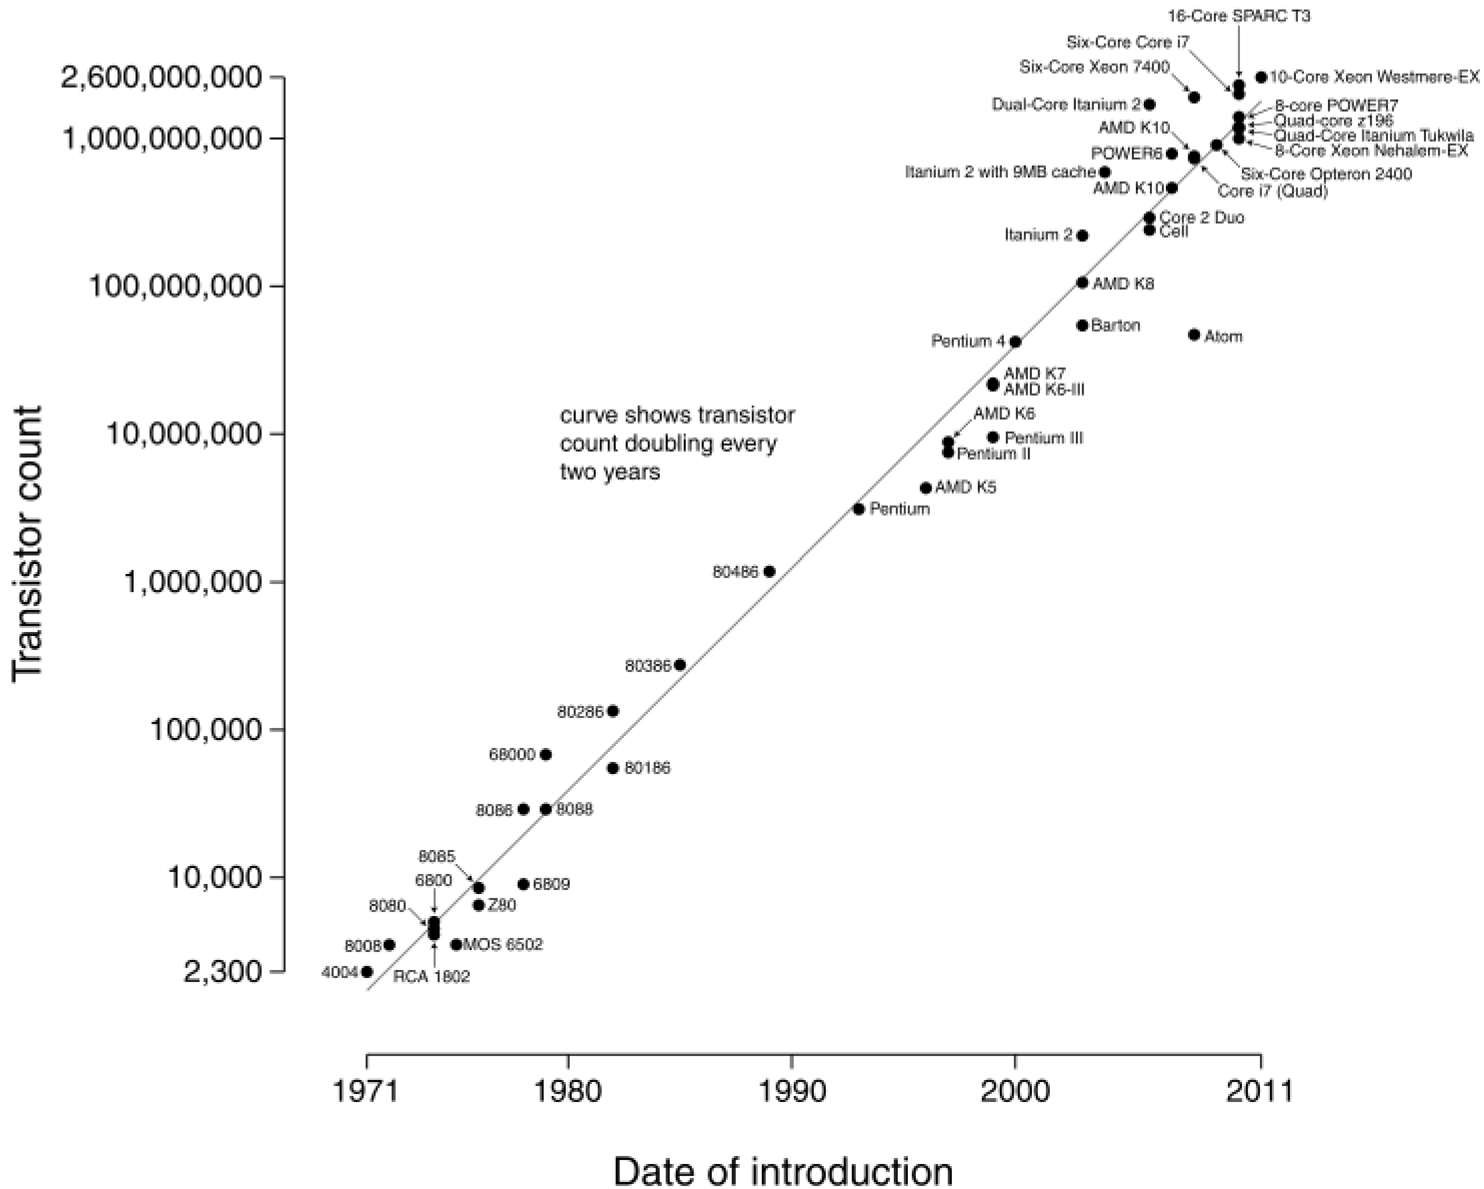

Top Ten Technical Approaches for Exascale

(Advanced Scientific Computing Advisory Committee)

→ 1. **Energy efficiency:** Creating more energy efficient circuit, power, and cooling technologies.
2. **Interconnect technology:** Increasing the performance and energy efficiency of data movement.

→ 3. **Memory technology:** Integrating advanced memory technologies to improve both capacity and bandwidth.

→ 4. **Scalable System Software:** Developing scalable system software that is power and resilience aware.
5. **Programming systems:** Inventing new programming environments that express massive parallelism, data locality, and resilience
6. **Data management:** Creating data management software that can handle the volume, velocity and diversity of data that is anticipated.

→ 7. **Exascale algorithms:** Reformulating science problems and refactoring their solution algorithms for exascale systems.
8. **Algorithms for discovery, design, and decision:** Facilitating mathematical optimization and uncertainty quantification for exascale discovery, design, and decision making.

