

SciDAC at NERSC

Scientific Discovery through
Advanced Computing
at the National Energy Research Scientific Computing Center

Horst D. Simon
Director, NERSC Center
Berkeley, California, USA

(SciDAC slides courtesy of Alan Laub, Director SciDAC, DOE)

August 6, 2003

Introduction – What is SciDAC?

- SciDAC is a pilot program for a “new way of doing science”
- first Federal program to support and enable “CSE” and (terascale) computational modeling and simulation as the third pillar of science (relevant to the DOE mission)
- spans the entire Office of Science (ASCR, BES, BER, FES, HENP)
- involves all DOE labs and many universities
- builds on 50 years of DOE leadership in computation and mathematical software (EISPACK, LINPACK, LAPACK, BLAS, etc.)

SciDAC Goals

- an INTEGRATED program to:
 - (1) create a new generation of **scientific simulation codes** that take full advantage of the extraordinary capabilities of terascale computers
 - (2) create the **mathematical and computing systems software** to enable scientific simulation codes to effectively and efficiently use terascale computers
 - (3) create a **collaboratory software environment** to enable geographically distributed scientists to work effectively together as a TEAM and to facilitate remote access, through appropriate hardware and middleware infrastructure, to both facilities and data

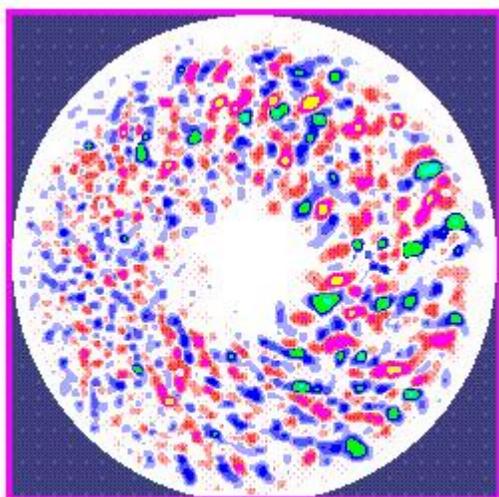
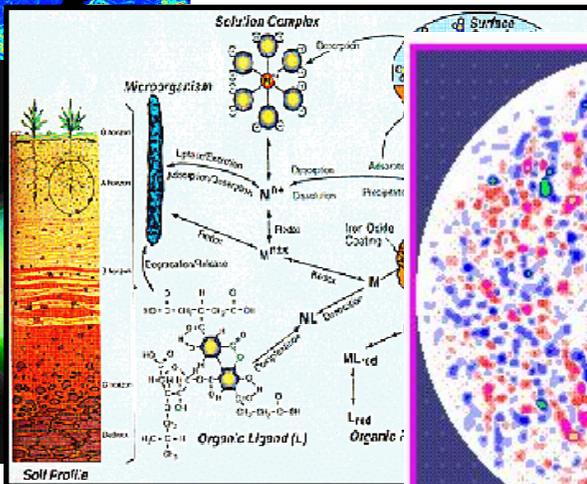
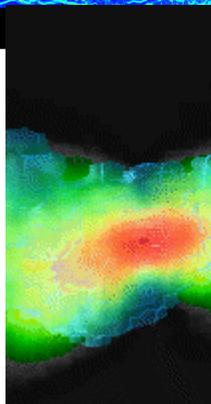
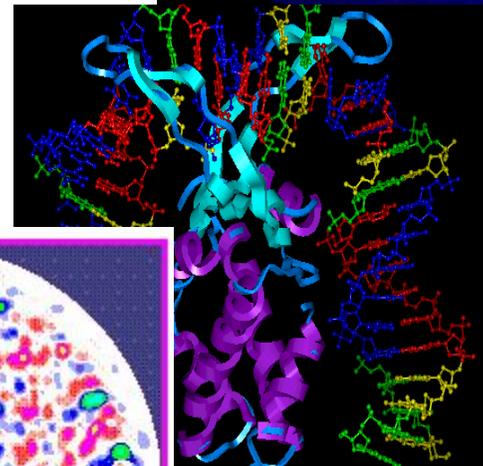
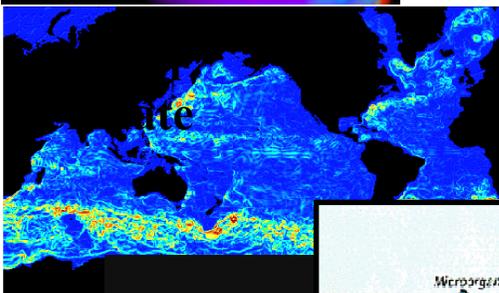
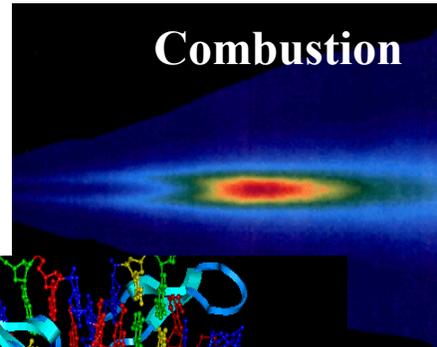
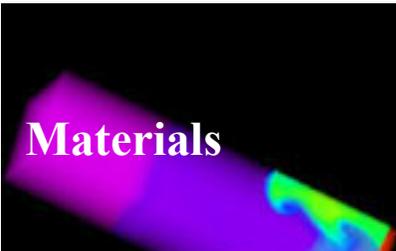
with the ultimate goal of advancing fundamental research in science central to the DOE mission

Scientific Computing – Third Pillar of Science

Many SC programs need dramatic advances in simulation capabilities to meet their mission goals

