ASCR-HEP Partnerships

Presented to HEPAP

by

Barb Helland
SC/Advanced Scientific Computing Research

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.

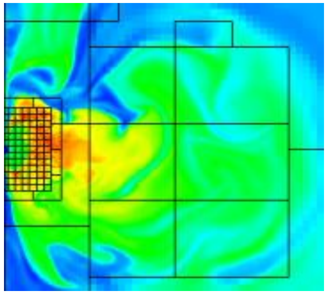

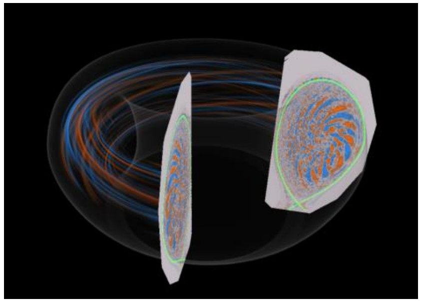

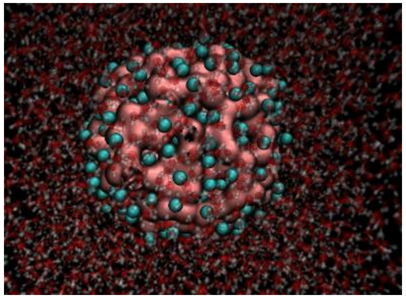

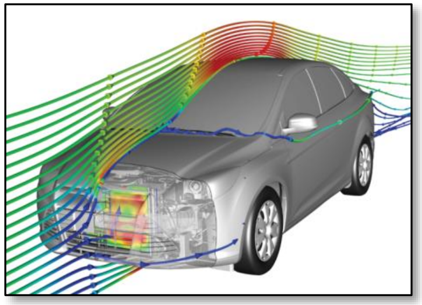

ASCR's Research

- Applied Mathematics
  - *Supports the research and development of applied mathematics models, methods and algorithms for understanding natural and engineered systems related to DOE's mission … with focus on discovery of new applied mathematics, for the ultra-low power, multicore-computing, and data-intensive future*
- Computer Science
  - *Supports the research and development of today's and tomorrow's leading edge computers tools for DOE science and engineering and extracting scientific information, discovery and insight from massive data from experiments and simulation…with focus on exascale implications.*
- Partnerships
  - *CoDesign and SciDAC partnerships to pioneer the future of scientific applications;*
- Next Generation Networks for Science
  - *Building upon results from Computer Science and Applied Mathematics, supports research and development of integrated software tools and advanced network services to use new capabilities in ESnet to integrate science workflows to network infrastructures and experiments and enable large-scale scientific collaborations.*

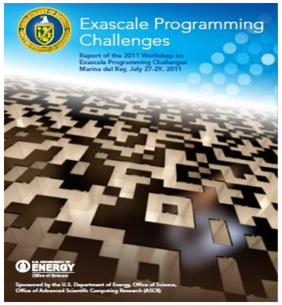

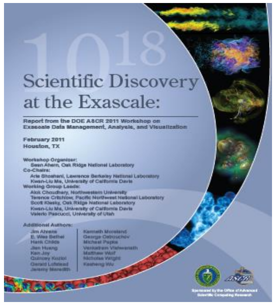

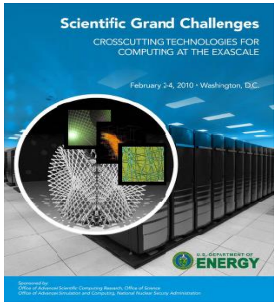

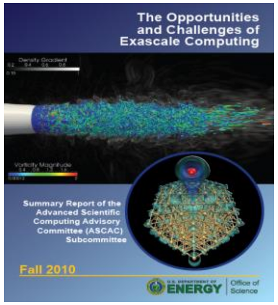

ASCR Partnerships Across the Office of Science

SciDAC
Scientific Discovery through Advanced Computing

- SciDAC focuses on the high end of high performance computing and addresses two challenges:

  to broaden the HPC community to increase the impact of high performance computing on the Department's missions, and

  to ensure that further progress at the forefront of HPC technology is rapidly assimilated into DOE applications.

- A decade of effort has enabled this program to simultaneously meet both of these challenges.

- SciDAC is leading U.S. industry into new ways to use computing to improve competitiveness.

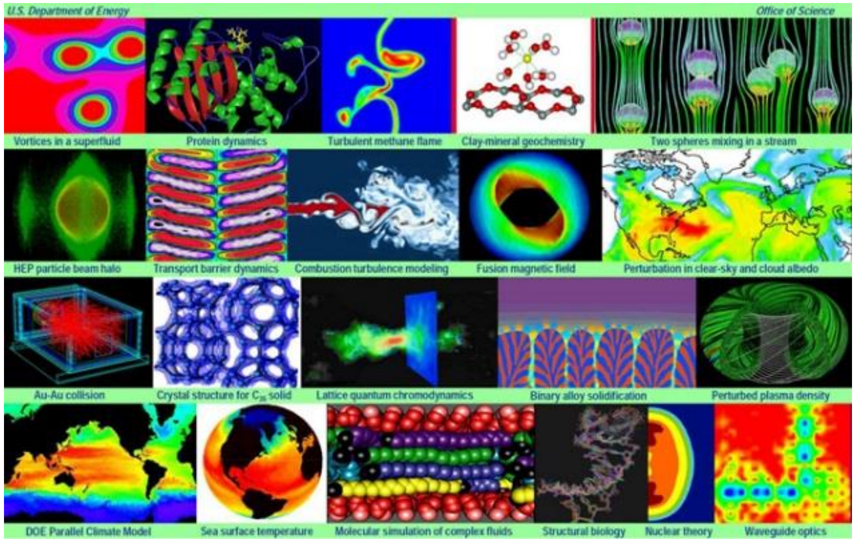

Scientific Discovery Through Advanced Computing (SciDAC)

Collaborations of applied mathematicians, computer scientists and domain scientists to advance scientific frontiers through modeling, simulation and analysis

	Time Frame	Activities	Result
SciDAC-1	2001 - 2006	Created the scientific software infrastructure for parallel computing; funded collaborations in DOE's domain sciences	Science at the Tera-scale; Grid computing tools
SciDAC-2	2006 - 2011	Added DOE science domains; enhanced university involvement; provided outreach to broader scientific community	Science to the Peta-scale; Grids for high-energy physics & earth systems modeling
SciDAC-3	2011 - 2016	Focus on science domains that directly impact energy; Move toward predictive simulations; Exploit capabilities of multi-core & emerging hybrid architectures; co-design	Science at scale (e.g. 2-20 pflops)
SciDAC-4	2016 - 2020	Focused on ASCR partnerships with Office of Science programs to efficiently use pre-exascale computing resources to advance their science goals.	Science at scale (e.g. 100-200 pflops)

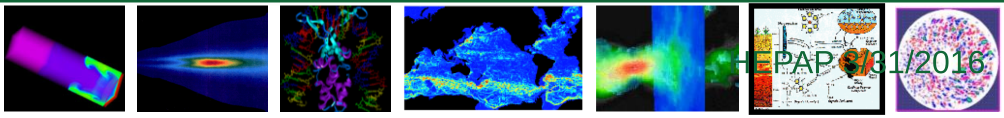

SciDAC 3

Biological and Environmental Research

- Computationally Efficient Schemes for BioGeochemical Cycles
- Multiscale Methods for Accurate, Efficient & Scale-Aware Models of the Earth System
- Predicting Ice Sheet and Climate Evolution at Extreme Scales

Basic Energy Sciences

- Advanced Modeling of Ions in Solutions, on Surfaces and in Biological Environments
- Developing Advanced Methods for Excited State Chemistry in the NWChem Software Suite
- Optimizing Superconductor Transport Properties through Large-Scale Simulation
- Scalable Computational Tools for Discovery and Design: Excited State Phenomena in Energy Materials
- Simulating the Generation, Evolution and Fate of Electronic Excitations in Molecular and Nanoscale Materials with First Principles Methods
- Discontinuous methods for massively parallel Quantum Molecular Dynamics: Li-ion interface dynamics from first principles

Fusion Energy Sciences

- Partnership for Edge Physics Simulation
- Plasma Surface Interactions

High Energy Physics

- **Community Project for Accelerator Science and Simulation**
- **Computation-Driven Discovery for the Dark Universe/ Computing the Sky**
- **Searching for Physics Beyond the Standard Model**
- **Optimizing High Energy Physics Data Management and Analysis Capabilities: the Exaflop/Exabyte Nexus**

Nuclear Physics

- A MultiScale Approach to Nuclear Structure and Reactions
- Computing Properties of Hadrons, Nuclei and Nuclear Matter from Quantum Chromodynamics
- Nuclear Computational Low Energy Initiative

ASCR – GEANT4 Partnership (2013-2014)

A SciDAC like partnership brought together computer scientists and physicists to enable computational physics on new architectures

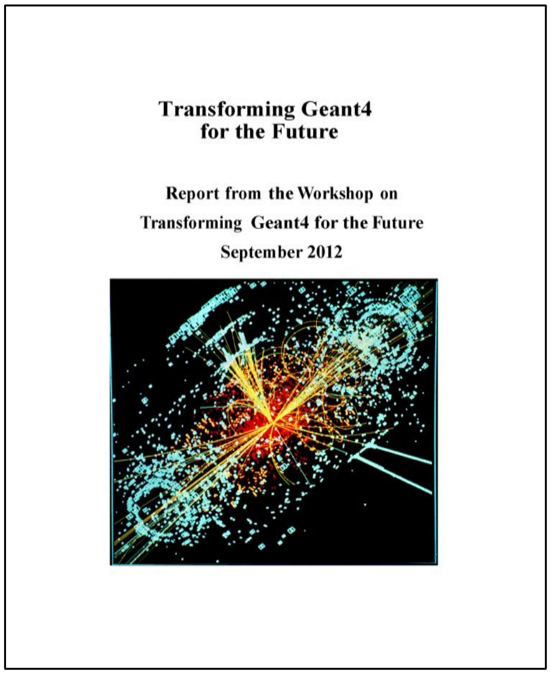
Transforming Geant4
for the Future

Report from the Workshop on
Transforming Geant4 for the Future
September 2012

- **Results**
  - Code profiling to uncover potential optimization opportunities
    - 1.5% single processor improvements by aggressive code transformations
    - Greater performance improvements will require extensive code re-writing
  - Cross-section code can improve memory locality with software caching
    - Preliminary analysis revealed significant performance improvement potential
    - Analysis of space-performance trade-off underway

Next Generation Workload Management and Analysis System for Scientific Data

DOE ASCR and HEP funded project funded 2012-2015

- Generalization of PanDA Workload Management System as meta application, providing location transparency of processing and data management, for HEP and other data-intensive sciences, and a wider exascale community
  - Pis: Alexei Klimentov*, Sergey Panitkin, Torre Wenaus and Dantong Yu, Brookhaven National Laboratory; Kaushik De, Gergely Zaruba, University, Texas at Arlington; Alexandre Vaniachine, Argonne National Laboratory

PanDA as workload manager

PanDA automatically chooses job execution site

- Multi-level decision tree – task brokerage, job brokerage, dispatcher
- Also predictive workflows – like PD2P (PanDA Dynamic Data Placement)

Site selection is based on processing and storage requirements

Why not use network information in this decision?