→ 9. **Resilience and correctness:** Ensuring correct scientific computation in face of faults, reproducibility, and algorithm verification challenges.
10. **Scientific productivity:** Increasing the productivity of computational scientists with new software engineering tools and environments.

http://science.energy.gov/~/media/ascr/ascac/pdf/reports/2013/report.pdf

Use of Co-Design in DOE Exascale Strategy

Approach: Use scientific workflow requirements to guide architecture and system software and use technology capabilities to design algorithms and software

- Three co-design centers focus on specific application domains:
  - **ExaCT**: Combustion simulation (uniform and adaptive mesh)
  - **ExMatEx**: Materials (multiple codes)
  - **CESAR**: Nuclear engineering (structures, fluids, transport)
- Create "proxy apps", next "proxy workflows"
  - Scaled down versions of "full" code
  - Selects parts/patterns from code to drive programming/architecture

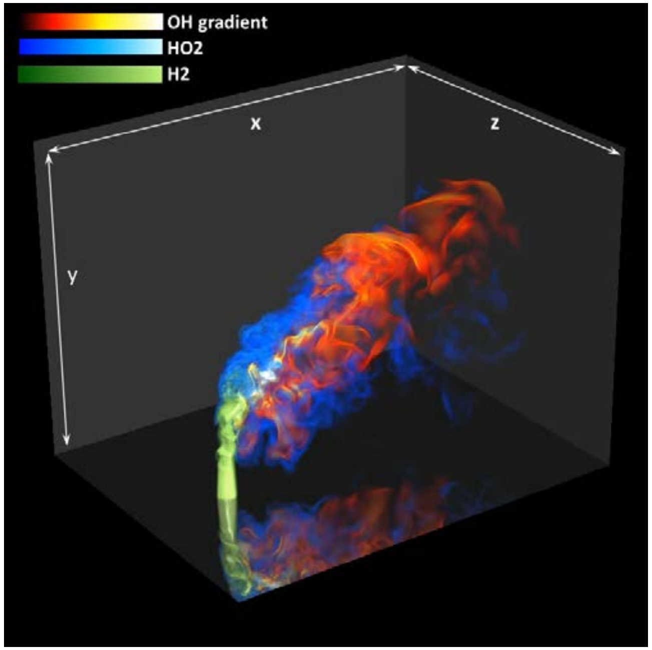

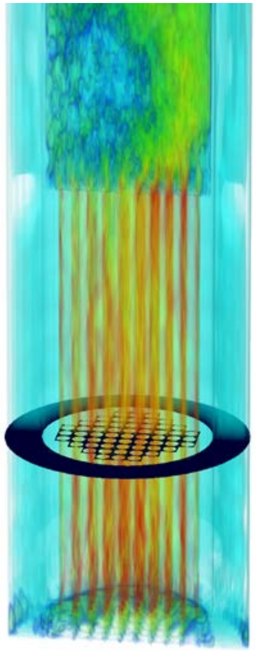

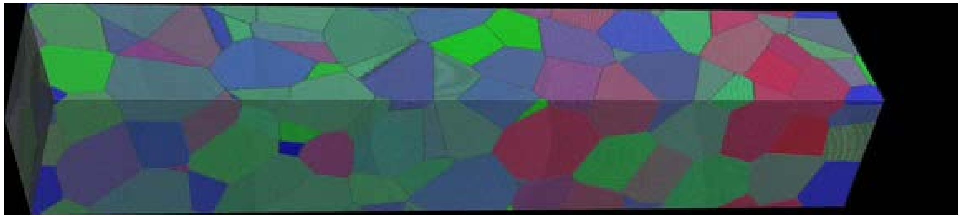

Research & Evaluation Prototypes

- **FastForward:** In FY 2012, Research and Evaluation Prototypes activity worked with NNSA to award $95M (total, including cost sharing, over two years) for innovative R&D on critical technologies – memory, processors and storage – needed to deliver next generation capabilities within an affordable energy footprint.
  - **Funded Projects:**
    - AMD: processors and memory for extreme systems;
    - IBM : memory for extreme systems;
    - Intel Federal: energy efficient processors and memory architectures;
    - Nvidia: processor architecture for exascale computing at low power; and
    - Whamcloud: storage and I/O (input/output) – *subsequently bought by Intel*.
- The FY 2015 increase takes FastForward research to the next level:
  - Lab/vendor partnerships (+$12,216K)
    - develop prototypes of the most promising mid-term technologies from the Fast Forward program for further testing.
  - Nonrecurring engineering (+$7,934K)
    - Incorporate near-term technologies from Fast Forward into planned facility upgrades at NERSC, ALCF and OLCF.

High Performance Computing (HPC) US Federal Government Investments

- US federal High Performance Computing (HPC) investments
  - Made pivotal investments in the computer industry at critical times
  - During stable times, no investment is required or requested
- Previous US Federal HPC investments Fueled major HPC advances
  - 1946 ENIAC: start of electronic digital computing
  - 1951 ERA-1101: technical computing
  - 1972 ILLIAC IV: parallel computing
  - 1970s-80s: CDC and Cray 1
  - 1993 Cray T3D: massively parallel computing
  - 2004 IBM BG: low power computing
- DOE Computational Facilities
  - ORNL Titan, Cray XK7, 27 PF
  - LLNL Sequoia, IBM BG/Q 20 PF
  - ANL Mira, IBM BG/Q, 10 PF
  - LBNL Edison, Cray XC30, 2 PF
  - LANL Cielo, Cray XE6, 1.1 PF

1 PF = 10^{15} floating point operations per second

International Competitors Have Ambitious Plans

China (goal: 2020)

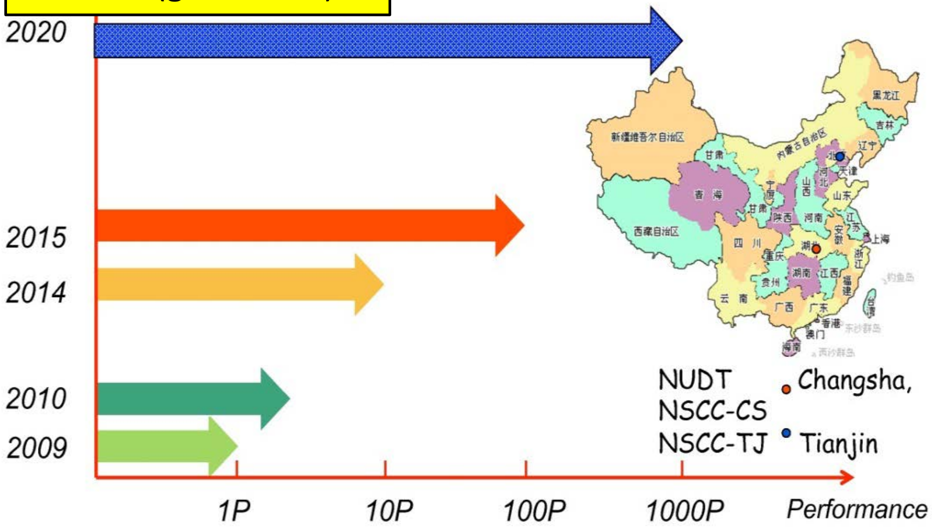

Japan (goal: 2020)

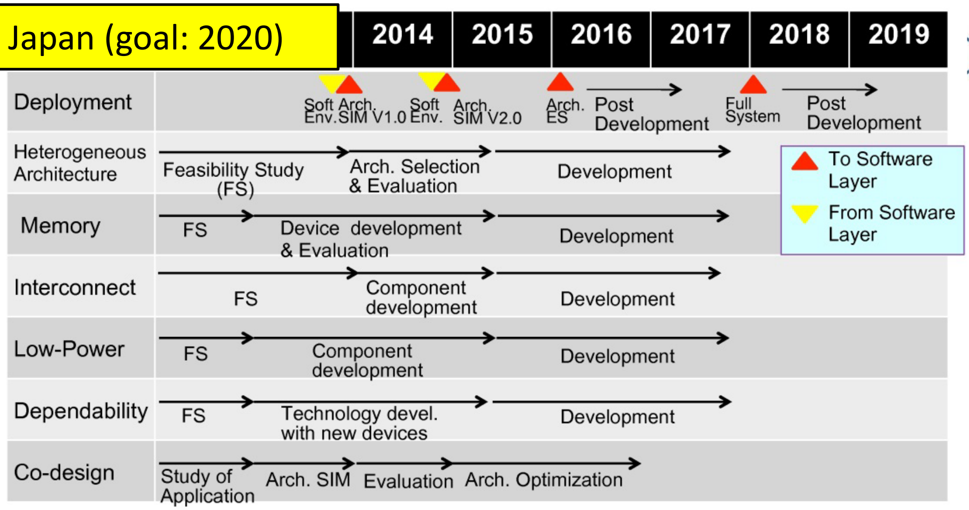

European Union (goal: 2020)