Materials

Combustion



Health Effects, Bioremediation

Components of Matter

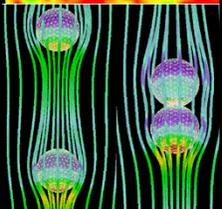
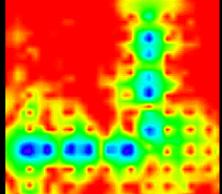
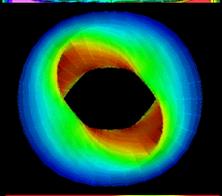
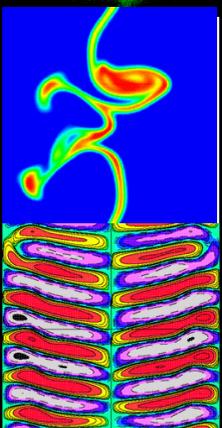
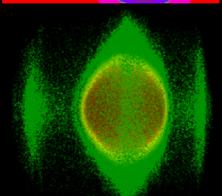
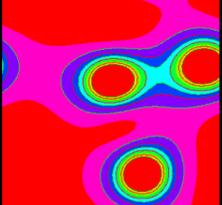
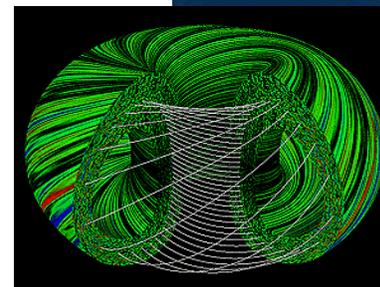
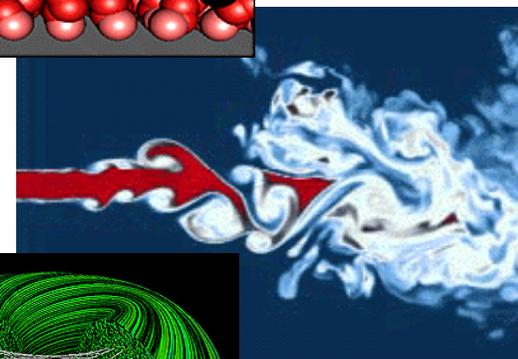
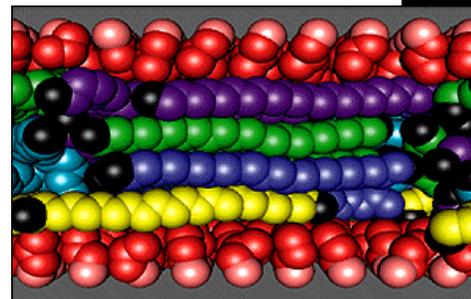
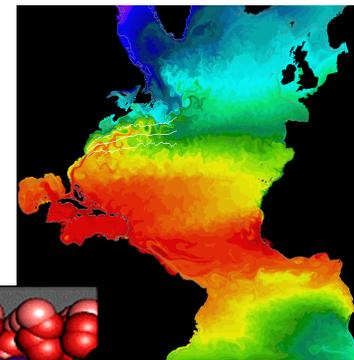
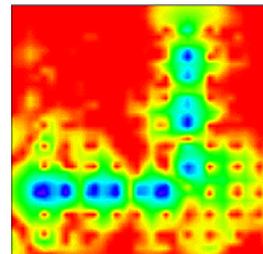
Subsurface Transport

Fusion Energy



SciDAC

- **Harness the power of terascale super-computers for scientific discovery:**
 - Form multidisciplinary teams of computer scientists, mathematicians, and researchers from other disciplines to develop a new generation of scientific simulation codes.
 - Create new software tools and mathematical modeling techniques to support these teams.
 - Provide computing & networking resources.



Addressing the Performance Gap through Software

Peak performance is skyrocketing

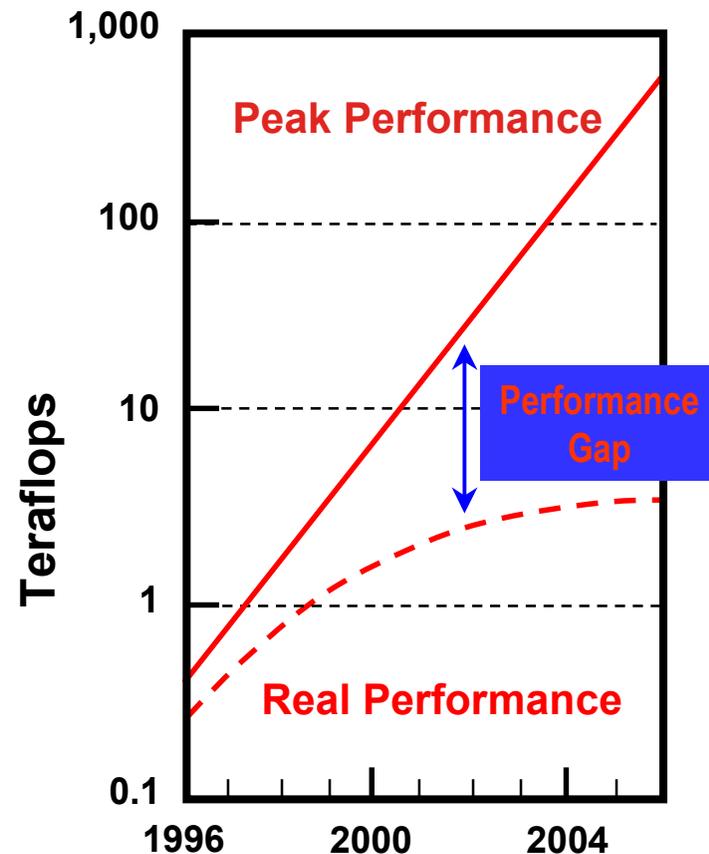
- In 1990s, peak performance increased 100x; in 2000s, it will increase 1000x

But ...

- Efficiency for many science applications declined from 40-50% on the vector supercomputers of 1990s to as little as 5-10% on parallel supercomputers of today

Need research on ...

- Mathematical methods and algorithms that achieve high performance on a single processor and scale to thousands of processors
- More efficient programming models for massively parallel supercomputers



SciDAC Focus on Software

Applications

Global Climate
 Computational Chemistry
 Fusion
 Magnetic Reconnection
 Wave-Plasma Interactions
 Atomic Physics for Edge Region
 High Energy/Nuclear Physics
 Accelerator Design
 QCD
 Supernova Research
 Neutrino-Driven Supernovae
 and their Nucleosynthesis
 Particle Physics Data Grid