Can we go even further – network provisioning?

Network knowledge useful for all phases of job cycle

➢ Workload partitioning for traditional and opportunistic resources

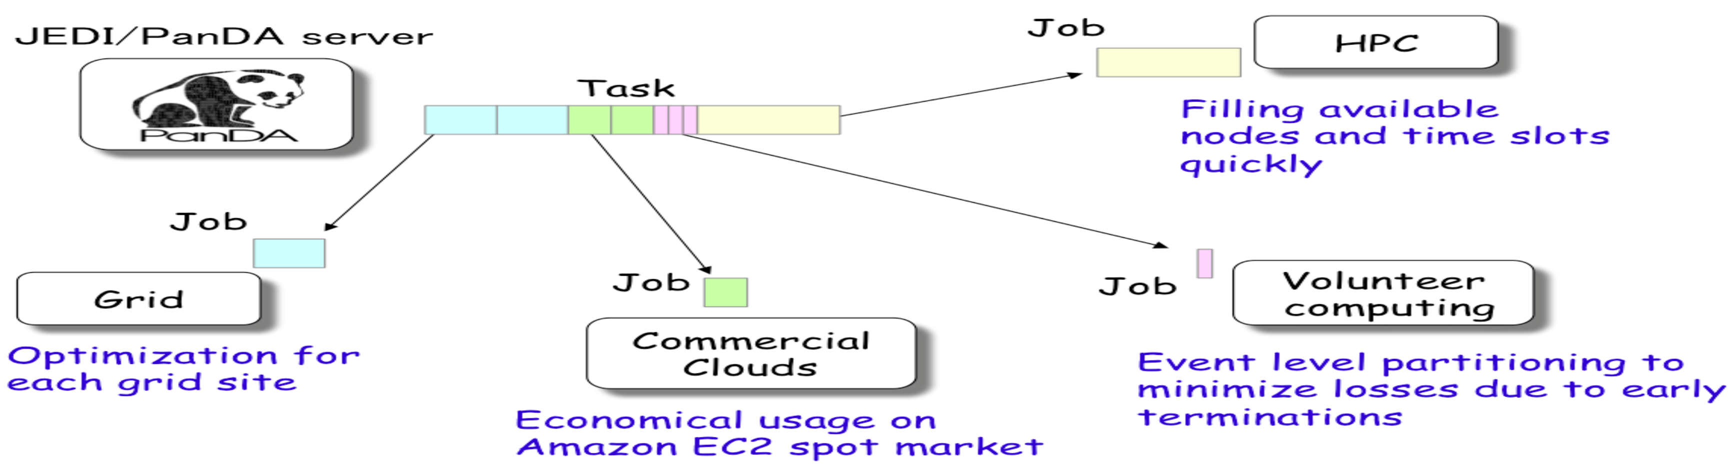

Recent Partnerships

SDN-Enabled Terabits Optical Networks for Extreme-Scale Science

- **ASCR 2015 Funding Opportunity Announcement from NGNS, closed July 2, 2015**
  – to develop a new generation of intelligent terabit optical network capabilities that will enable new modalities of network-intensive science.
  – To address increasing network complexity with intelligent and automated technologies that simplify network usage for scientists, automate complex network control and management functions, and improve the productivity and performance of network-intensive science.

SENSE: SDN for End-to-End Networks for Extreme-Scale Science

Inder Monga [Lead-PI] (ESnet), Phil Demar (FNAL), Harvey Newman (Caltech),
Linda Winkler (ANL), Tom Lehman (UMD/MAX), Damain Hazen (NERSC)

----- Mar 2016 – Feb 2019 -----

(SCR/HEP Jointly Funded)

Kickoff Meeting held at FNAL on March 16-17, 2016

Goal

- **Leverage Software Defined Network (SDN) to simplify federated network complexities with smart services to enable a new generation of network-aware data-intensive applications.**

Impacts

- **Present geographically distributed scientific resources to users as subsystems of a local facility,**
- **Simplifies the distribution complex datasets over federated autonomous multi-domains networks,**
- **Initiate the foundation of near real-time distributed computing and data analytics.**

SENSE SDN Control Plane Architecture for End-to-End Orchestration

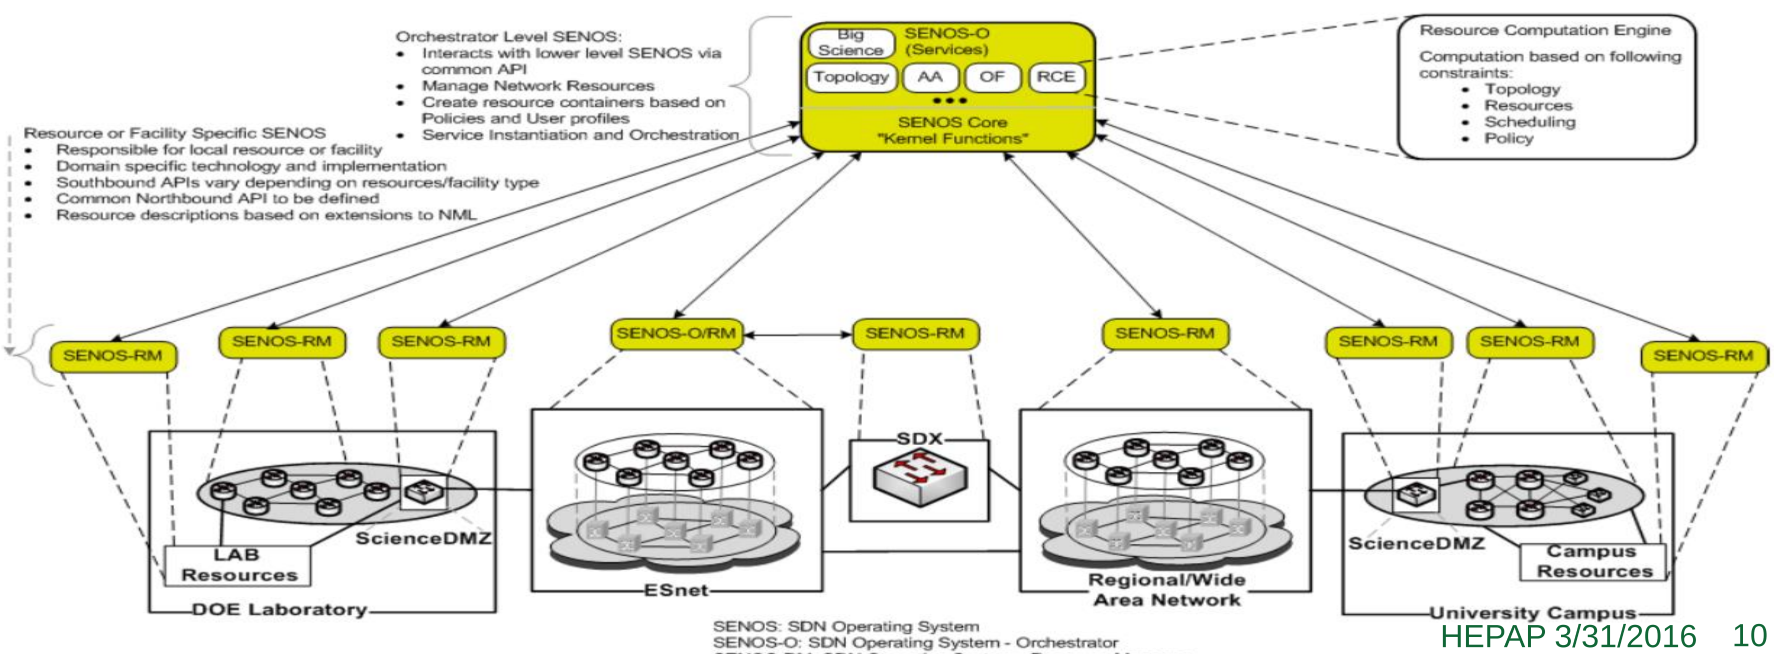

Figure 1. SENOS End-to-End Orchestration

ASCR-HEP - International Pilot Project

- **International collaborative project involving DOE and NSF and potentially agencies from Italy, France, Germany, UK and Canada (Planning Stage)**
- **ASCR – HEP partnership to build infrastructure for large international experiments such as the ATLAS and CMS particle physics experiments associated with the LHC, and the cosmic frontier surveys including DESI and LSST.**
- **Topic 1. Innovative Workflow Solutions.**
  - To address the needs of large scale scientific collaborations associated with major international facilities
  - Enable computing, supporting software, data management and access, networking, and people .
- **Topic 2. High-Performance Data Analysis and Algorithms.**
  - To anticipate and potentially overcome the scientific data analysis challenges that impede the discovery reach of large-scale, collaborative science programs.
  - Focus on real-time data analysis techniques, scalable and adaptive algorithmic approaches, and advances that foster the coupled analysis of scientific data from experiments, observations, and simulations.
- **Status**
  - Discussions with Italy, France, Germany, UK and Canada to produce the final draft of the solicitation
  - Hold international workshop in Europe during summer 2016.
  - Release solicitation during fall 2016.

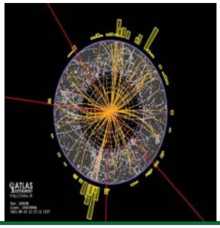

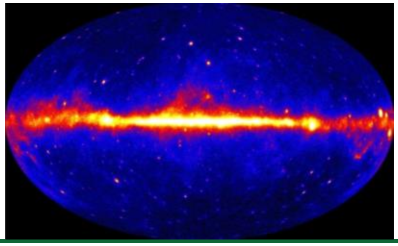

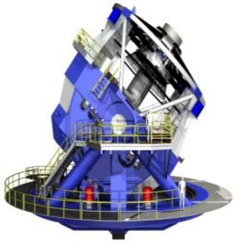

ASCR's Facilities

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 5,000 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
    - Open to national and international researchers, including industry
    - Not constrained by existing DOE or Office of Science funding or topic areas
    - Allocations based on rigorous peer and computational reviews
  - Approximately 300 users and 25-30 projects at each center, each year

Linking it all together – Energy Sciences Network (ESnet)

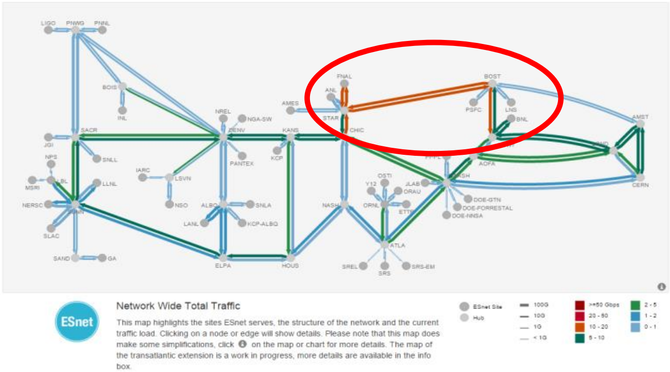

Path to the Future – Research & Evaluation Prototypes

- Includes Computational Science Graduate Fellows – to meet future workforce needs

OHEP and ESnet: 30 years of collaboration.