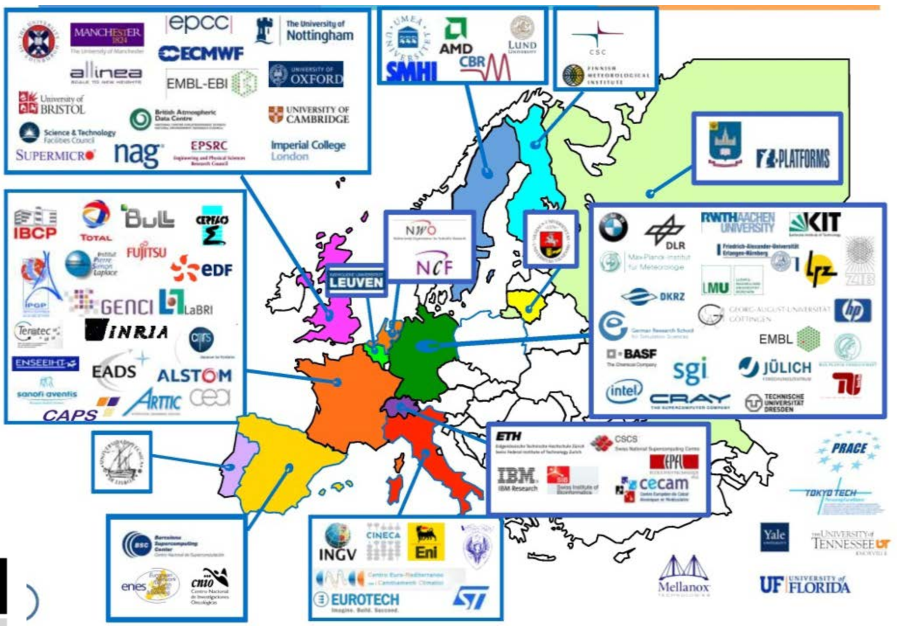

Other countries:
Russia
South Korea
India
Singapore
Australia

EU: J-Y. Berthou, IESP Workshop, Kobe, Japan, April 2012
China: Y. Lu, IESP Workshop, Kobe, Japan, April 2012
Japan: M. Kondo, IESP Workshop, Kobe, Japan, April 2012

Exascale Computing

- **What isn't Exascale computing?**
  - Exaflops Linpack Benchmark Computer
  - Just a billion floating-point arithmetic units packaged together
- **What is Exascale computing?**
  - 1,000X performance over a "Petaflops" system (exaflops sustained performance on complex, real-world applications)
  - Similar power and space requirements as a Petaflops computer
  - High programmability, generality, and performance portability
- **Required areas of investment**
  - New system designs / execution models
  - Enabling technologies (such as, node & interconnects)
  - Scalable algorithms
  - System Software

Leadership and Production Computing Facilities

Titan:
- Peak performance of 27.1 Petaflops
- 18,688 Hybrid Compute Nodes
- 8.9 MW peak power

Mira:
- Peak performance of 10 Petaflops
- 49,152 Compute Nodes
- 4.8 MW peak power

Edison XC30:
- Peak performance 2.4PF
- 5,576 Compute Nodes
- 2.1 MW peak power

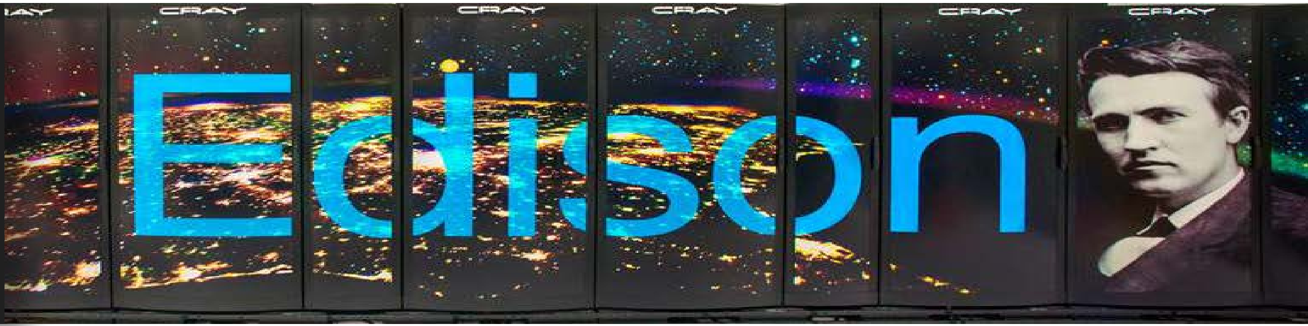

ESNET: World’s First Continental 100 Gbs Production Network

Network Specifications:

- 16 Alcatel Lucent routers, plus new Ciena optical platform
- Network can grow to 88 independent 100G channels
- Deploying 100G production connections at ANL, BNL, FNAL, LBNL, LLNL, NERSC, ORNL in next 6 months
- Improved fiber and optical diversity
- >99.99% availability to Labs in CY2012

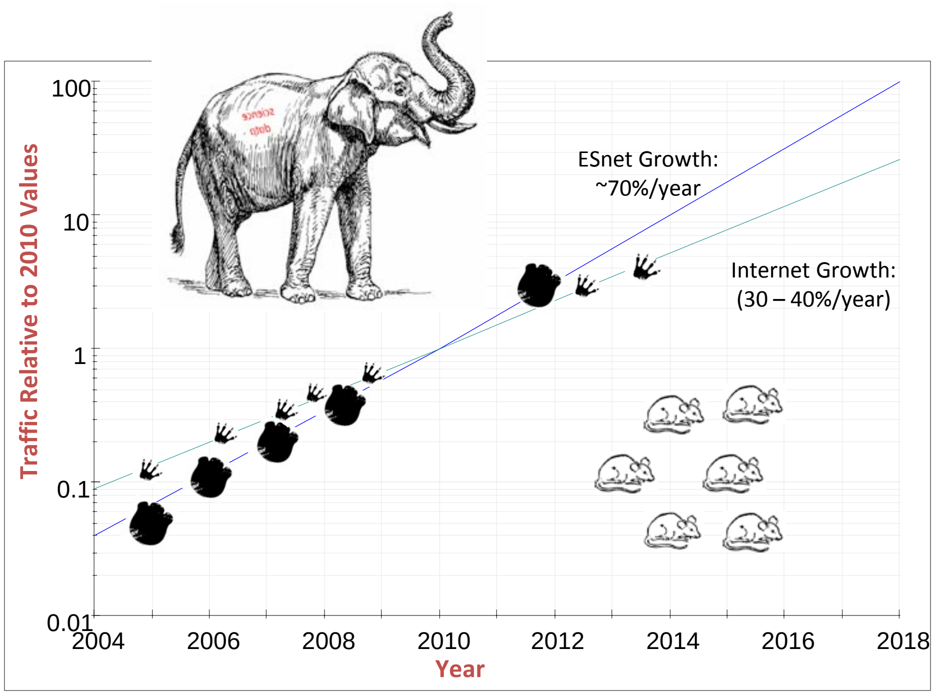

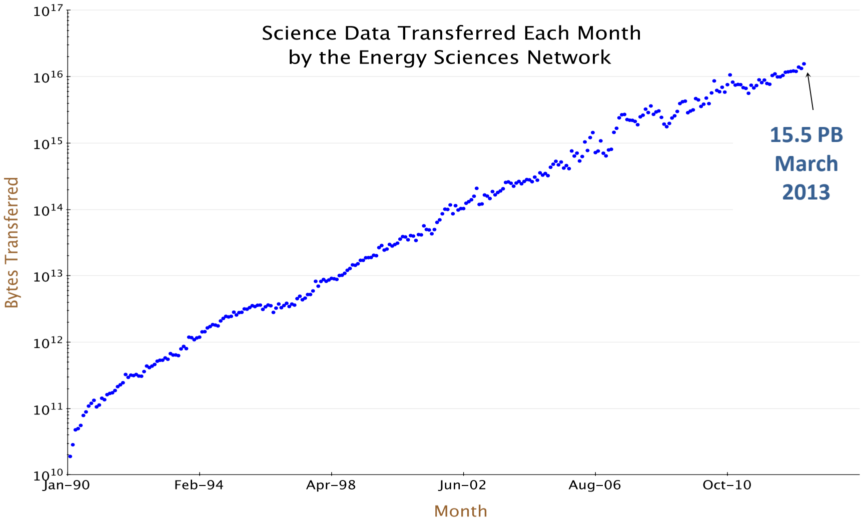

ESnet grows twice as fast as the commercial Internet, due largely to *elephant flows* and data intensive science.