Computer Science

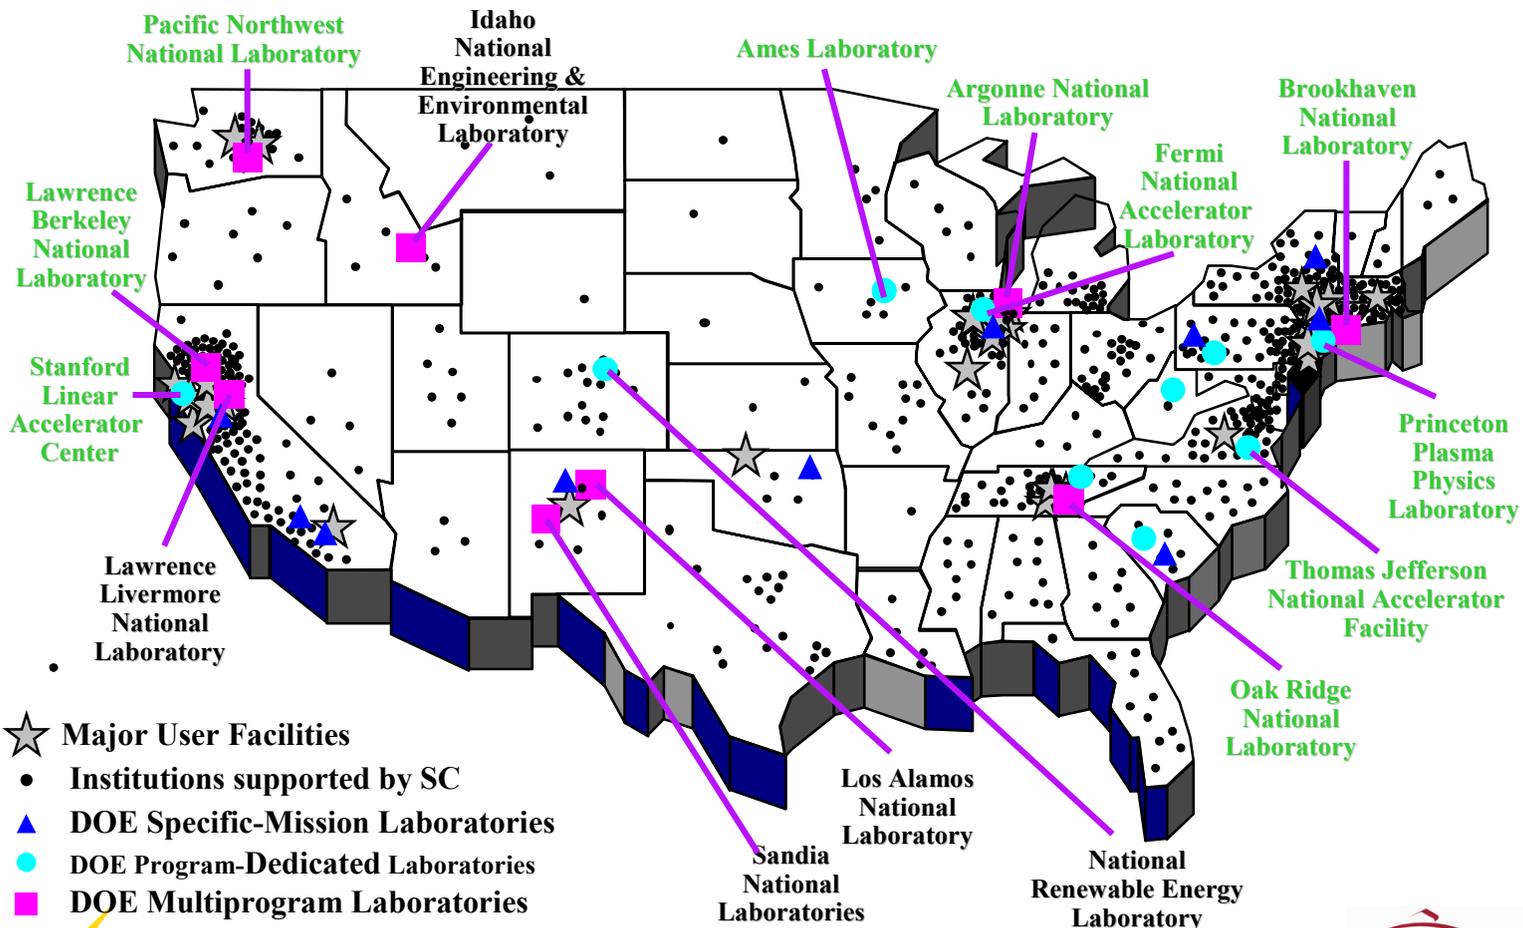
Scalable System Software
 Common Component Architecture
 Performance Science and Engineering
 Scientific Data Management

Mathematics

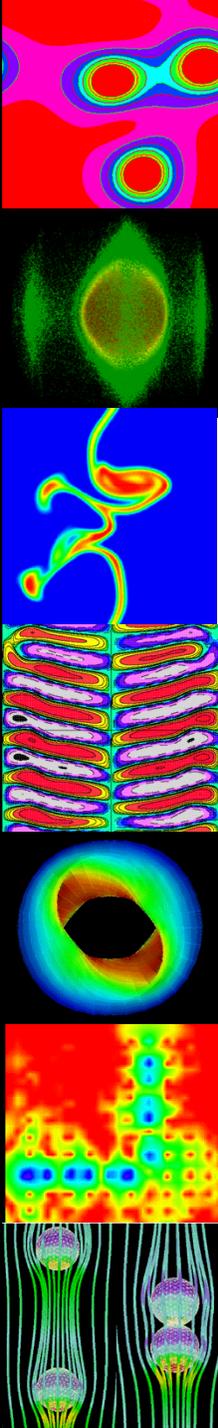
PDE Solvers/Libraries
 Structured Grids/AMR
 Unstructured Grids

7 Integrated Software
 Infrastructure Centers (ISICs)
 were established in FY01
 (3 in Berkeley)

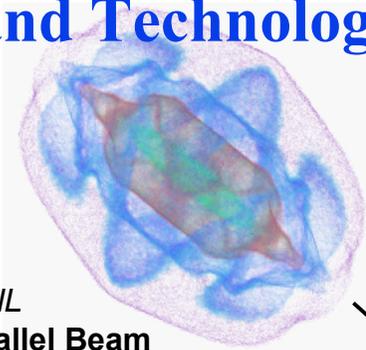
Science in the 21st Century is Distributed!



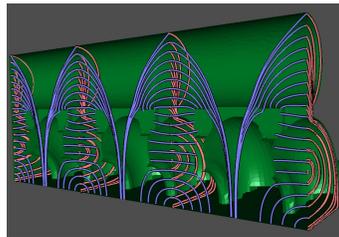
- ★ Major User Facilities
- Institutions supported by SC
- ▲ DOE Specific-Mission Laboratories
- DOE Program-Dedicated Laboratories
- DOE Multiprogram Laboratories



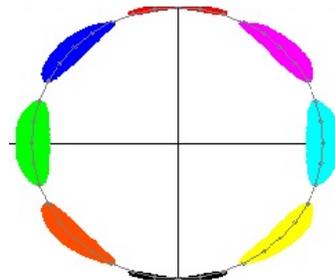
Typical SciDAC Application Project: Advanced Computing for Twenty-First Century Accelerator Science and Technology



LBL
Parallel Beam Dynamics Simulation



UC Davis
Particle & Mesh Visualization

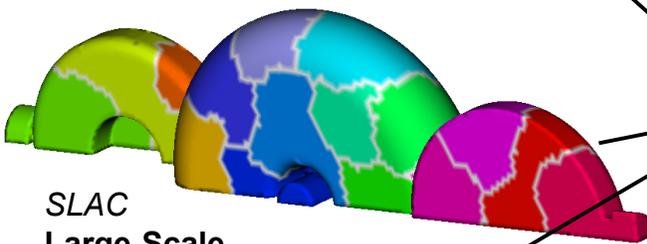


FNAL, BNL
High Intensity Beams in Circular Machines

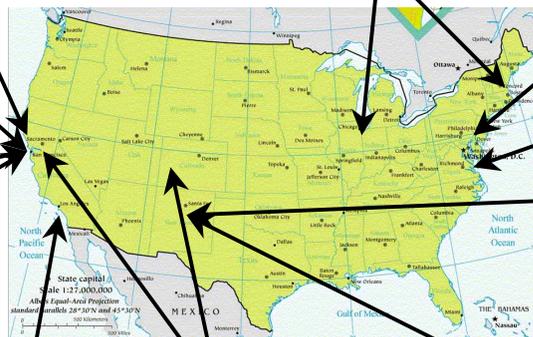
$$M = e^{if_2} e^{if_3} e^{if_4} \dots$$

$$N = A^{-1} M A$$

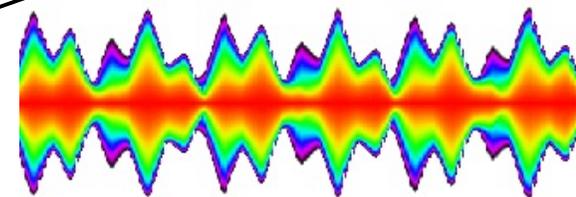
U. Maryland
Lie Methods in Accelerator Physics