1. "What we can do on **LANs** today is indicative of what we wish to be able to do on **wide area networks**."

2. "Just as we expect a **computer** to perform as if we are the **only user**, we expect the **network** to give that same appearance."

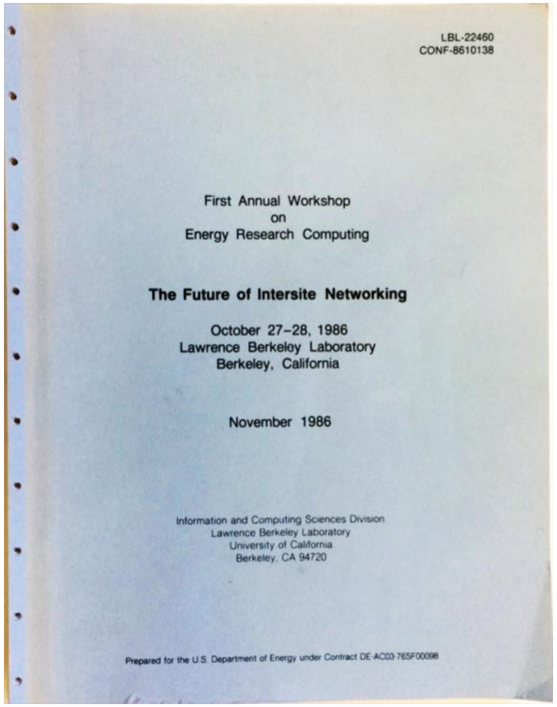
LBL-22460
CONF-8610138

First Annual Workshop
on
Energy Research Computing

The Future of Intersite Networking

October 27–28, 1986
Lawrence Berkeley Laboratory
Berkeley, California

November 1986

Information and Computing Sciences Division
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

1st ANNUAL WORKSHOP ON ENERGY RESEARCH COMPUTING

Harvey B. Newman
256-48 HEP
California Inst. Tech.
Physics Dept.
Pasadena, CA 91125
(818) 356-6656
NEWMAN@CITHEX.BITNET,
NEWMAN@CITHEX.CALTECH.EDU,

B. J. Helland
Ames Laboratory
Iowa State University
Ames, IA 50011
(515) 294-3086
HELLAND@ALISUVAX.BITNET,
HELLAND@ISUL.MFENET,

William Johnston
Bldg 50B-3238
Lawrence Berkeley Laboratory
1 Cyclotron Road
Berkeley, CA 94720
WEJOHNSTON@LBL.ARPA,

Stephen Wolff
Room 533
National Science Foundation
1800 G Street, N.W.
Washington, D.C. 20550
(202) 357-9717
STEVE@BRL.ARPA,

ESnet Goes Global: Extension to Europe

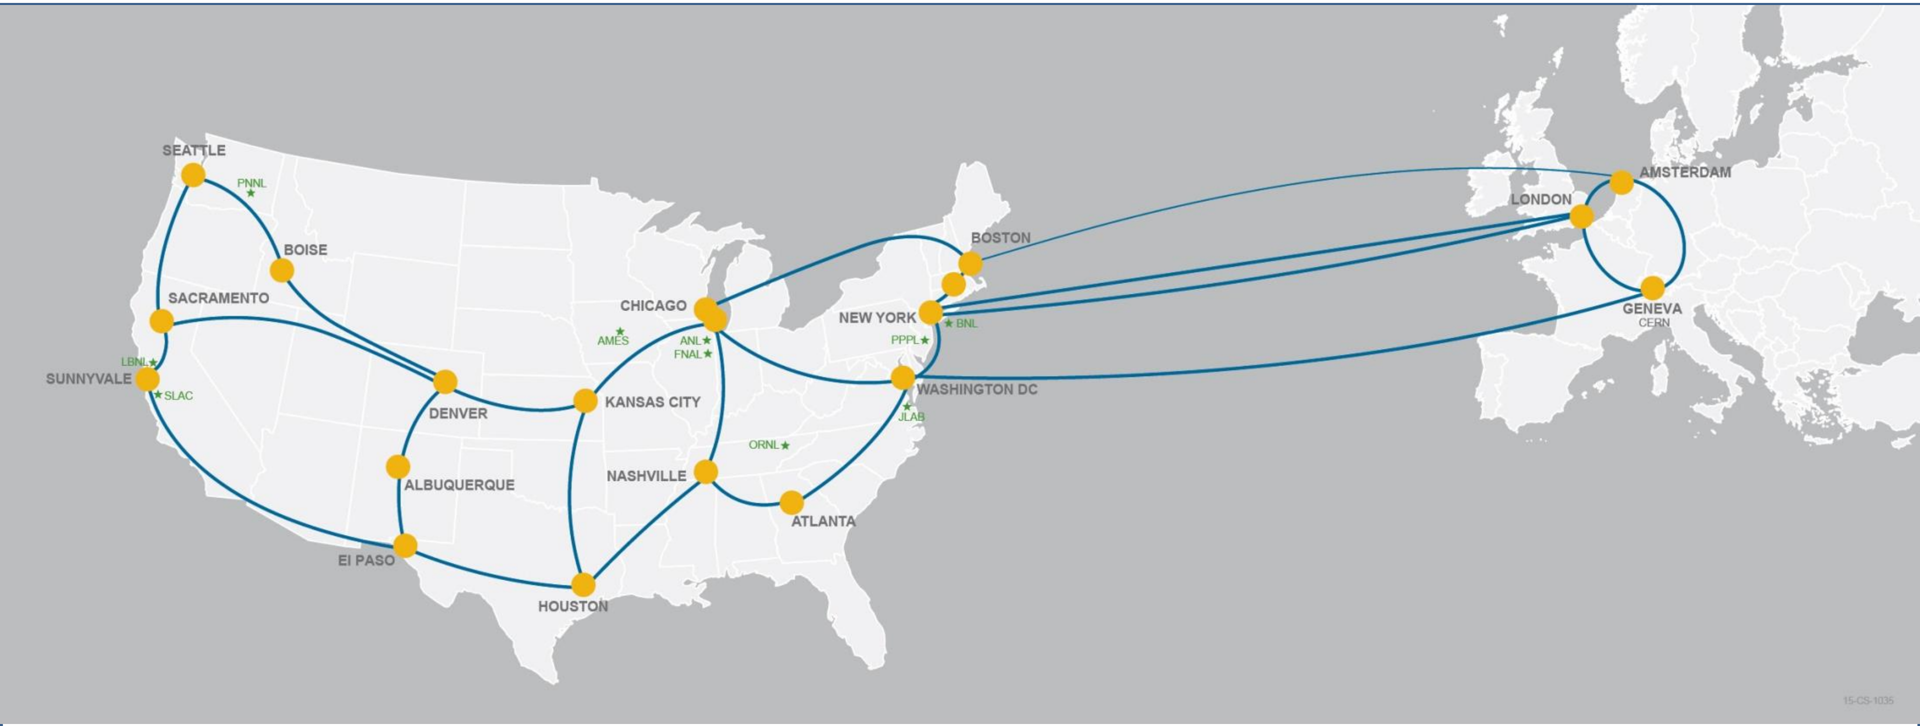

- 25% of all DOE traffic goes to/from Europe
- 3x100+ Gbps across the Atlantic with redundant paths to serve all DOE missions
- 10x increase over current transatlantic network for Large Hadron Collider (LHC) experiment
- Was Ready in January 2015 to support LHC Run 2