Distribution of Users at the ~30 SC Facilities 2013

Nearly ¾ of users do their work at ASCR or BES facilities

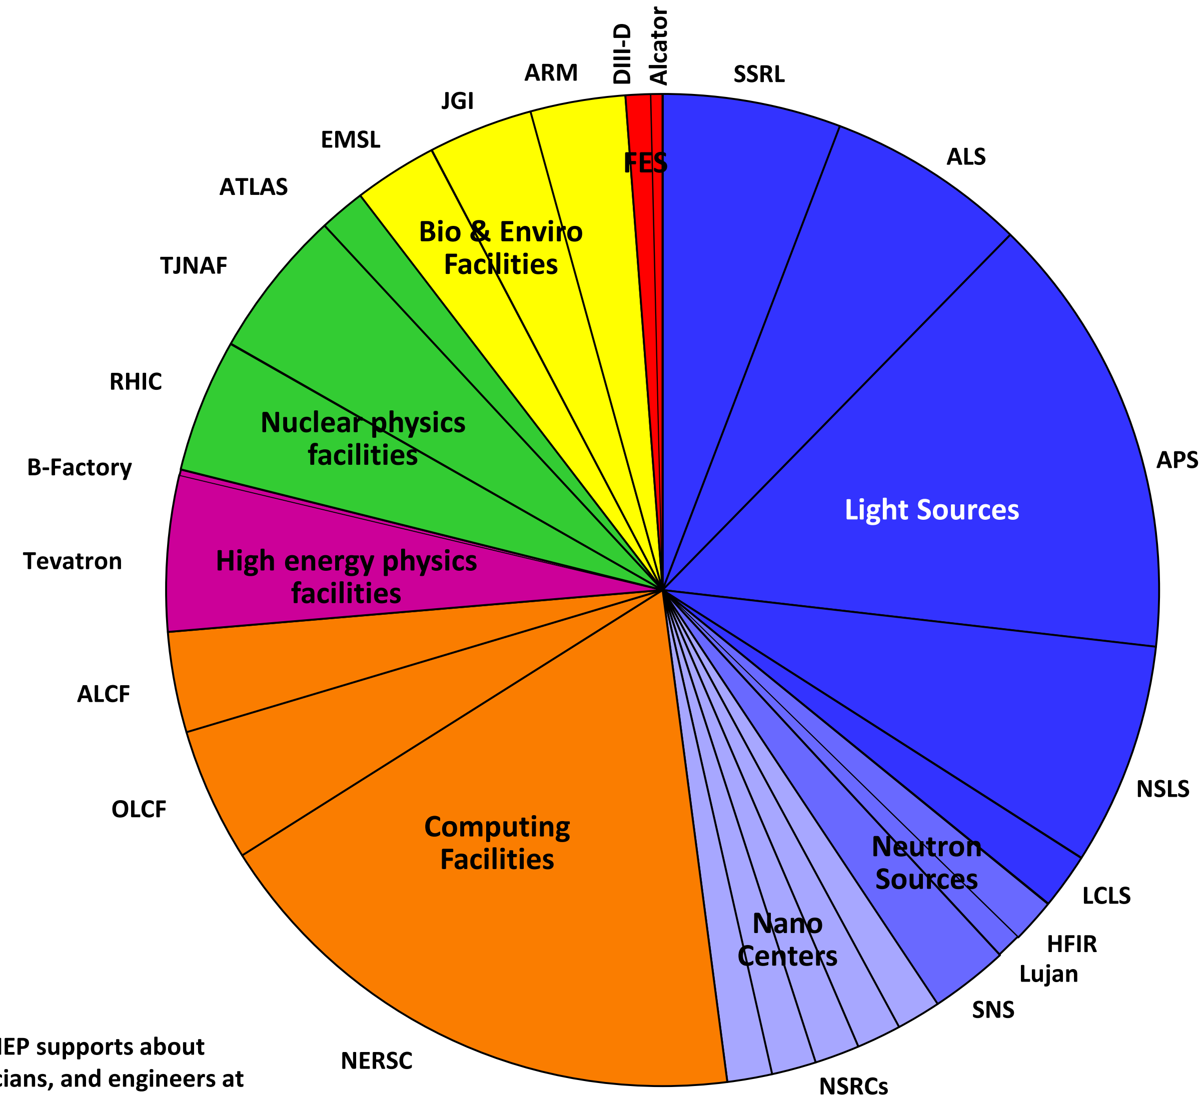

Does not include LHC; HEP supports about 1,700 scientists, technicians, and engineers at the LHC.