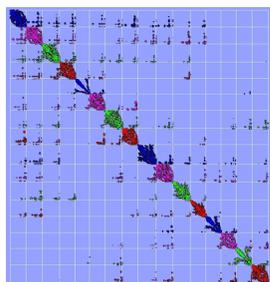
SLAC
Large-Scale Electromagnetic Modeling



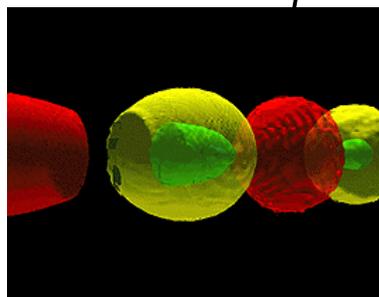
Jefferson Lab.
Coherent Synchrotron Radiation Modeling



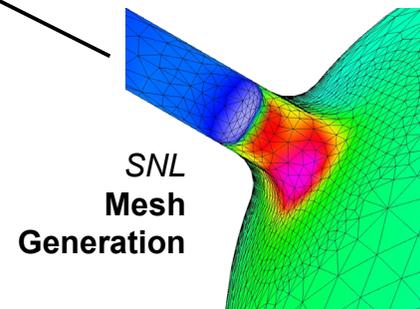
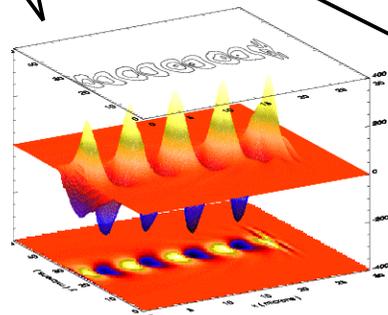
LANL
High Intensity Linacs, Computer Model Evaluation



Stanford, NERSC
Parallel Linear Solvers & Eigensolvers



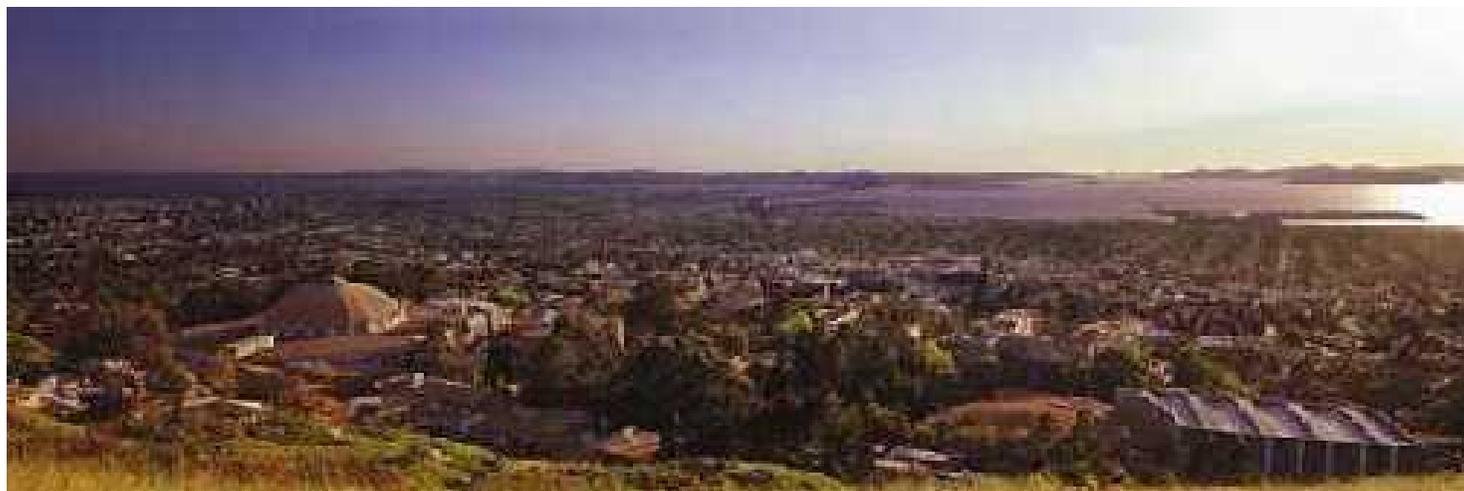
UCLA, USC, UCB, Tech-X, U. Colorado
Plasma-Based Accelerator Modeling



SNL
Mesh Generation

NERSC Center Overview

- Funded by DOE, annual budget \$28M, about 65 staff
- Supports open, unclassified, basic research
- Located in the hills next to University of California, Berkeley campus
- close collaborations between university and NERSC in computer science and computational science