Increased Esnet Traffic

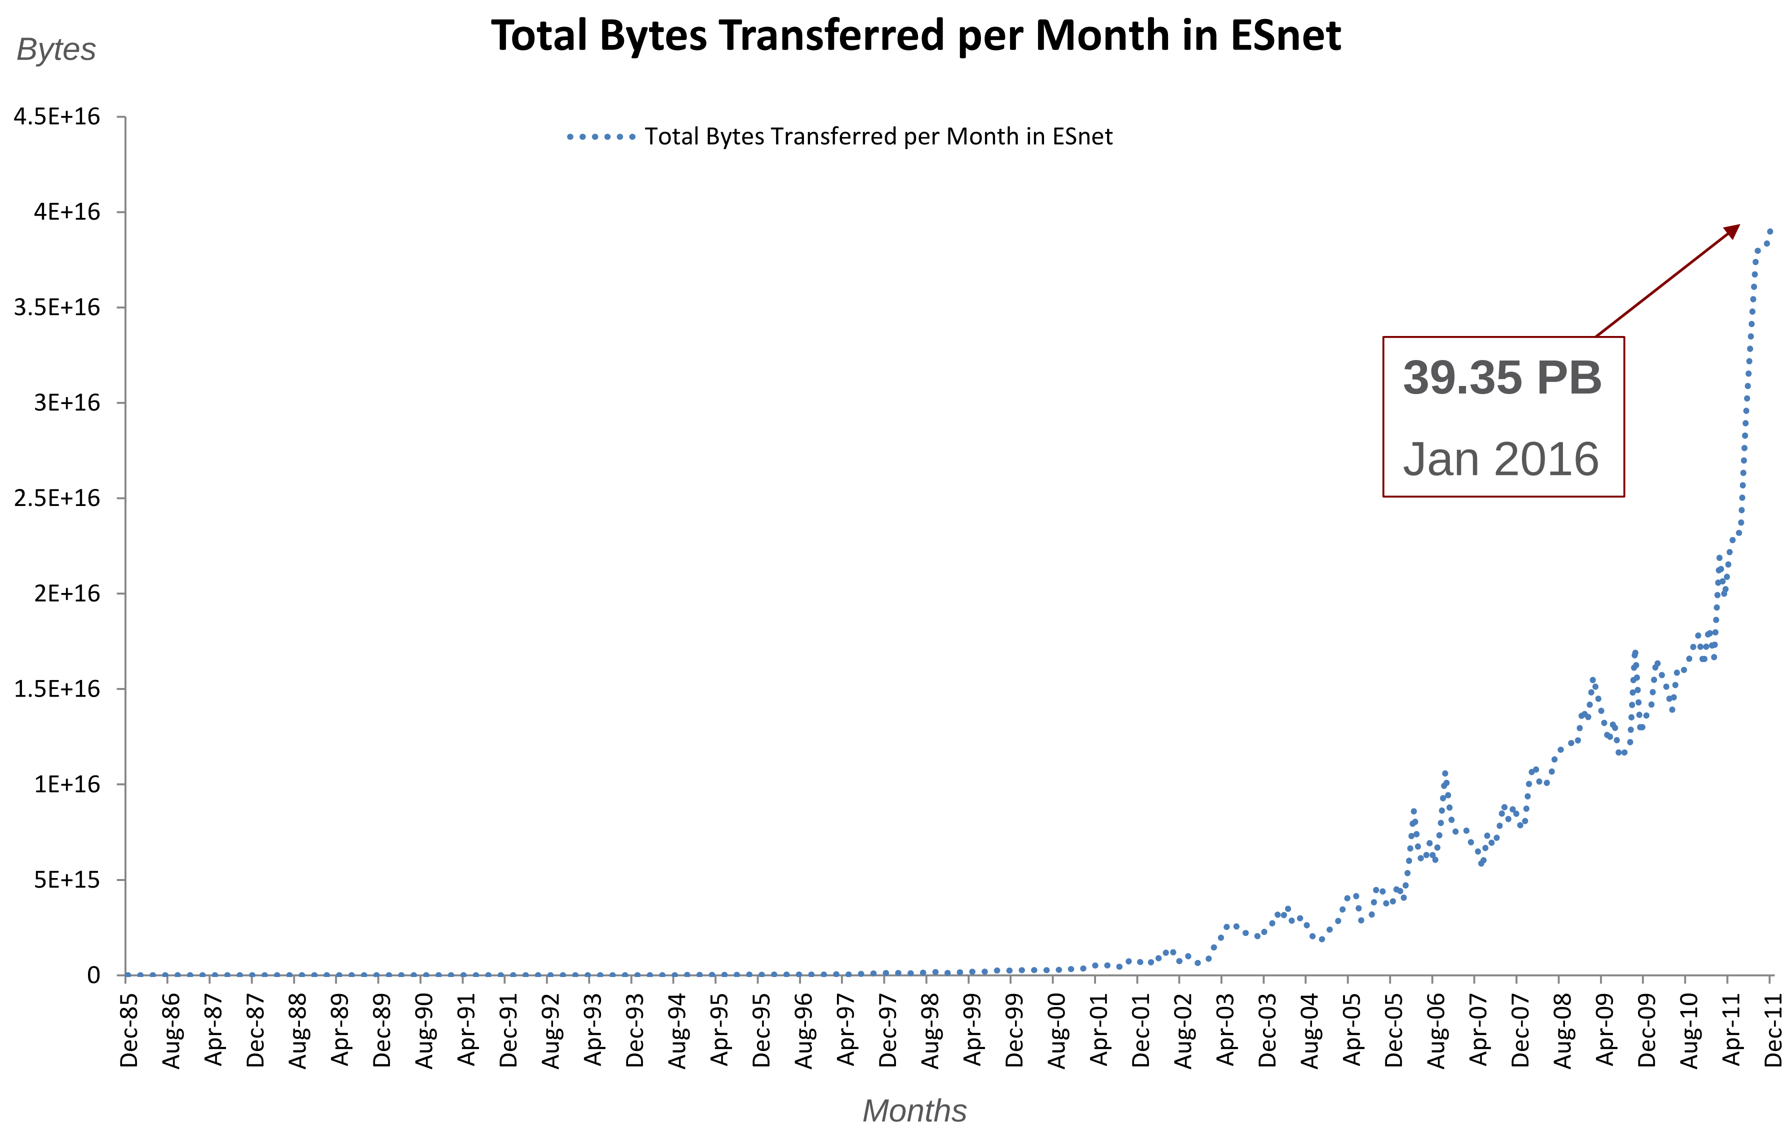

Components of ASCR's Allocation Policy

- **Facility directors will retain up to 10% of allocable hours to support pilot or startup projects, to support code scaling and for petascale computer science and performance metrics research.**
- **The majority of (60-85%) of available processor hours at NERSC will be for researchers working on SC-funded or SC relevant projects.**
  – NERSC issues Call for Proposals and Headquarters Program Management determines allocations
- **The majority (60-85%) of available processor hours on the Leadership Computing resources will be allocated through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program.**
  – LCF Directors will be responsible for conducting joint reviews and select projects for their facility.
- **ASCR retains between 5-30% of resources at NERSC, ALCF and OLCF for the ASCR Leadership Computing Challenge**
  – ASCR issues Open Call for proposals, conducts reviews and selects projects
  – For projects critical to DOE Missions and to broaden HPC/LCF community

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**

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	200	>8.5	180
Peak Power (MW)	2	9	4.8	< 3.7	13.3	1.7	13
Total system memory	357 TB	710TB	768TB	**~1 PB DDR4 + High Bandwidth Memory (HBM)+1.5PB persistent memory**	**> 2.4 PB DDR4 + HBM + 3.7 PB persistent memory**	**676 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	~4,600 nodes	>3,200 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®

PanDA Continued: Compute & Data Integrated Complex

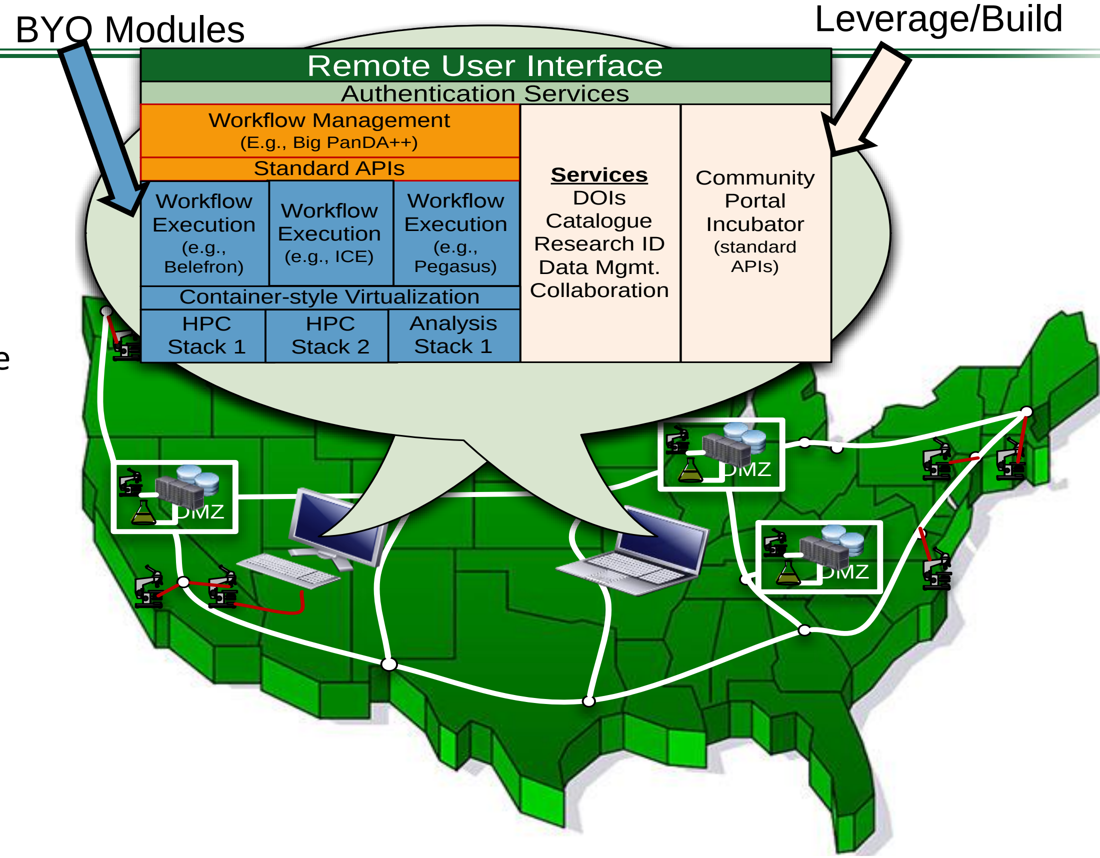

Diverse set of compute & data resources co-located with DOE/SC experimental facilities

- Take advantage of locality to minimize data movement and latency
- Co-location of experimental and computing staff to build partnerships and advance innovation
- Each facility can customize workflows to meet the needs of the experimental and computational users

Built on top of a common remote authentication & execution framework

- Provides common APIs across centers to enable users to build additional services
- Enables complex workflows across centers
- Geographically dispersed infrastructure and data replication for critical services

Vision: Universal workload management system for distributed workflows

- Demonstrate ASCR “facilities of the future” integration with experimental science facilities
- Demonstrate at scale on Titan for HEP applications (2016-2017)
- Set up PanDA instance on OLCF in 2017
- Demonstrate non-HEP application use cases
  - Candidates: molecular dynamics, bioinformatics, and neuroscience
- Demonstrate multiple applications on PanDA@Titan (2017-2018)
- Refactor PanDA for workflows and usage patterns identified from user requirements
- Ensure PANDA is interoperable with existing/evolving workflow systems

Shifter for HPC & HEP

- **NERSC in partnership with the *Forum for Computational Excellence* has brought Docker-like container technology to its HPC systems through a new software package known as Shifter on Cori and Edison.**
- **Vast majority of LHC computing depends on CVMFS – previous efforts were challenging on HPC.**
- **Using shifter delivered full CVMFS software stack to Cori and Edison for ALICE, ATLAS, and CMS and avoided tarring up chunks of software and copying them to every node.**
- **Delivers > 1 TB software stack to 1000's of nodes, negligible initialization time and creates an environment that "just works".**

SHIFTER

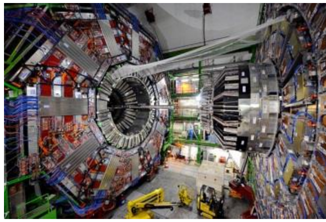

***ExDAC* has enabled Shifter technology in the Cosmic Frontier for DES and LSST-DESC and in partnership with the FCE at ANL has provided this capability to MicroBooNE and SPT.**

μBooNE

SOUTH POLE TELESCOPE

DARK ENERGY SURVEY

LSST DESC Dark Energy Science Collaboration

High Energy Physics

Elementary particles and fields

Director's Discretionary Program
Sergey Panitkin, BNL
Big PanDA Project
Allocated hours: 500,000

Science Objectives and Impact

- Investigate properties of Higgs boson and thereby the origin of mass
- Integrate scientific workflow management from the ATLAS experiment at the Large Hadron Collider on OLCF for big data availability, processing, relocation, and job coordination
- Explore elementary constituents of matter and understand their interactions at the most fundamental level

ATLAS event display of a candidate Higgs particle decay to two electron-pair event. This event was recorded by ATLAS on 18-May-2012, 20:28:11 CEST. The tracks and clusters of the two electron pairs are colored red and blue, respectively.