Fusion Simulations Are a Major Element of NERSC Computing

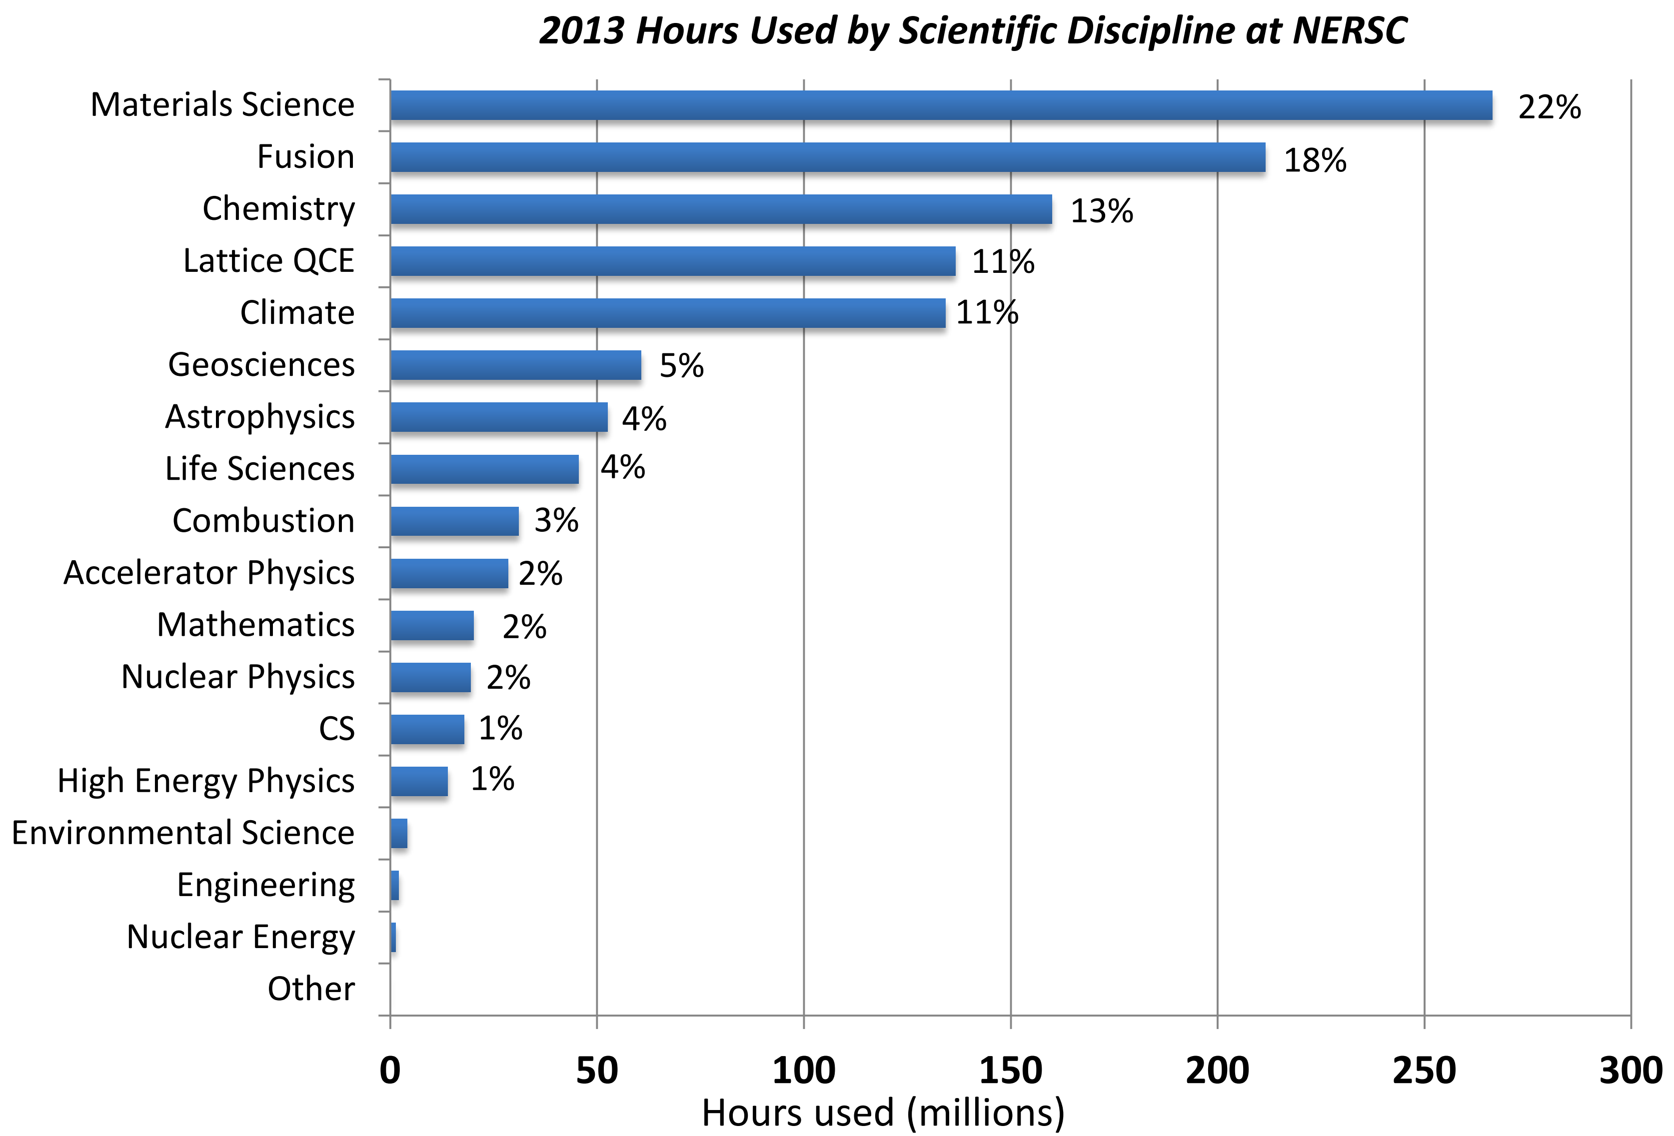

Top Fusion Codes at NERSC

Breakdown of Hours used at NERSC in 2013 by Application

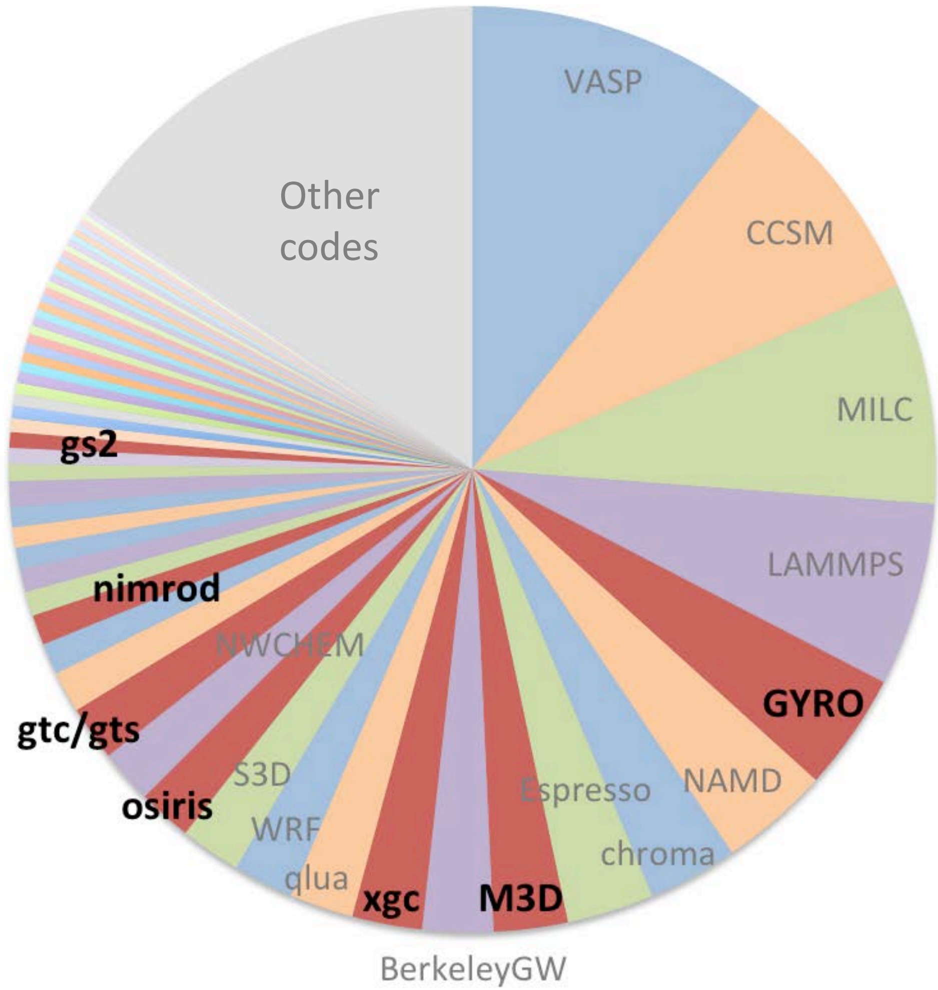

- Fusion codes are highlighted in BOLD
- Fusion codes make up ~18% of the NERSC workload

INCITE: Innovative and Novel Computational Impact on Theory and Experiment

INCITE promotes transformational advances in science and technology through large allocations of computer time, supporting resources, and data storage at the Argonne and Oak Ridge Leadership Computing Facilities (LCFs) for computationally intensive, large-scale research projects.