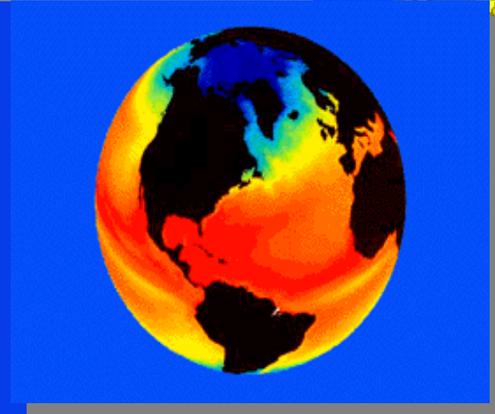
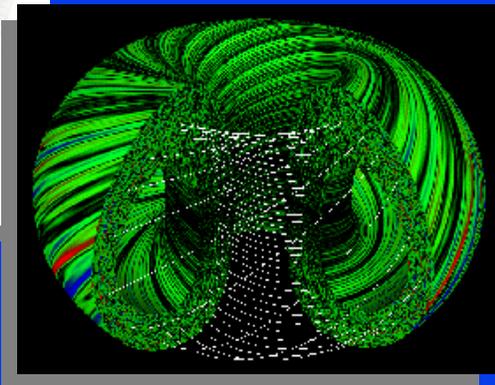
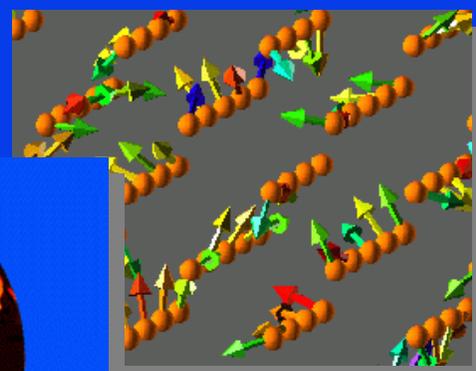
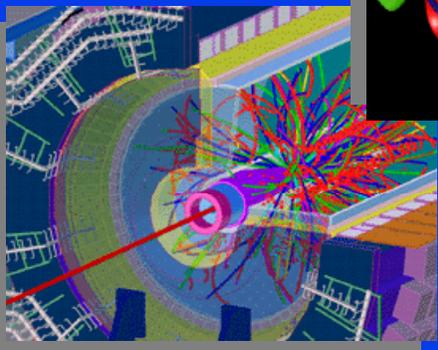
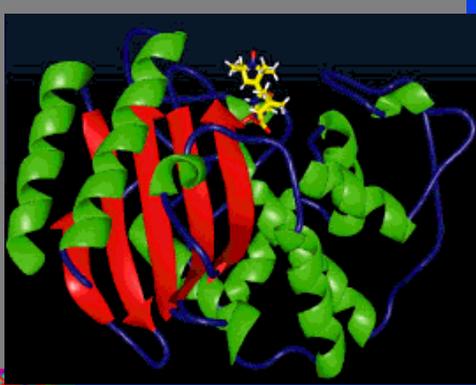


National Energy Research Scientific Computing Center

- Serves all disciplines of
the DOE Office of Science

- ~2000 Users in ~400 projects

- 20% of
allocations
to SciDAC

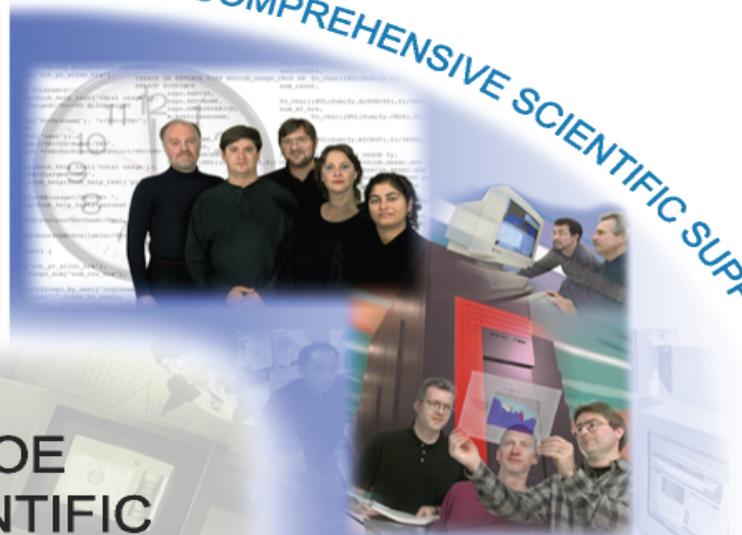


Components of the Next-Generation NERSC

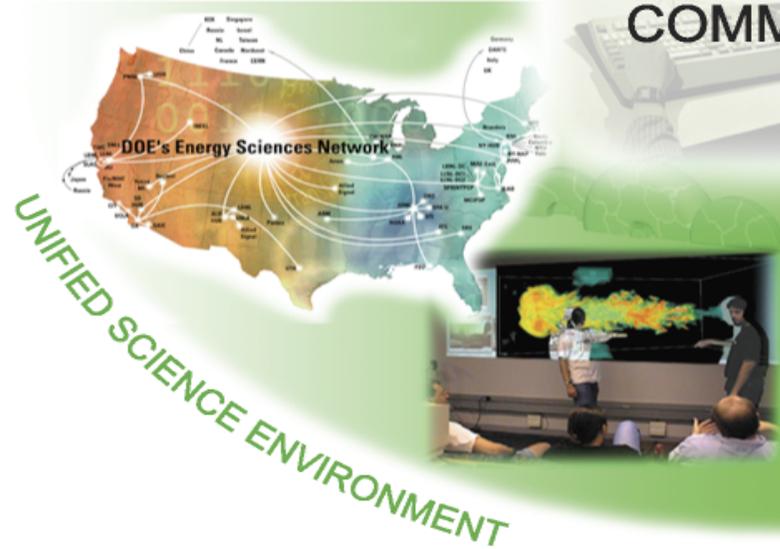
HIGH-END SYSTEMS



COMPREHENSIVE SCIENTIFIC SUPPORT



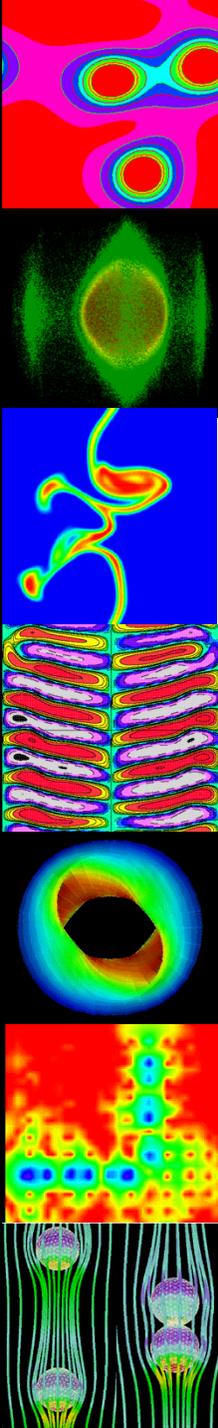
DOE SCIENTIFIC COMMUNITY



UNIFIED SCIENCE ENVIRONMENT



INTENSIVE SUPPORT FOR SCIENTIFIC CHALLENGE TEAMS



NERSC 3 (Seaborg) Upgrade to 10 Tflop/s Completed

- **System Characteristics:**
 - 416 16 way Power 3+ nodes with each CPU at 1.5 Gflop/s
 - *380 for computation*
 - 6,656 CPUs – 6,080 for computation
 - Total Peak Performance of 10 Teraflop/s
 - Total Aggregate Memory is 7.8 TB
 - Total GPFS disk will be 44 TB
 - *Local system disk is an additional 15 TB*
 - Combined SSP-2 measure is 1.238 Tflop/s
 - In production now; largest unclassified system in the U.S.