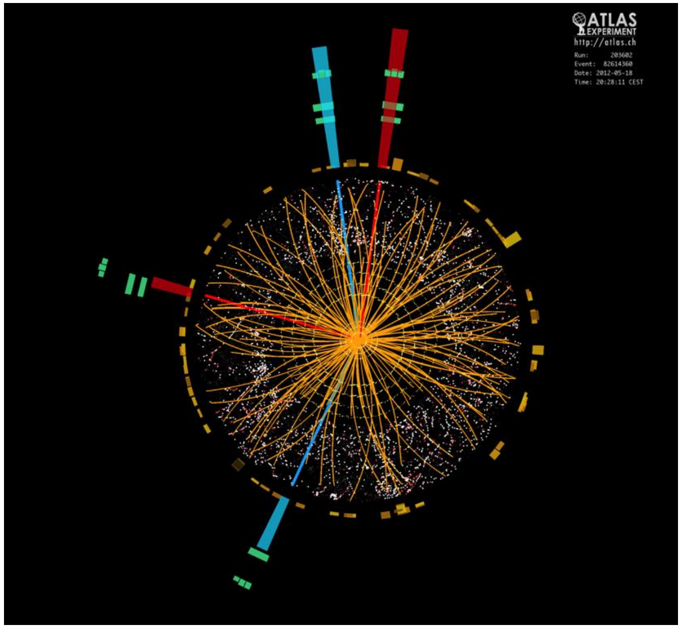

OLCF Contribution

- First implementation of PanDA (Production and Data Analysis) workflow management at a Leadership Computing Facility
- Demonstration of potential for efficient PanDA backfill of draining queues in preparation for large leadership-class jobs

Technical Achievements

- **Demonstration of efficient simulations (e.g., Geant4) and analyses essential for estimation of backgrounds, efficiencies, and systematic errors associated with experimental LHC data**
- **Integrated implementation conforming to OLCF's operational best practices and infrastructure**

ATLAS Collaboration, "Observation of a new particle in the search for a Standard Model Higgs boson with the ATLAS detector at the LHC," Phys. Lett. B716 (2012).

Select ALCF & HEP Collaborations

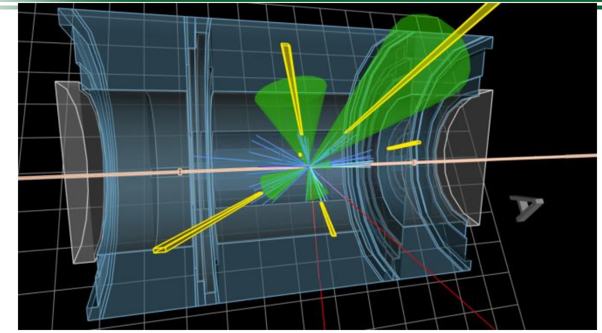

Simulation of LHC Events

Tom LeCompte, ANL

Adapted grid workflow and improved Alpgen performance for Mira. Working on scaling and performance of Sherpa.

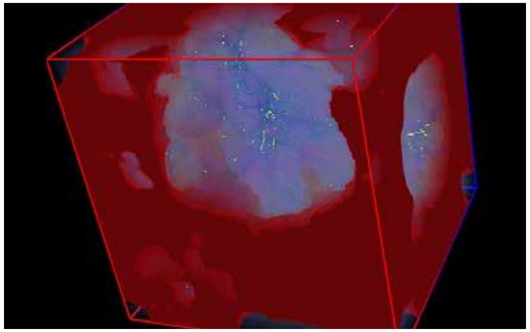

Cosmic Reionization

Nickolay Gnedin, FANL

Restructured I/O to enable capability scale runs

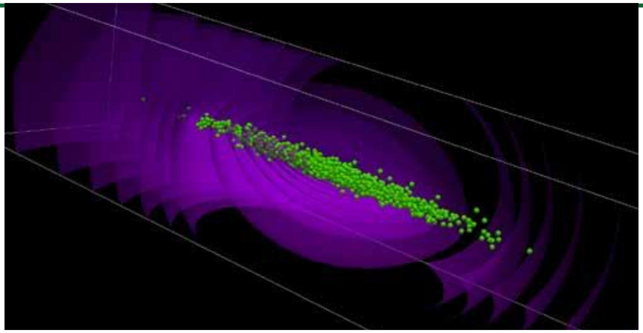

Design of Fermilab and Cern accelerators

James Amundson, FANL

Transitioning production code to new compiler base enabling much better applications feature support and performance

Cosmological Simulation for Large-Scale Sky Surveys

Salman Habib, ANL

Working on advancing physics, I/O, and analysis. Developing new capabilities, enhancing performance, designing and testing future platforms.

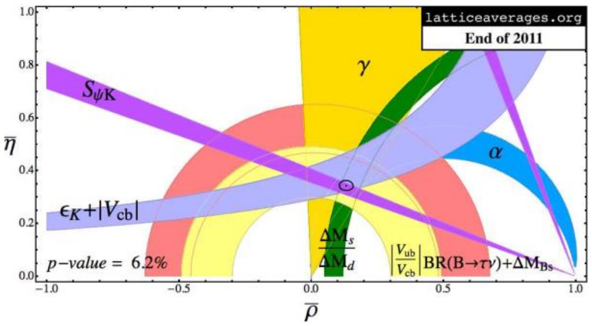

USQCD

Multiple INCITE & ALCC awards - Mackenzie, Kuti, & Orginos

Active member of USQCD, co-PI on SciDAC. Porting of communications layer to Blue Gene, optimization of code base. Developing new lattice framework targeting DSL for easier, well optimized code

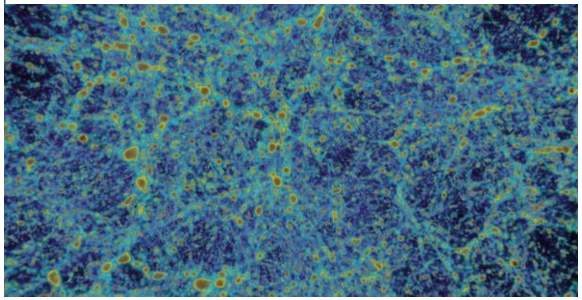

Astronomers Observe a Supernova Colliding with its Companion Star

Scientific Achievement

NERSC machine-learning data analysis pipeline allowed the first ever observation of light from a supernova slamming into a neighboring star.

Significance and Impact

Although thousands of Type Ia supernovae have been found, this kind of explosion had been described but never before observed.

The intermediate Palomar Transient Factory (iPTF) analysis pipeline found the light signal from the supernova just hours after it ignited in a galaxy about 300M light years away from Earth. iPTF depends on NERSC computation and global storage resources.

Nature | Vol521 | May 21, 2015
The Atrophysical Journal 708:1025–1031, 2010 January 10

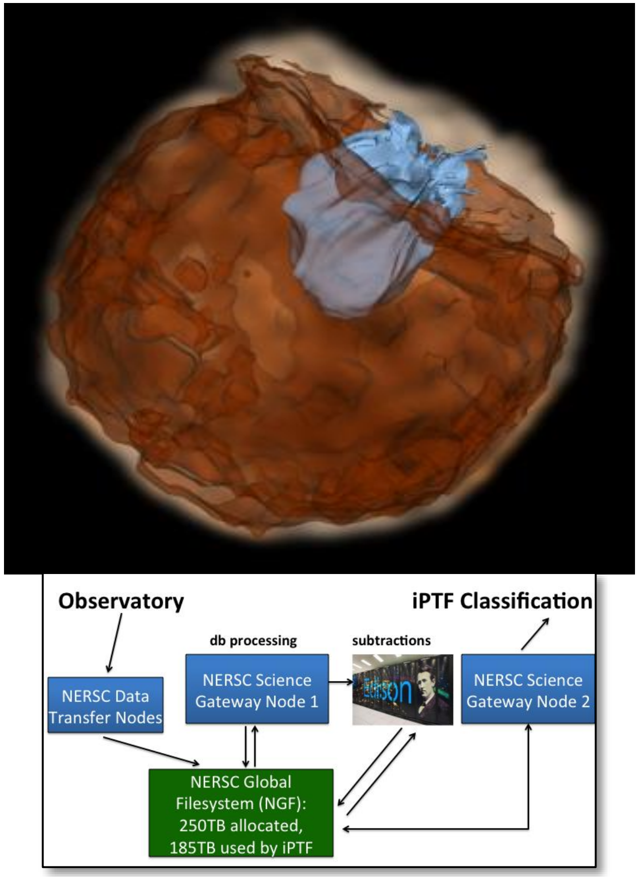

Simulation of the expanding debris from a supernova explosion (shown in red) running over and shredding a nearby star (shown in blue). Image credit: Daniel Kasen, Berkeley Lab/ UC Berkeley. Inset shows NERSC contribution to iPTF processing pipeline

What’s Next?

National Strategic Computing Initiative

Executive Order Signed July 29, 2015

Executive departments, agencies, and offices participating in the NSCI shall pursue five strategic objectives (whole of government approach):

1) **Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.**
2) **Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing**.
3) Establishing over the next 15 years, a viable path forward for future HPC systems in the Post-Moore's-Law Era to advance beyond traditional lithographic scaling of devices.
4) **Increasing the capacity and capability of an enduring national HPC ecosystem, employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, and workforce development.**
5) **Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the U.S. commercial, government, and academic sectors.**

NSCI Agency Responsibilities

There are three lead agencies for NSCI: the Department of Energy (DOE), the Department of Defense (DOD), and the National Science Foundation (NSF).

- The DOE Office of Science and DOE National Nuclear Security Administration will execute a joint program focused on advanced simulation through a capable exascale computing program emphasizing sustained performance on mission relevant applications and on analytic computing to support its missions and post-Moore's Law HPC capability.
- NSF will play a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development.
- DOD will focus on data analytic computing to support its mission.

These responsibilities leverage the historical roles each of the lead agencies have played in pushing the frontiers of high-performance computing, and will keep the nation on the forefront of this strategically important field. The lead agencies will also work with the foundational research and development agencies and the deployment agencies to support the objectives of the NSCI to address the wide variety of needs across the Federal Government.

