U.S. Department of Energy

Office of Science

View from Germantown

Michael Strayer

Associate Director, Advanced Scientific Computing Research

March 3, 2009

U.S. Department of Energy

Office of Science

ASCR Vision

First in Computational Science

"***Best in class in advancing science and technology through modeling and simulation***"

Facilities

Enabling Technologies

Computational Partnerships

Michael Strayer, ASCAC Presentation March 15-16, 2006

Where are we today?

U.S. Department of Energy

Office of Science

Delivering the Science

Scientific Discovery and the Role of High End Computing

U.S. Department of Energy

Office of Science

Top 10 Computational Science Accomplishments

Titles in Blue – SciDAC: Titles in Black - INCITE

Rank	Title
1	**Modeling the Molecular Basis of Parkinson's Disease (Tsigelny)**
2	**Discovery of the Standing Accretion Shock Instability and Pulsar Birth Mechanism in a Core-Collapse Supernova Evolution and Explosion (Blondin)**
3	**Prediction and Design of Macromolecular Structures and Functions (Baker)**
4	**Understanding How Lifted Flame Stabilized in a Hot Coflow (Yoo)**
5	**New Insights from LCF-enable advanced kinetic simulations of global turbulence in fusion systems (Tang)**
6	**High Transition Temperature Superconductivity: A High-Temperature Superconducting State and a Pairing Mechanism in 2-D Hubbard Model (Scalapino)**
7	**PETSc: Providing the Solvers for DOE High-Performance Simulations (Smith)**
8	**Via Lactea II, A Billion Particle Simulation of the Dark Matter Halo of the Milky Way (Madau)**
9	**Probing the properties of water through advanced computing (Galli)**
10	**First Provably Scalable Maxwell Solver Enables Scalable Electromagnetic Simulations (Kovel)**

U.S. Department of Energy

Office of Science

DCA++ Achieves 1.3 Petaflops

2008 Gordon Bell Prize Winner

- High T_c Superconducting in Cuprates
  - 2-D Hubbard Model
  - Study Materials with Disorders/Impurities
  - First petaflop application
  - Spurred community debate
  - Inspired SNS experiment
- DCA++
  - Monte Carlo Method
  - **10X Speedup** by Scientific Computing Group at OLCF through:
    - Delaying memory intensive operations (reorder barriers)
    - Mixed Precision arithmetic (move fewer bits per flop)

Doug Scalapino

Philip Anderson

More Information: Thomas Schulthess ASCAC presentation 1:45-2:30 pm today

U.S. Department of Energy

Office of Science

Linear Scaling Code Unlocks New Frontiers

2008 Gordon Bell Special Prize Winner

- LS3DF
  - novel decomposition and patching scheme to do *ab initio* DFT calculations for large systems.
    - Electrostatic Energy computed globally
    - Quantum Mechanical Energy computed by dividing the whole system into small fragments, calculating the energies of these fragments, and then combining the separate fragment energies
    - Patching scheme tailored to cancel out the artificial boundary effects caused by division of the system into smaller fragments.
  - **400 times faster** than other algorithms for 14,000 atoms
  - **Potential applications in nanoscale materials for next generation solar cells**

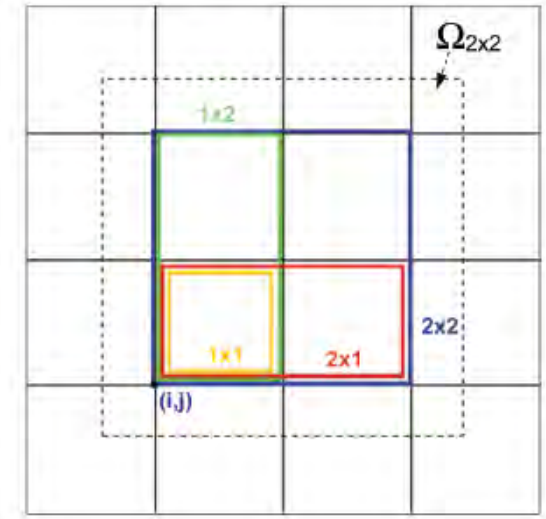

More Information: Juan Meza ASCAC presentation 2:45-3:30 pm today

U.S. Department of Energy
Office of Science

Recognitions

Presidential Early Career Award for Scientists and Engineers

PECASE Award: Bert Debusschere
Stochastic Dynamical Systems: Spectral Methods for the Analysis of Dynamics and Predictability