2012 INCITE Allocations

Awarded 1.67 billion hours
for CY 2012

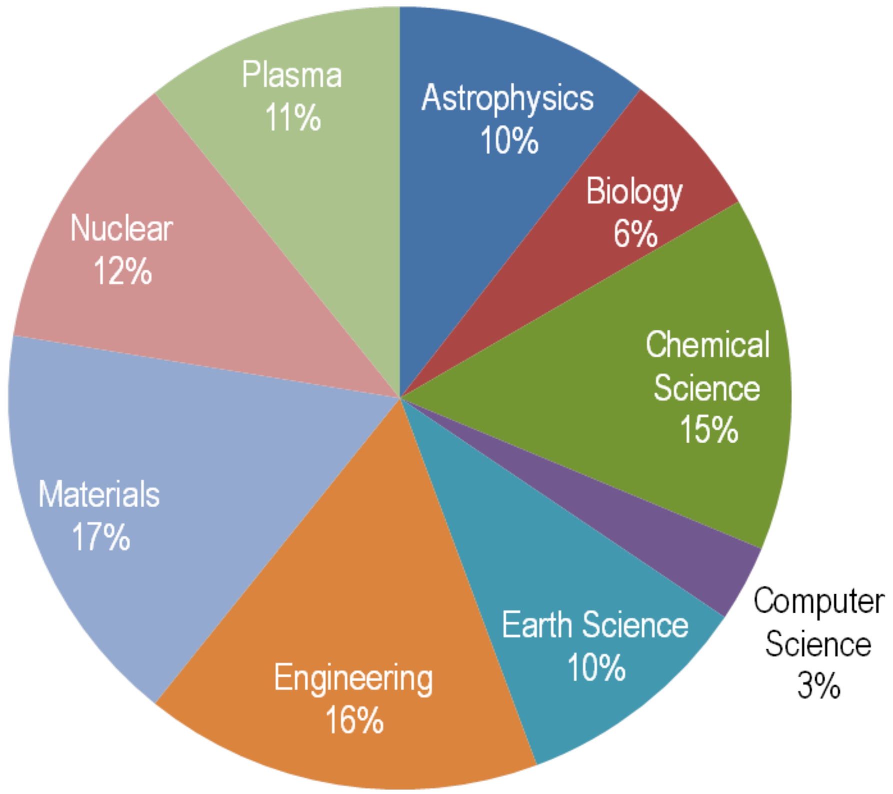

28 new projects and 32 renewal projects

The Bleeding 'Edge' of Fusion

INCITE Program
PI: CS Chang
Princeton Plasma Physics Laboratory
Hours used: 102 million

Science Objectives and Impact

- **Predictions validated and new insight gained through analysis of DIII-D tokamak results**
- **First-principles simulation of edge physics on Titan can increase understanding of ITER**

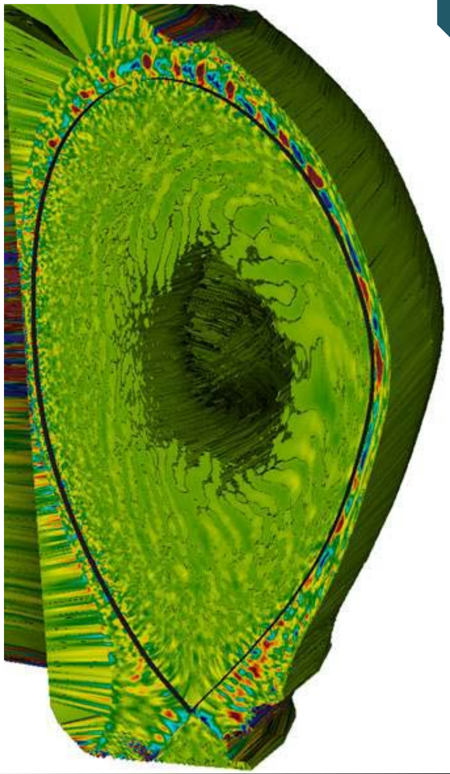

This figure shows the turbulence front from the plasma edge in the DIII-D reactor being spread inward in multiscale interaction with the evolving background profile under the central heat source. Eventually, the whole volume becomes turbulent, with the spatial turbulence amplitude distribution being just enough to produce the outward heat transport to expel the centrally deposited heat to the edge. This is how the plasma profile, the heat source and the turbulence self-organize.
- Visualization by David Pugmire ORNL)

Performance & OLCF Contribution

- Gyrokinetic, particle-in-cell code
  - Fortran, MPI based
- OLCF provided the necessary HPC power; Up to 90% of Titan's max capability has been utilized for the study
- OLCF liaison contributed to algorithmic improvement, which made the XGC1 code use GPUs and CPUs efficiently, with linear scalability to maximal size of Titan and provided viz support

Science Results

- Simulated nonlinear coherent turbulence structures (called "blobs") in the plasma edge of DIII-D reactor
- Identified the momentum source and its inward transport process
- Predicted the divertor heat load distribution

All for the first time in first-principles calculations of a tokamak reactor

Delivering Capabilities that Keep the U.S. Competitive
"ASCR inside"

A few ASCR Technologies and the Companies that Use them

- **MPICH** – Message passing library
  "MPICH's impact comes from the fact that since it is open source, portable, efficient, and solid, most computer vendors have chosen it as the foundation of the MPI implementation that they supply to their customers as part of their system software." - Rusty Lusk, MPICH consortia
  ""MPICH is critical to the development of the F135 engine, which will power America's next-generation Joint Strike Fighter," - Robert Barnhardt, VP, Pratt & Whitney
- **Fastbit** – Search algorithm for large-scale datasets
  "FastBit is at least 10 times, in many situations 100 times, faster than current commercial database technologies" – Senior Software Engineer, Yahoo!
- **OSCARS** - On-demand virtual network circuits
  "It used to take three months, 13 network engineers, 250 plus e-mails and 20 international conference calls to set up an inter-continental virtual circuit. With OSCARS and collaborative projects, we can establish this link in 10 minutes." - Chin Guok, ESnet network engineer
- ***perfSONAR*** - network performance monitoring
  "These tools give us better visibility into the network, allowing us to troubleshoot performance issues quickly." -- Internet2 Network Performance Workshop participant

Innovation through Industrial Partnerships

BOEING	P&G	Ford	United Technologies Research Center	BOSCH	Smart Truck
Aircraft design	**Consumer products**	**Engine cycle-to-cycle variation**	**Jet engine efficiency**	**Li-ion batteries**	**Long-haul truck fuel efficiency**
Simulating takeoff and landing scenarios improved a critical code for estimating characteristics of commercial aircraft, including lift, drag, and controllability	Leadership computing and molecular dynamics software advanced understanding of chemical processes that can limit product shelf life	Emerging model of engine cyclic variation will apply thousands of processors to a challenging problem	Accurate predictions of atomization of liquid fuel by aerodynamic forces enhance combustion stability, improve efficiency, and reduce emissions	New classes of solid inorganic Li-ion electrolytes could deliver high ionic and low electronic conductivity and good electrochemical stability	Simulations reduced by 50% the time to develop a unique system of add-on parts that increases fuel efficiency by 7−12%