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 mid 2020s**

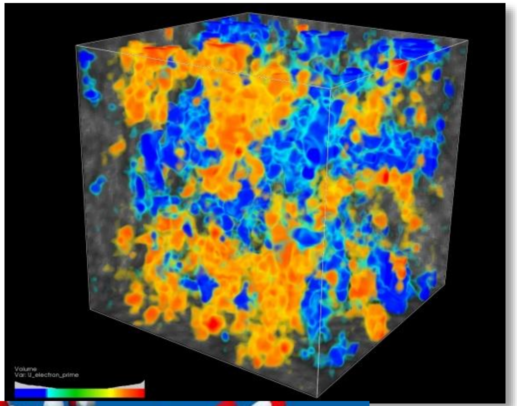

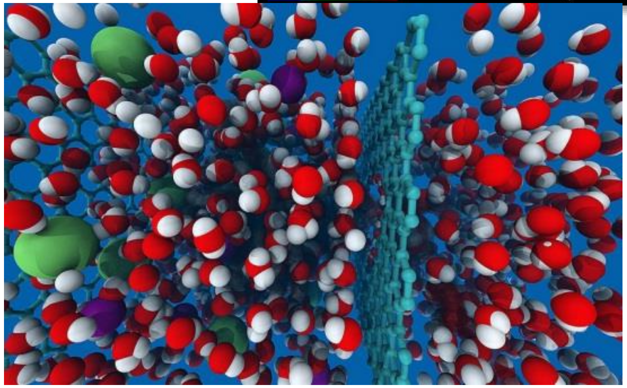

Challenges

- **Increasing Performance / Value**
- **Efficiency**
  - Energy efficiency: reduce energy per operation (pJ/op)
  - Hardware efficiency: massively parallel architectures, down to the processor level
  - Software efficiency: effectively exploit H/W parallelism
    - Current efficiencies on conventional machines are < 10% for many real-world applications
- **Memory / Storage**
  - Make effective use of data movement (this is the dominant energy cost)
- **Reliability**
  - Successfully complete execution through system failures
- **Productivity**
  - Programming environment that makes HPC machines accessible to everyone
  - Reduce time to solution
- **Cost / Affordability**

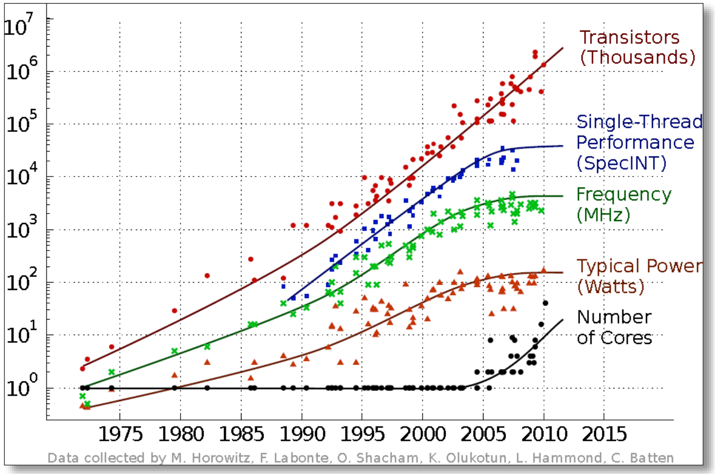

From Giga to Exa, via Tera & Peta

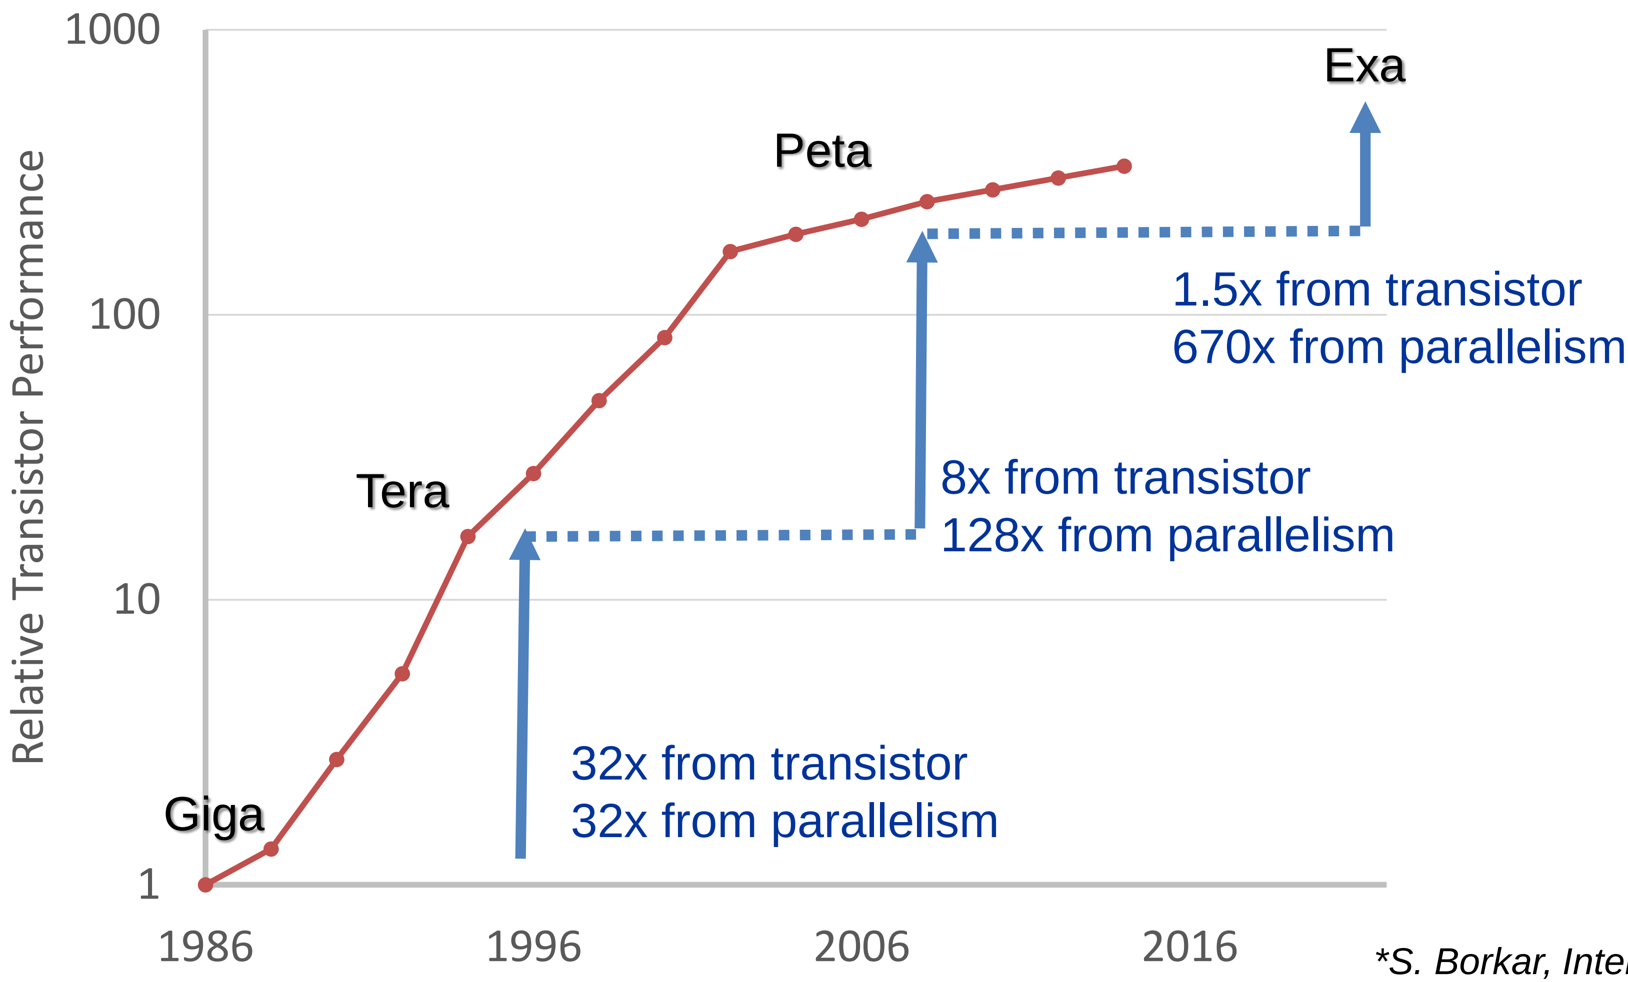

*S. Borkar, Intel

1. **Performance from parallelism—not just node performance**
2. **Data movement will increase**

Specialization

- **Integrating current and future processing technology into the computing fabric – heterogeneous processing**