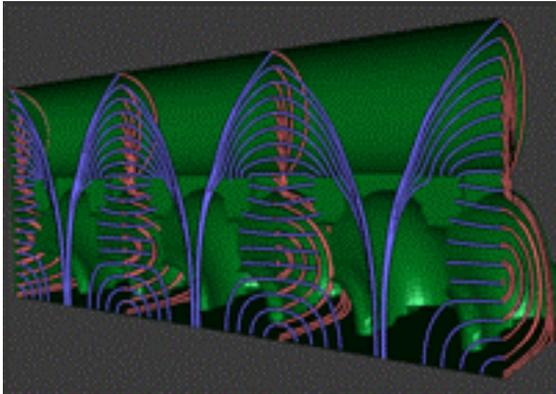
Comparison with other systems

	NERSC	ASCI White	ESC	Cheetah ORNL	PNNL (mid 2003)
Nodes	416	512	640	27	960
Processors	6656	8192	5120	864	1900
Peak Performance (Tflop/s)	10	12	40	4.5	11.4
Memory	7.8	4	10	1	6.8
Disk	60	150	700	9	53+234

PNNL system available in Q3 CY2003; 53 TB SAN + 234 TB local disk

SSP = sustained systems performance (NERSC applications benchmark)

SciDAC Project: Accelerator Science



- PI: **Robert Ryne, Berkeley Lab**
- Current Requirements:
 - 1.6 million MPP hours
 - large memory: up to 2 TB
 - 64-bit MPI

- visualize and post process up to 3 TB of data
- NERSC Provided:
 - 3 TB scratch space
 - consulting support for large memory management and performance analysis
 - CVS support and web hosting

Accelerator Science (cont.)

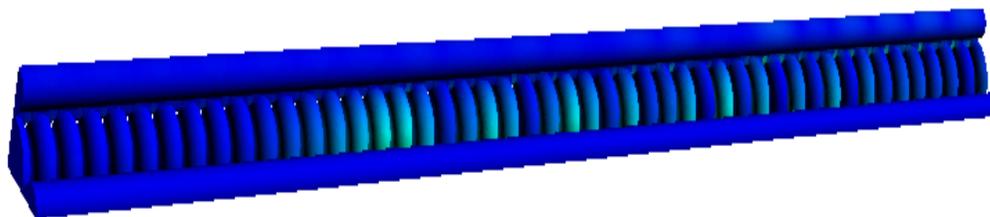
- Science Results:
 - understand beam heating for PEP-II (SLAC) upgrade
 - help design the Next Linear Collider accelerating structure
 - understand emittance growth in high intensity beams
 - study laser wakefield accelerator concepts for future accelerator design
- Future Requirements (3 years):
 - 15-20 million MPP hours
 - 5+ TB scratch space
 - continued consulting support

Applied Math. Contribution to Accelerator SciDAC: Large-scale Eigenvalue Calculations

- Calculates cavity mode frequencies and field vectors.
 - Finite element discretization of Maxwell's equations gives rise to a generalized eigenvalue problem.
 - When losses in cavities are considered, eigenvalue problems become complex (and symmetric).
 - NERSC, Stanford collaboration.
 - *Parry Husbands, Sherry Li, Esmond Ng, Chao Yang (NERSC/TOPS+SAPP).*
 - *Gene Golub, Yong Sun (Stanford/Accelerator).*



Individual cells used
 in accelerating
 structure



Omega3P model of a 47-cell section of the 206-cell Next Linear Collider accelerator structure

Future Applied Math. Contributions

- **SuperLU:**
 - Improve the interface with PARPACK.
 - Parallelize the remainder of the symbolic factorization routine in SuperLU – guaranteeing memory scalability, and making the exact shift-invert algorithm much more powerful.
 - Fill-reducing orderings of the matrix.

- **Need to improve the Newton-type iteration for the correction step, as well as the Jacobi-Davidson algorithm:**
 - SuperLU has its limitations: memory bottleneck.
 - Future plans include joint work (LBNL+Stanford) on the correction step.
 - o *Iterative solvers.*
 - o *Preconditioning techniques.*

SciDAC is first Full Implementation of Computational Science and Engineering (CSE)

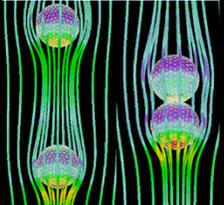
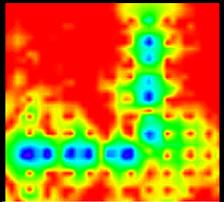
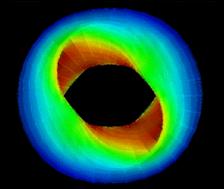
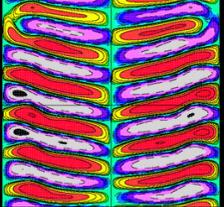
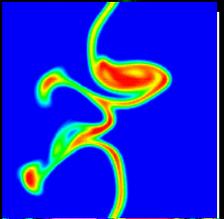
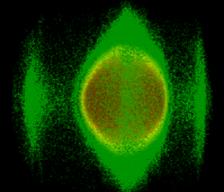
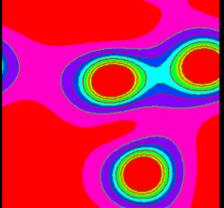
- CSE (or CSME) is a widely accepted label for an evolving field concerned with the science of and the engineering of systems and methodologies to solve computational problems arising throughout science and engineering
- CSE is characterized by
 - Multi - disciplinary
 - Multi - institutional
 - Requiring high end resources
 - Large teams
 - Focus on community software
- CSE is not “just programming” (and not CS)
- Ref: Petzold, L., *et al.*, Graduate Education in CSE, *SIAM Rev.*, 43(2001), 163-177

The Future:

more resources

better integration

next level simulation science



Future Simulation Capability Needs

Application	Simulation Objective	Sustained Capability (Tflops)	Significance
Climate Science	Calculate chemical balances in atmosphere, including clouds, rivers, and vegetation.	> 50	Provides U.S. policymakers with leadership data to support policy decisions. Properly represent and predict extreme weather conditions in changing climate.
Magnetic Fusion Energy	Optimize balance between self-heating of plasma and heat leakage caused by electromagnetic turbulence.	> 50	Underpins U.S. decisions about future international fusion collaborations. Integrated simulations of burning plasma crucial for quantifying prospects for commercial fusion.
Combustion Science	Understand interactions between combustion and turbulent fluctuations in burning fluid.	> 50	Understand detonation dynamics (e.g. engine knock) in combustion systems. Solve the "soot" problem in diesel engines.
Environmental Molecular Science	Reliably predict chemical and physical properties of radioactive substances.	> 100	Develop innovative technologies to remediate contaminated soils and groundwater.
Astrophysics	Realistically simulate the explosion of a supernova for first time.	>> 100	Measure size and age of Universe and rate of expansion of Universe. Gain insight into inertial fusion processes.