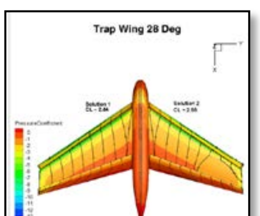

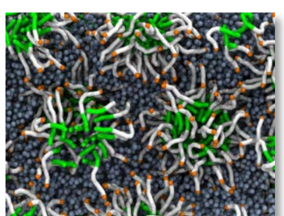

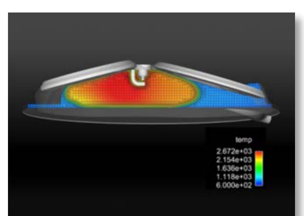

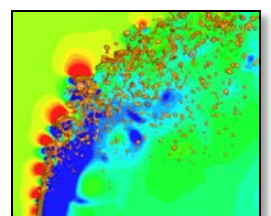

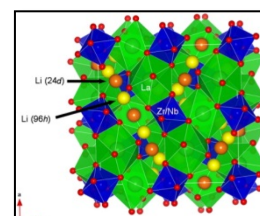

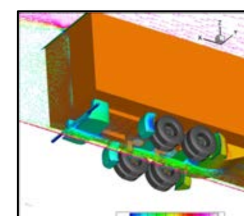

U.S. DEPARTMENT OF ENERGY

Mission: Extreme-Scale Science Data Explosion

Genomics

Data Volume increases to 10 PB in FY21

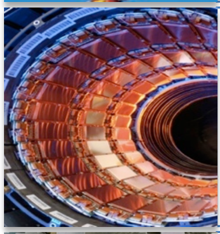

High Energy Physics (Large Hadron Collider)

15 PB of data/year

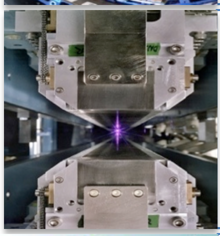

Light Sources

Approximately 300 TB/day

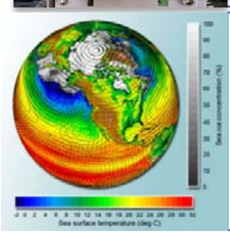

Climate

Data expected to be hundreds of 100 EB

Driven by exponential technology advances

Data sources

- Scientific Instruments
- Scientific Computing Facilities
- Simulation Results
- Observational data

Big Data and Big Compute

- Analyzing Big Data requires processing (e.g., search, transform, analyze, …)
- Extreme scale computing will enable timely and more complex processing of increasingly large Big Data sets

LCLS Partnership
Taking Snapshots of the Catalytic Reaction in Photosystem II

- X-ray Free-electron laser, 10^{17}-fold greater peak brightness than ALS 5.0.2
- 50 fs X-ray pulses above the Mn absorption edge
- Diffract before destroy approach
- Simultaneously detect X-ray emission

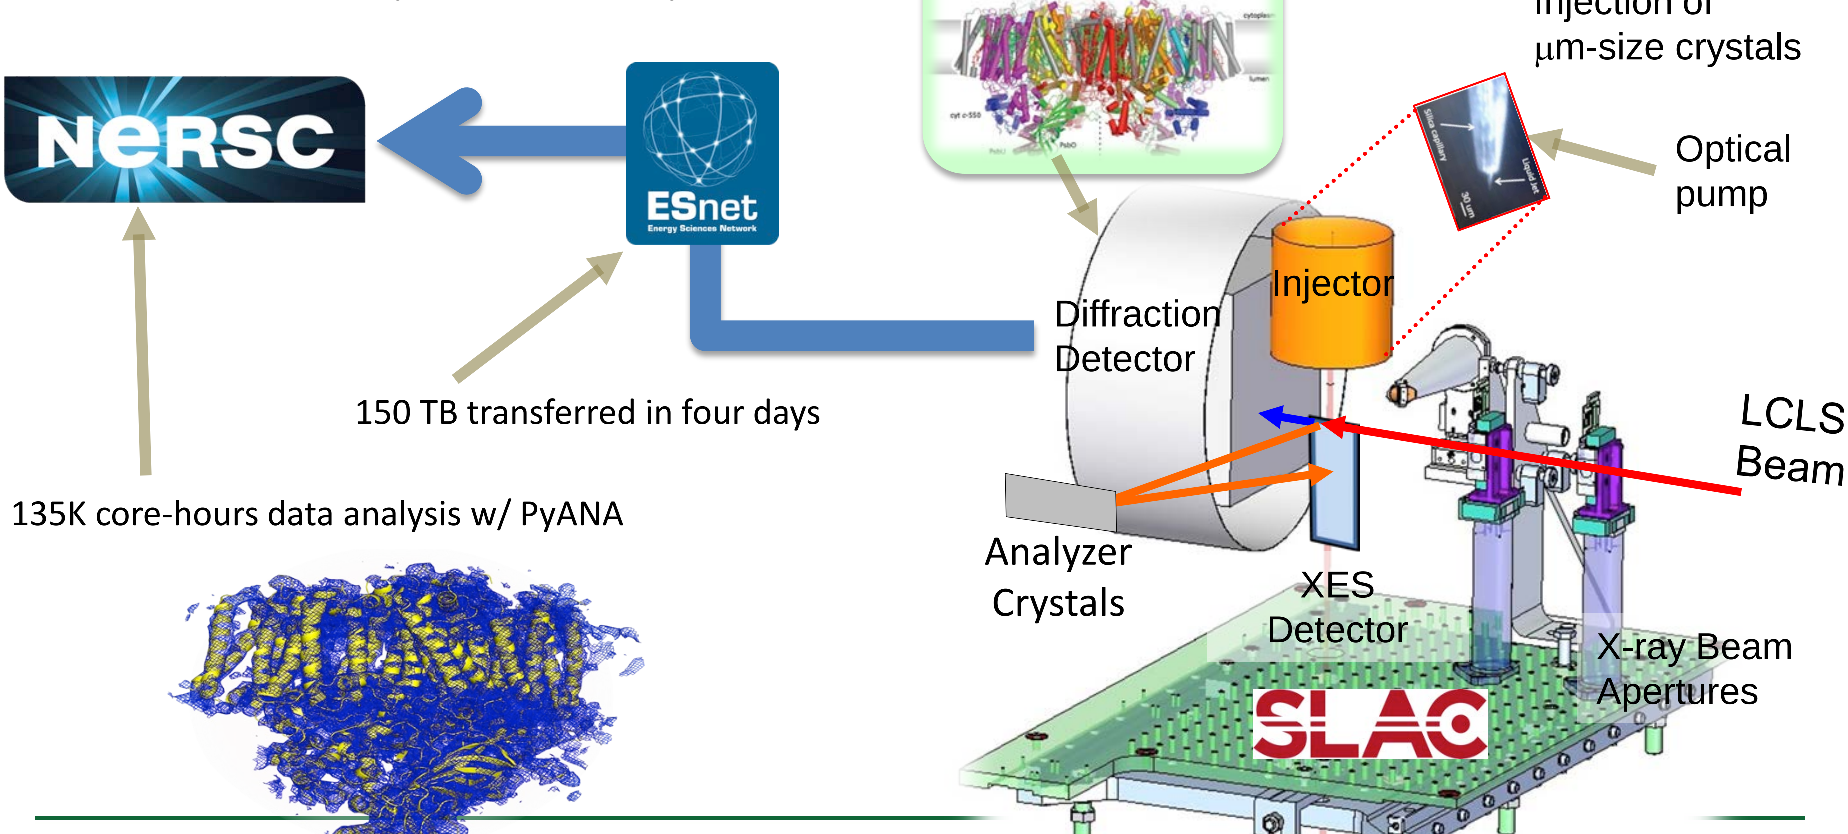

Kern et al (2012) PNAS 109: 9721, Sierra et al (2012) Acta Cryst D68: 1584, Mori et al (2012) PNAS 109: 19103

ASCR Partnerships Across the Office of Science

- SciDAC
  - Partnerships facilitate transfer of ASCR research results into SC applications
  - Third generation, 18 partnerships across SC and 4 SciDAC Institutes