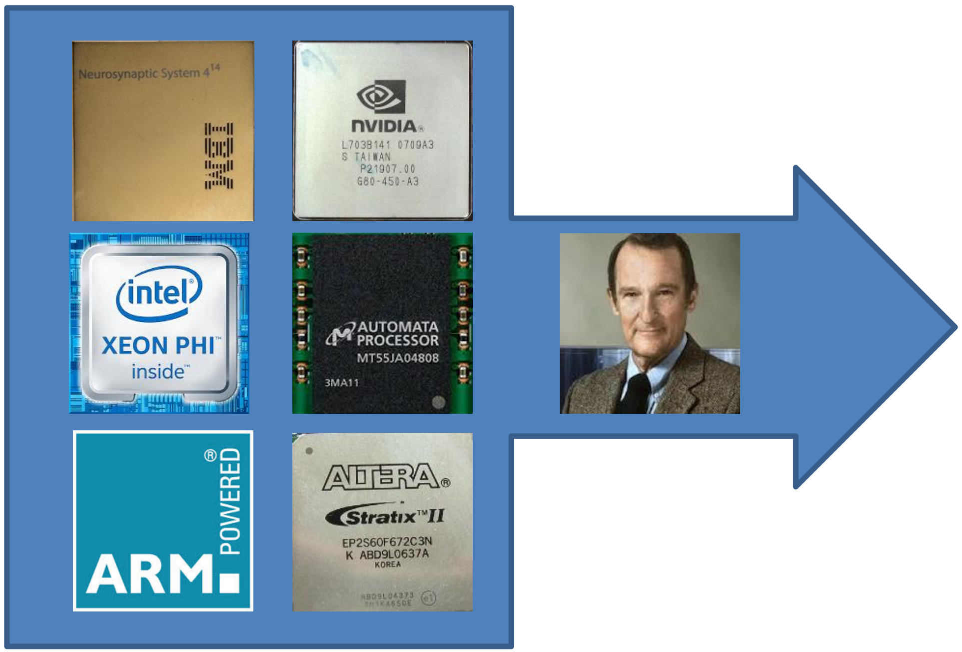

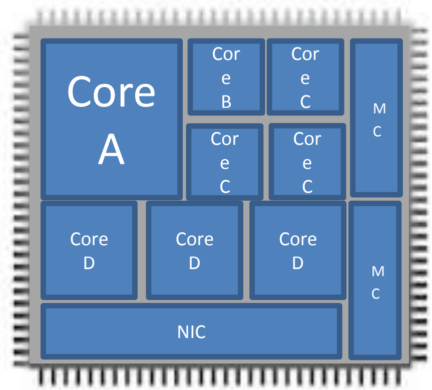

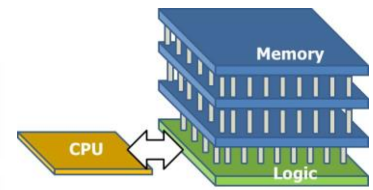

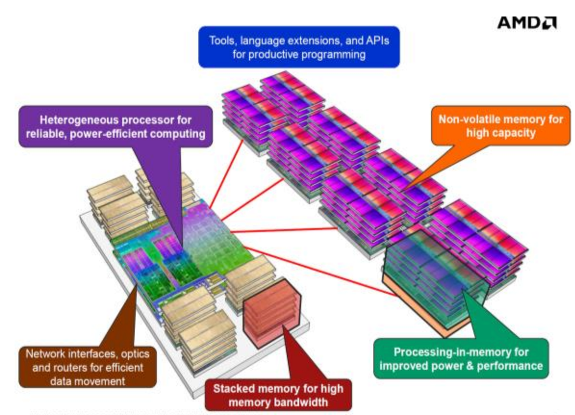

U.S. DEPARTMENT OF ENERGY | Office of Science

ECI Strategy

- SC & NNSA Partnership
- Develop next generation of applications
- Exploit co-design process, driven by the full application workflow
- Develop exascale software stacks
- Partner with and fund vendors to transition research to product space
- Integrate applications, acquisitions, and research and development
- Collaborate with other government agencies and other countries, as advantageous

Exascale Computing Project

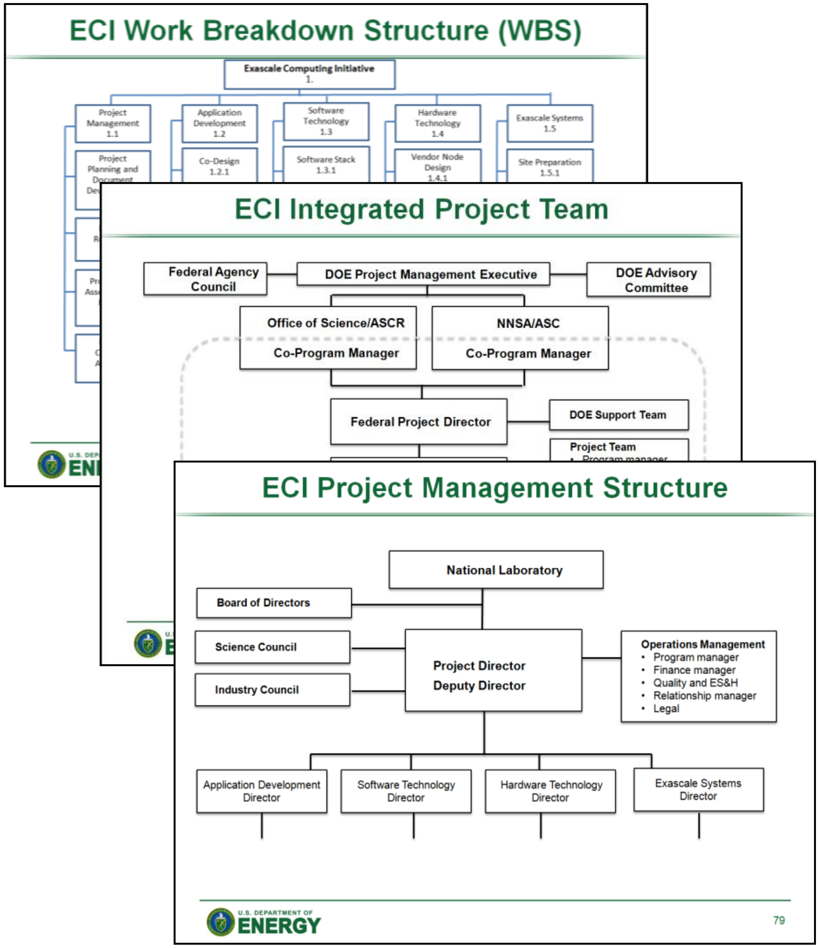

- The exascale project will follow established DOE review and decision protocols for its execution
- A project office has been established at ORNL with representation from the major participating laboratories (ANL, LANL, LBNL, LLNL, ORNL, SNL)
- An Integrated Project Team (IPT) has been established, 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
- **A new Exascale Computing subprogram has been established in ASCR and includes only those activities required for the delivery of exascale computers**

ECP Timeline

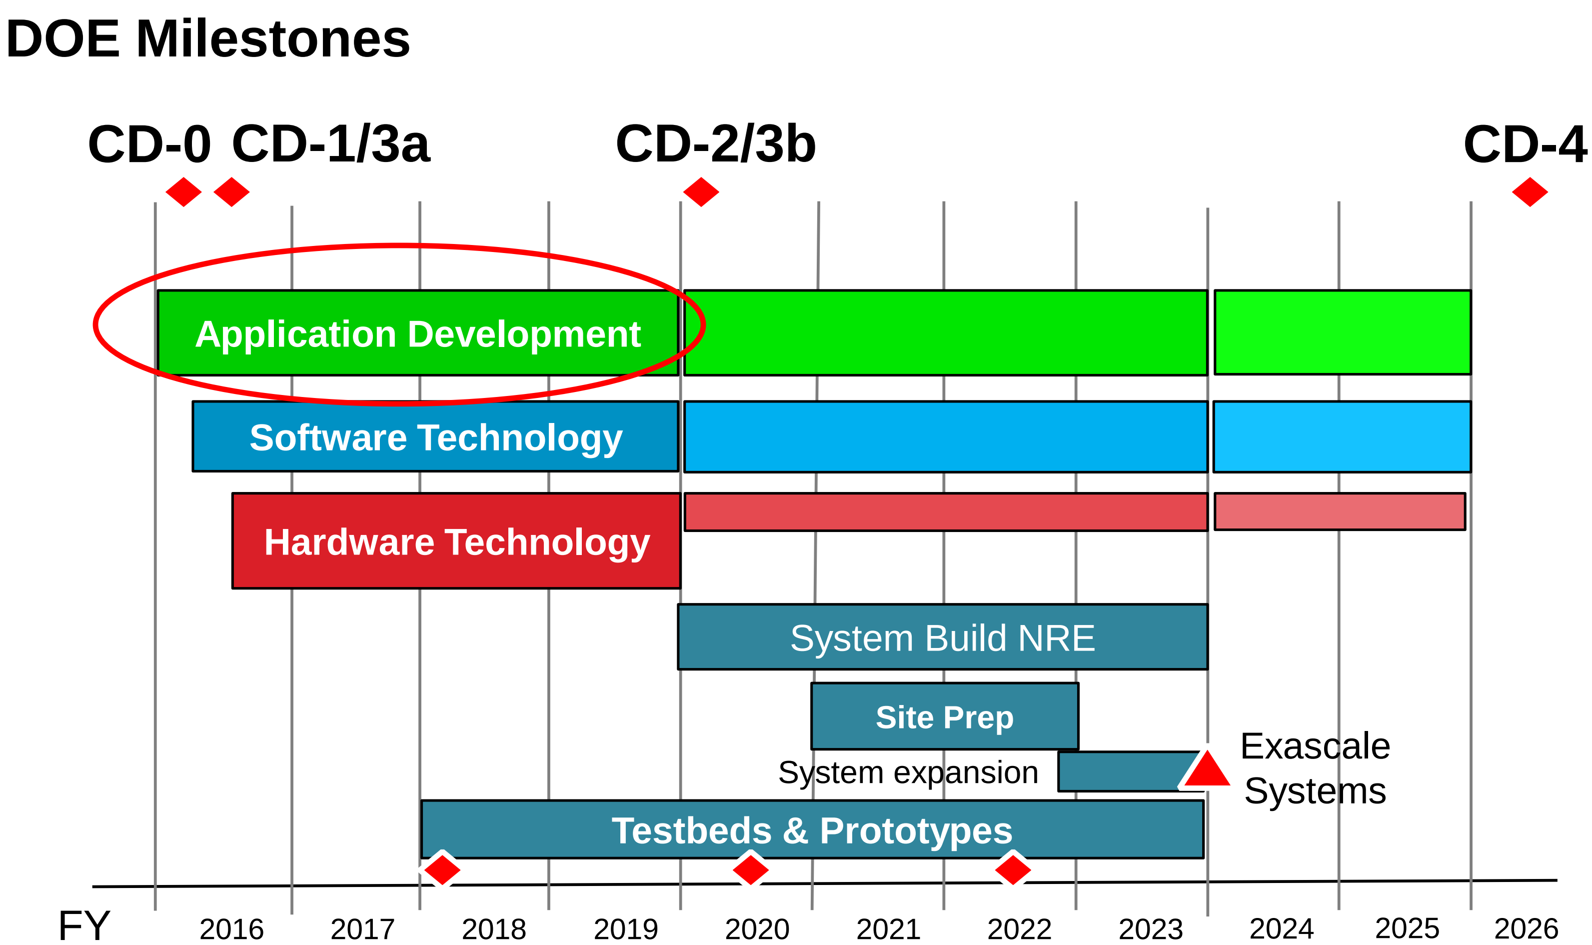

Exascale Application Needs Gathering Activities

- **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 --- **135 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 --- **114 responses received**

DOE National Laboratories Responses: 135
Academic Responses: 94
Industry Responses: 8
Foreign Response: 2
Others: 5