Examples of Science Needs

(www.ultrasim.info)

- **Final Release Documents**

- **ASCR**

- *Reasserting U.S. Leadership in Scientific Computation*
- *Coping with the Ultrascale Tsunami of Scientific Data*

- **BER**

- *Accelerating Climate Prediction*

- **BES**

- *Computational Design of Catalysts*
- *Autoignition and Control of “Flameless” Combustion*
- *The Fundamentals of Soot Birth and Growth*
- *Accelerating the Revolution in Computational Materials Science*
- *Simulation Real-World Combustion Devices*
- *Computational Nanoscience on Earth Simulator Class Machines: A Revolution in Materials Science*

- **FES**

- *Fueling Design Optimization via Supercomputing – Fusion Energy Sciences*
- *U.S. Leadership in Scientific Computation – Fusion Energy Sciences*

Examples of Science Needs (cont'd)

(www.ultrasim.info)

- **Working Documents**

- **ASCR**

- *UltraNet – enabling scientific insight*
 - *Ultrascale Visualization – gleaming insight through scientific visualization*

- **BER**

- *Computational Environmental Molecular Science*
 - *Grand Challenges in Computational Structural and Systems Biology*
 - *Benefits of an Earth Simulator Class Machine for U.S. Climate Science*
 - *Realizing the Potential of the Genome Revolution: Facilities for 21st Century Systems Biology Science*
 - *Computational Structural Genomics*
 - *Protecting the Nation's Groundwater*

- **BES**

- *Turbulence and “Self-Accelerated” Combustion*
 - *Impact of Earth Simulator-Class Computers on Computational Nanoscience and Materials Science*
 - *Computational Materials Science*

Examples of Science Needs (cont'd)

(www.ultrasim.info)

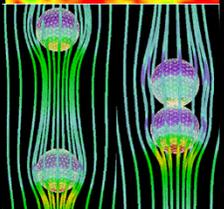
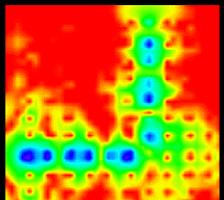
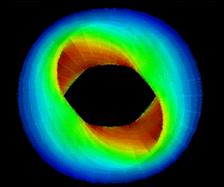
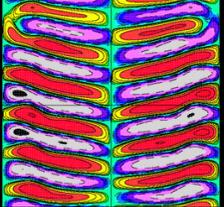
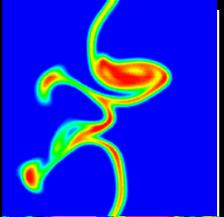
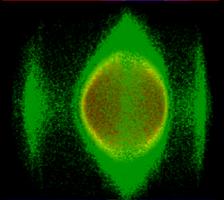
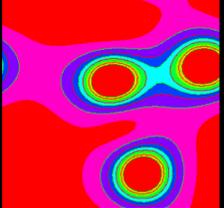
- **Working Documents (cont'd)**

- **FES**

- *Computational Fusion Energy Research: The Need for New Levels of Supercomputing*

- **HENP**

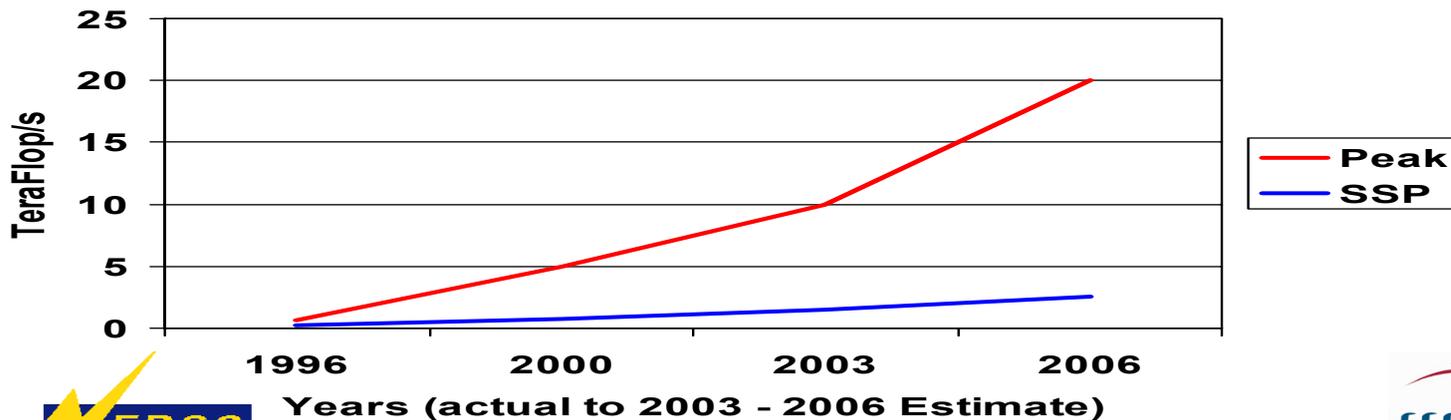
- *An Astrophysics Response to the Challenge of the Earth Simulator*
 - *High-Energy and Nuclear Physics*
 - *Ultrascale Computing in Accelerator Science and Technology: Scientific Opportunities and Impact*



Future Resource Needs: The Divergence Problem

- The requirements of high performance computing for science and engineering and the requirements of the commercial market are diverging.
- The commercial cluster of **SMP** approach is no longer sufficient to provide the highest level of performance
 - Lack of memory bandwidth
 - High interconnect latency
 - Lack of interconnect bandwidth
 - Lack of high performance parallel I/O
 - High cost of ownership for large scale systems