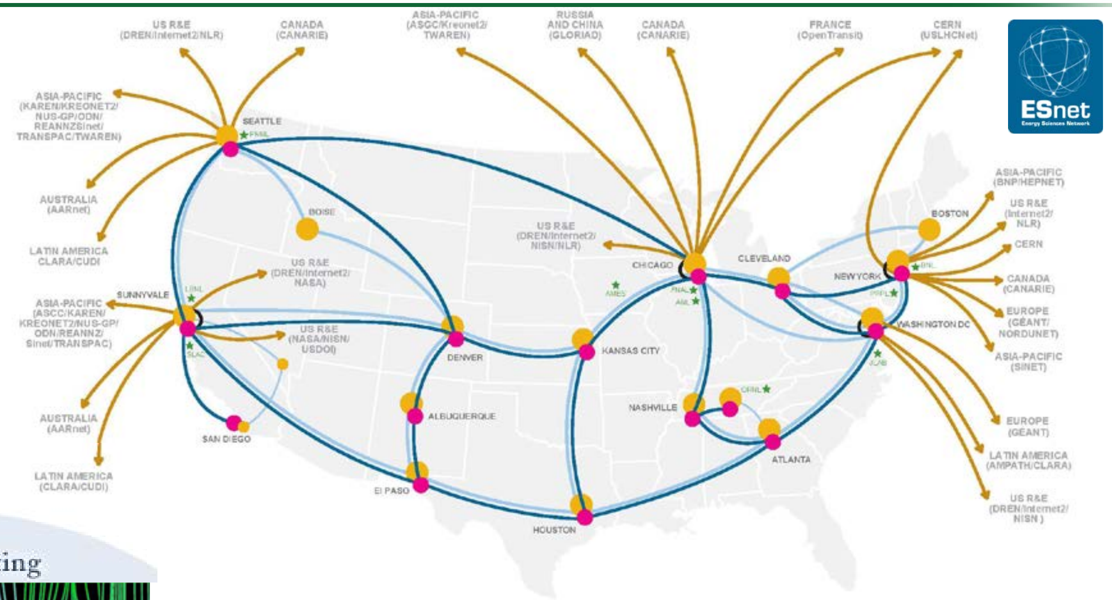

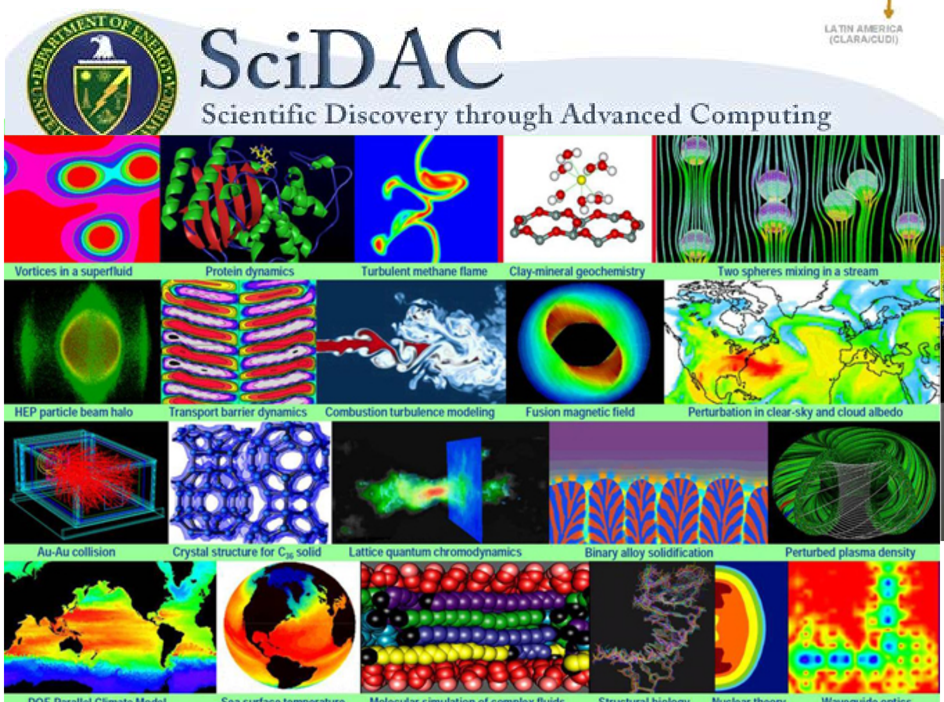

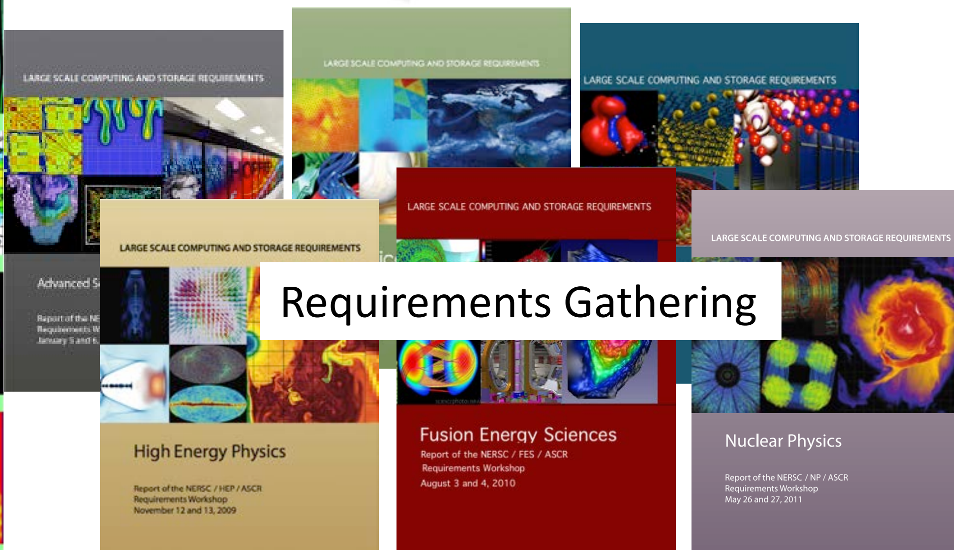

Requirements Gathering

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

Scientific Discovery through Advanced Computing (SciDAC): Multiscale Integrated Modeling for Fusion Energy Science

Funding Opportunity Number: DE-FOA-0001096
Announcement Type: Initial
CFDA Number: 81.049

Award sizes will range from $750,000 per year to $1,250,000 per year.

Issue Date: 03/06/2014
Letter of Intent Due Date: 03/26/2014 at 4:59 PM Eastern Time
A Letter of Intent is strongly encouraged
Pre-Application Due Date: Not Applicable
Application Due Date: 05/02/2014 at 4:59 PM Eastern Time

Scientific Discovery through Advanced Computing (SciDAC): Multiscale Integrated Modeling for Fusion Energy Science

The Office of Fusion Energy Sciences (FES) and the Office of Advanced Scientific Computing Research (ASCR), Office of Science, U.S. Department of Energy (DOE), announce their interest in receiving applications from collaborative groups of investigators for developing an integrated simulation capability for fusion energy science. More specifically, applications are solicited for the development of advanced multiphysics and multiscale integrated simulation capabilities for magnetically confined plasmas addressing problems of direct relevance to burning plasma science and ITER. While developing a full Whole Device Modeling (WDM) simulation capability is beyond the scope of this FOA, this is intended to be a first step toward this goal. Responsive applications are expected to integrate the most critical physical processes across all relevant regions and on all relevant temporal and spatial scales, using an appropriately justified combination of first principles models and high physics fidelity reduced models. Simulation codes should be able to exploit the massive concurrency of the SC leadership class computing facilities and not merely their high capacity. Applications focused solely on the development of computational frameworks are not responsive to this FOA. However, since advanced computational frameworks are essential for enabling and facilitating the coupling and integration of component modules, allocation of resources to adapt, maintain, upgrade, and extend existing frameworks, including those developed by the Fusion Simulation Prototype Centers or “proto-FSPs”, is permissible provided they satisfy the above stated requirement of exploiting the capabilities of the SC leadership computing facilities.

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