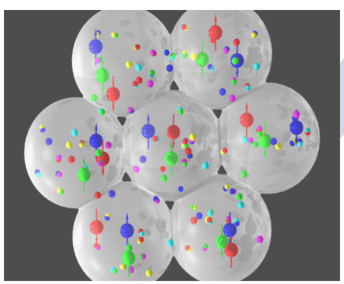

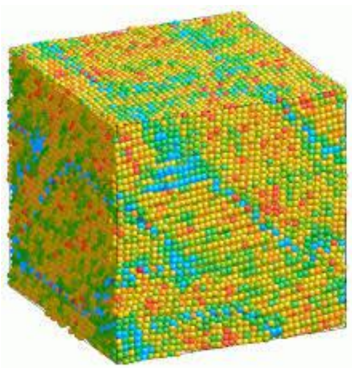

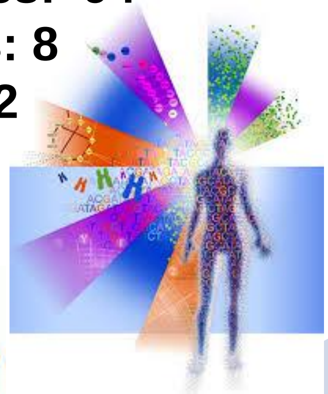

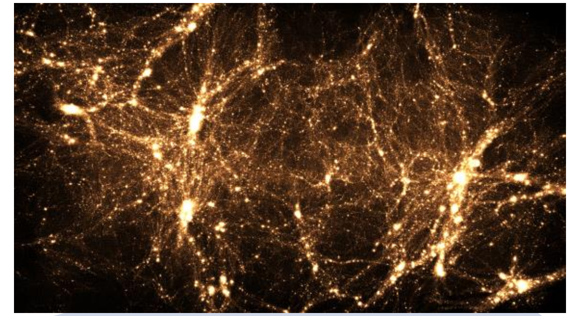

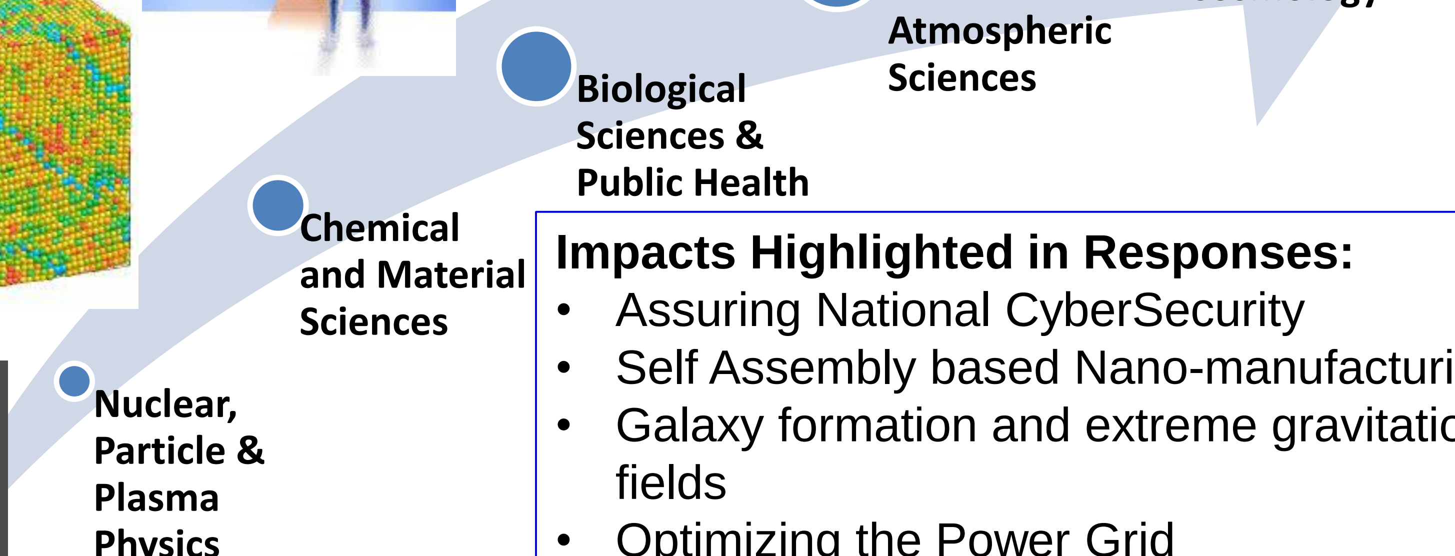

Impacts Highlighted in Responses:

- Assuring National CyberSecurity
- Self Assembly based Nano-manufacturing
- Galaxy formation and extreme gravitational fields
- Optimizing the Power Grid
- Realistic Hypersonic flow for Flight Vehicles
- High Resolution Atmospheric & Climate Models
- N-by-N Comparison of All Patients in US
- Time-sensitive simulations of global scale epidemics

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
- **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.

Proposed Schedule

June 10-12,2015	HEP
November 3-5 2015	BES
January 27-29, 2016	FES
March 29-31, 2016	BER
June 15-17 2016	NP
Target: Sept. 2016	ASCR

ASCR – HEP Requirements Review Findings and Observations

- **The HEP community has historically leveraged HPC capabilities primarily to address theory and modeling tasks. Energy Frontier-related experimental efforts are now also beginning to employ HPC systems. The estimated computational needs for these experiments far exceed the expected HEP investment in computational hardware that will be required to execute the envisioned HEP science program by 2025.**

- **The 2025 timeline will define a phase transition in terms of HEP data flows. Three new facilities will come online and the landscape will change dramatically. The high-luminosity LHC will be operating, producing data samples at unprecedented scales. Fermi National Accelerator Laboratory's (Fermi's) flagship Deep Underground Neutrino Experiment (DUNE) will be operating, as will the next-generation dark energy survey experiment, the Large Synoptic Sky Telescope (LSST). Data sizes produced each year could be 200 times greater than what is being produced by today's operating experiments. In addition, these new experiments will each require a simulation program that dwarfs what we are doing today in order to take advantage of the expected improvement in statistical precision."**

Exascale Requirements Review for High Energy Physics draft final report, March 2016, Salman Habib, Robert Roser, co-chairs

HEP Requirements Review High-Level Key Finding

Long-range planning between HEP and ASCR will be required to meet HEP's research needs. To best use ASCR HPC resources, the experimental HEP program needs (1) an established, long-term plan for access to ASCR computational and data resources, (2) the ability to map workflows to HPC resources, (3) the ability for ASCR facilities to accommodate workflows run by collaborations potentially comprising thousands of individual members, (4) to transition codes to the next-generation HPC platforms that will be available at ASCR facilities, (5) to build up and train a workforce capable of developing and using simulations and analysis to support HEP scientific research on next-generation systems.

Exascale Requirements Review for High Energy Physics draft final report, March 2016, Salman Habib, Robert Roser, co-chairs

http://arxiv.org/abs/1603.09303

Beyond Moore's Law: Opportunities for New Partnerships

NSCI Objectives

Executive departments, agencies, and offices participating in the NSCI shall pursue five strategic objectives:

1) Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
3) **Establishing over the next 15 years, a viable path forward for future HPC systems in the Post-Moore's-Law Era to advance beyond traditional lithographic scaling of devices.**
4) Increasing the capacity and capability of an enduring national HPC ecosystem, employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, and workforce development.
5) Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the U.S. commercial, government, and academic sectors.

Post-Moore's Law Computing: What comes after exascale?

- CMOS lithographic feature sizes are approaching fundamental limits
  - Currently at 22 nm (Intel Knight's Corner)
  - 14 nm is projected for ~2016 (Intel Knight's Landing)
  - 10 nm is under development (Intel, Nvidia)
    - However, gate lengths may be smaller than 6 nm – corresponding gate dielectric thickness may reach a monolayer or less
  - The Intel roadmap reaches beyond 10 nm (7 nm and 5 nm) but may be unattainable
    - Non-silicon extensions of CMOS, e.g., using III-V materials or nanotubes/nanowires or non-CMOS technologies, including molecular electronics, spin-based computing, single-electron devices, and graphene have been proposed
    - At scales less than 10 nm, quantum tunneling is expected to become significant
  - Capital costs for tooling are increasing dramatically as feature sizes shrink
- Options:
  - Computing using superconducting technologies ← **IARPA historical lead**
  - Quantum computing
  - Neuromorphic computing
  - Probabilistic computing
  - ???

Considerable R&D required

NSTC QIS Interagency Working Group Activity

Chartered in October 2014

Membership: DOC, DOD, DOE, NSF, NSA, ODNI, OMB, OSTP

- **Survey of recent USG activity and investment**
- **Informal survey of international activity**
- **Workshop and RFI to solicit industry views on:**
  - Identify opportunities for research and development, emerging market areas, barriers to near-term and future applications, workforce needs
  - Approximately 77 workshop attendees (28 industry, 34 government, 14 university, foundation 1); 8 written responses to RFI
- **Initial recommendations presented to NSTC Physical Sciences Subcommittee – July 10, 2015**
- **Interim report provided to NSTC Committee on Science – December 2015**
- **Briefing to PCAST – January 14, 2016**
- **Public release of report is in preparation**

Quantum Information Science is Broad and Multidisciplinary

- **Spans multiple disciplines**
  - Computer science, applied mathematics, networking, information science
  - High-energy physics (advances in quantum theory)
  - Materials sciences (new materials)
- **Has many, conceivable applications**
  - Quantum simulation/computing
  - Sensors
  - Cryptography
  - Communications, networking
  - Metrology/measurement, accurate timekeeping, …
  - …
- **Quantum theory is evolving**
  - new understanding of entanglement and information theory
- **Practical quantum computers require new technologies**
  - New materials for quantum devices
  - qubit manipulation, error correction
  - System integration
  - Software
  - …

Quantum Computing Community Engagement

- Workshop on Materials Opportunities for Quantum Computing,* October 7-8, 2014, LANL
- **Grand Challenges at the Intersections of QIS, Particle Physics, and Computing, December 11, 2014, joint DOE/HEP & ASCR**
- NNSA Workshop on Applications of Quantum Computing, February 5-6, 2015, SNL (organized by LANL)
- **DOE/ASCR Workshop on Quantum Computing in Scientific Applications, Date: February 17-18, 2015**
- Workshop on Beyond Exascale: Qubits for Quantum Computing,* ORNL, August 20-21, 2015
- **Round Table on "Quantum Sensors at the Intersection of Fundamental Science, Advanced Computing, and QIS," February 25, 2016**
- Also … NSF and NIST community events (2015, 2016)

**Community organized*

Post-Moore R&D: ASCR's FY 2017 Request

- **Request: $10,000,000**
- **Challenge:**

  Begin research into computing technologies and associated applied mathematics and computer science to prepare for the generation of scientific computing beyond exascale.

- **FY 2017 Objective:**

  Initiate two activities, predicated on recent community engagement:

  - Development of two to three testbeds, which would support ASCR, BES, and HEP-based algorithm development activities (Research & Evaluation Prototypes)
  - Initiate R&D in quantum algorithms and quantum measurement, with the goal of developing technologies that could impact next-generation SC scientific facilities, across BES and HEP. The proposed activity is consistent with and coordinated with the broader, USG-wide, interagency QIS activity, specifically NIST and NSF (Partnerships)

And that is just the beginning.

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/