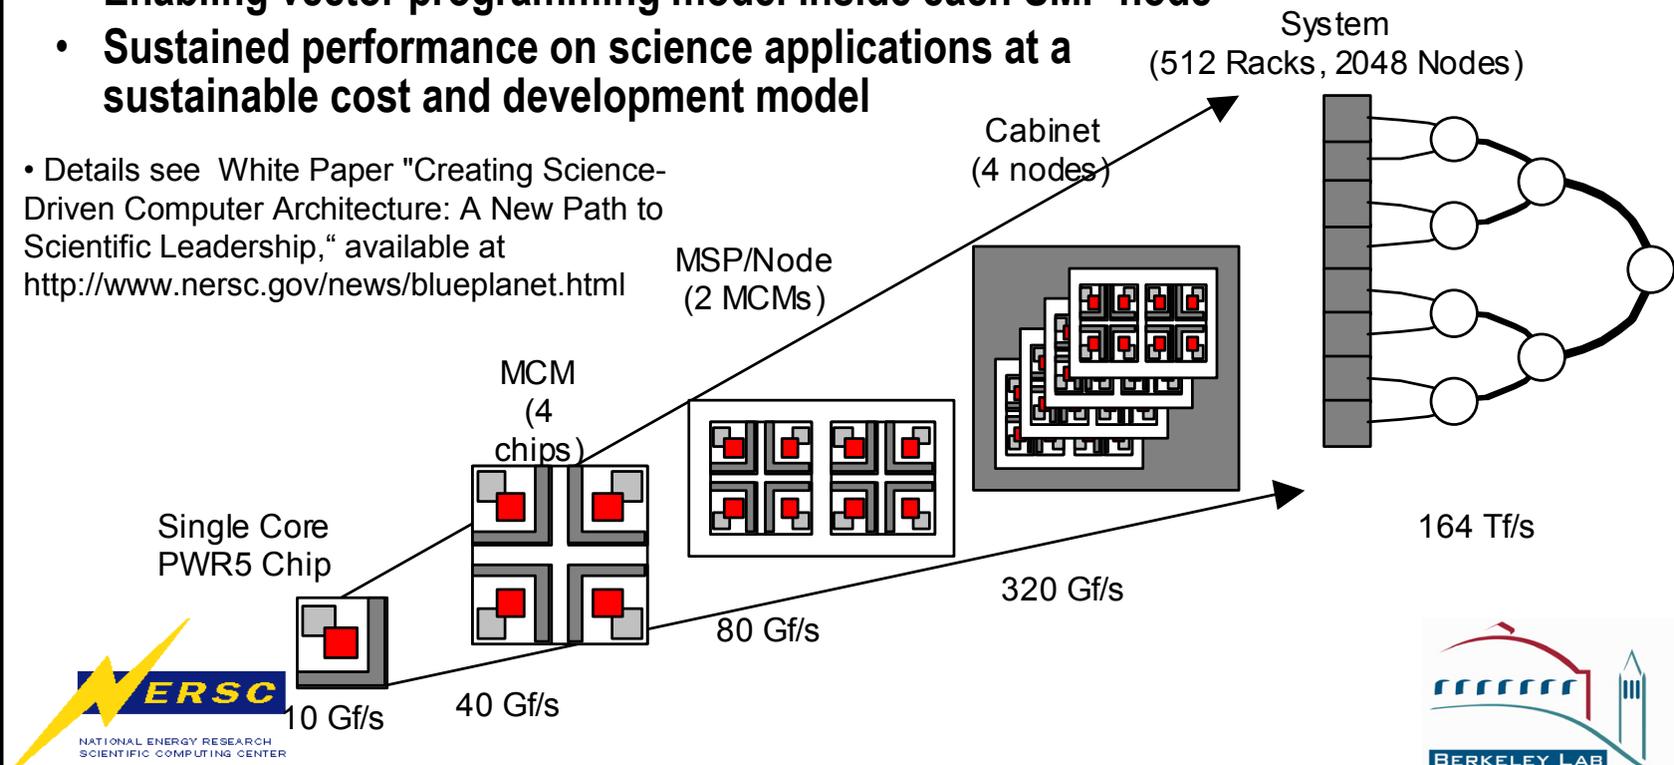
Divergence



Blue Planet: A Conceptual View

- **Increasing memory bandwidth – single core**
 - 8 single cpus are matched with memory address bus limits for full memory bandwidth
- **Increasing switch bandwidth – 8-way nodes**
- **Decreased switch latency while increasing span**
- **Enabling vector programming model inside each SMP node**
- **Sustained performance on science applications at a sustainable cost and development model**

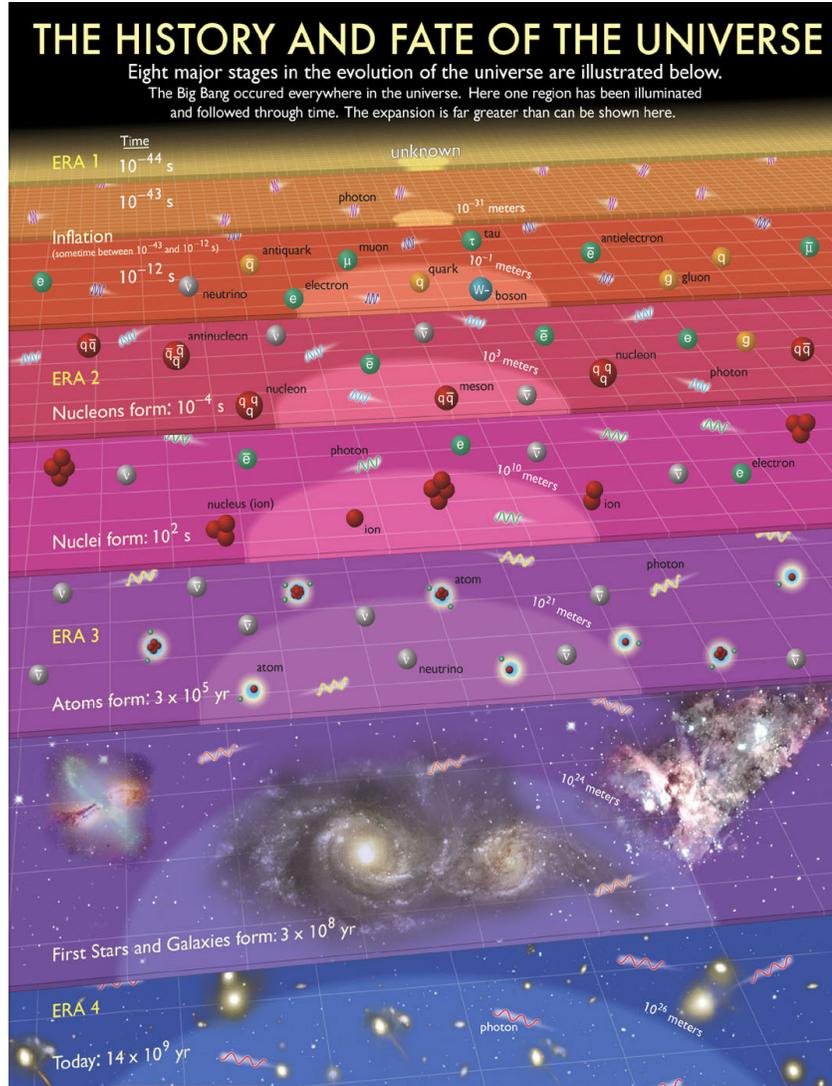
• Details see White Paper "Creating Science-Driven Computer Architecture: A New Path to Scientific Leadership," available at <http://www.nersc.gov/news/blueplanet.html>



Simulation Science by 2012: Cosmic Simulator

Science driven vision of a computational framework in 2010.

The Cosmic Simulator is the concept of providing an integrated framework in which component simulations can be linked together to provide a coherent, end-to-end, history of the Cosmos.



Cosmic Simulation

- Now conceptually possible to compute and simulate the physical history of the Universe
- A defined set of stages which correspond to major phase changes of basic physics and matter
- One stage's outputs provides the next's inputs
- Each stage is today comparable to large scale SciDAC project
- Challenging physics and computation issues
 - Petascale computing
 - Analysis of data from new space borne telescope require grid infrastructure
- Valuable simulated data sets for comparison with observations and research
 - Data management of petascale data sets

The Future

more resources:

**no limits seen to growth in demand for
supercomputer resources seen**

better integration:

**computational science and engineering will
become recognized as discipline**

next level simulation science:

**large scale simulation environments will
emerge that allow computer simulation at
unprecedented scale**