- *"For introducing rigorous, mathematical methods capturing stochastic uncertainties in computational biology and providing a framework for simulation-based discovery; and for service to the Sandia Diversity Council and Foreign National Networking Group."*

Bert Debusschere, Technical Staff Sandia National Laboratory

U.S. Department of Energy

Office of Science

Recognitions

National Academy of Engineering

James Sethian, head of the LBNL Mathematics Group and a professor of mathematics at UC Berkeley

Applied Mathematics researcher, James Sethian elected to National Academy of Engineering

- Honored "for the development of efficient methods of tracking moving interfaces."
- These techniques have a wide range of applications, including problems in fluid mechanics, combustion, manufacturing of computer chips, computer animation, image processing, structure of snowflakes and the shape of soap bubbles.

U.S. Department of Energy

Office of Science

Delivering the Software Foundation

Software Developed under ASCR Funding

Programming Models

Active Harmony

ARMCI
ATLAS
Berkeley UPC Compiler
Charm++
Fountain

FT-MPI
Global Arrays
Kepler
MVAPICH
OPEN-MPI
OpenUH
PVM

Development/ Performance Tools

BABEL
Berkeley Lab Checkpoint Restart (BLCR)
Dyninst API
Fast Bit
Goanna
HPCtoolkit

Jumpshot
KOJAK
MPIP
MRNet
Net PIPE
OpenAnalysis
PAPI
ROSE
ScalaTrace
STAT
TAO

TAU
Hpcviewer

Math Libraries

ACTS COLLECTION

ADIC
Hypre
ITAPS Software Suite
LAPACK
Mesquite

MPICH2
OpenAD
OPT++
PETSc
ROMIO
ScaLAPACK
Sparskit-CCA
Trilinos

System Software

Cluster Command & Control
High-Availability OSCAR HA-OSCAR
LWK-Sandia
PVFS
ZeptoOS

Collaboration

enote

Visualization /Data Analytics

BeSTMan
Parallel netCDF
Virtual Data Tool Kit

Miscellaneous

Libmonitor

U.S. Department of Energy

Office of Science

Base Programs
FY 2008-09 Request for Proposals

- **Multiscale Mathematics and Optimization for Complex Systems**
  - Letters of Intent
    - 426 one-page Letters of Intent (LOIs) were submitted by March 3
      - Optimization (138 LOIs): 76 led by Labs, 62 led by universities
      - Multiscale Mathematics (288 LOIs): 186 led by Labs, 102 led by universities
    - From LOIs, ASCR encouraged 114 full proposals due by April 28
  - Optimization peer-review panel convened June 10-11
    - **38** proposals: multilevel, stochastic, dimension reduction, mixed integer, inverse problems
  - Multiscale Mathematics peer-review panels convened June 23-24 and June 25-26, 2008
    - **35** proposals: hybrid methods, geoscience, fluids and plasmas, materials, data-model fusion
    - **25** proposals: uncertainty and sensitivity analysis, stochastic methods, additional topics

- **Petascale Tools**
  - 97 proposals received representing 34 projects
  - Topics included High Performance tools, Correctness tools, Development Environment and Scalable Infrastructure
  - Review August 26-27, 2008

- **Next Generation Networking for Science**
  - 40 proposals received
  - Topics included high performance networking and middleware R&D
  - Review on-going

New Awards from all solicitations are held up because of the Continuing Resolution

U.S. Department of Energy

Office of Science

SciDAC Projects
Mid term Reviews 2009

Upcoming Reviews

Science Applications (SA)s and Science Application Partnerships (SAP)s

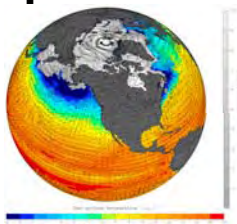

BER-ASCR *April 9-10, 22-24* (Climate, Subsurface)

HEP-ASCR-NP-BES *April 21-22* (Accelerator Modeling)

NNSA-ASCR *May 6-8, 15* (Materials/Chemistry, Turbulence)

FES-ASCR *May 14* (Fusion)

Centers for Enabling Technology (CETs) & Institutes *April 20-29*

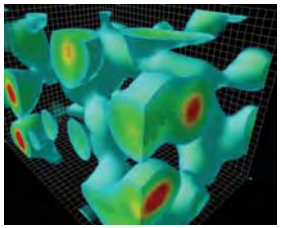

Distributed Systems *May 11-13*

Completed Reviews:

ASCR-HEP-NP-NSF (Distributed Systems)
HEP-NP-NNSA-ASCR (Astrophysics, QCD, Nuclear physics)

U.S. Department of Energy

Office of Science

Innovative and Novel Computational Impact on Theory and Experiment (INCITE)

2009 INCITE Awards: Allocations by Discipline

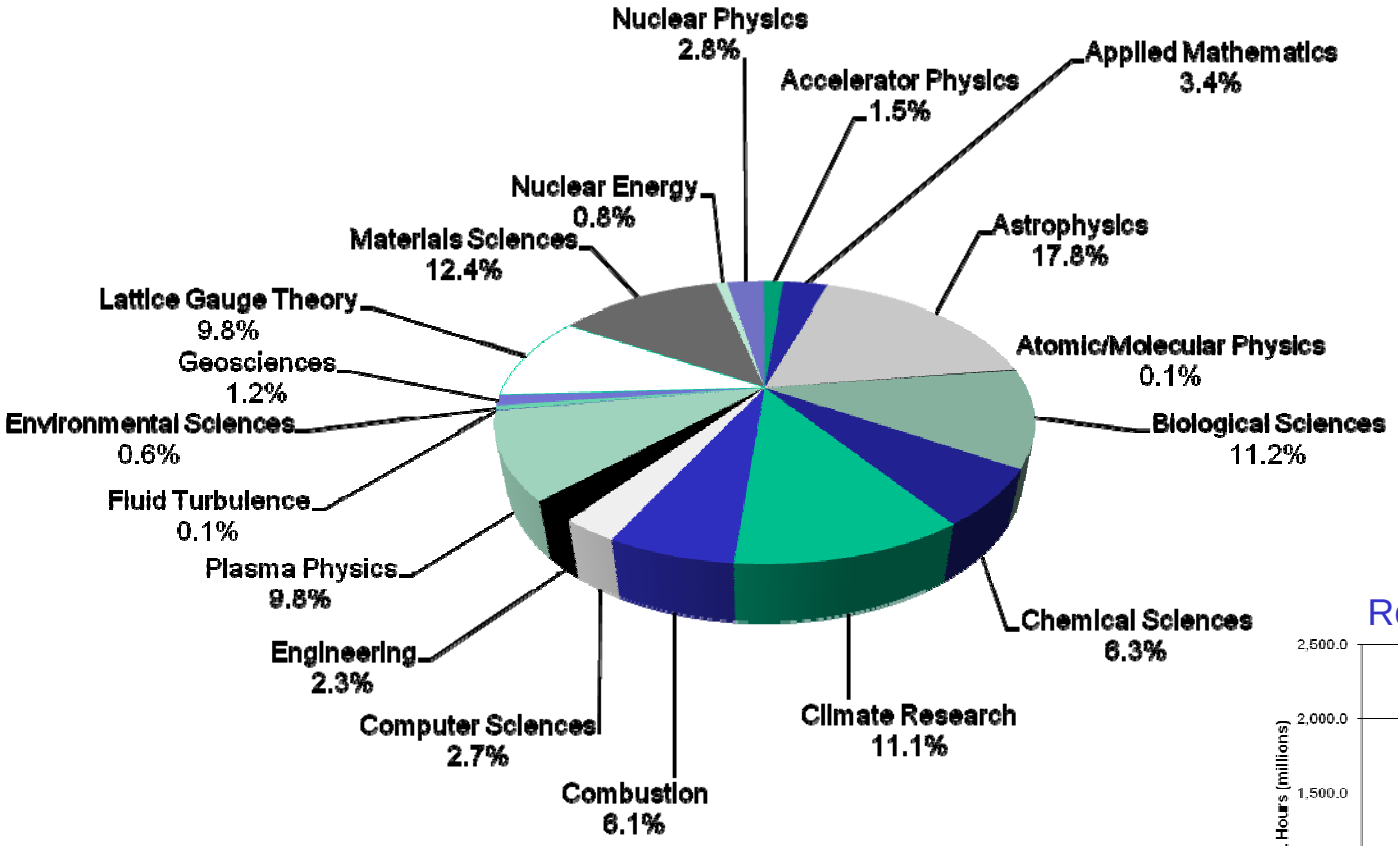

INCITE Trends
Requests and Allocations

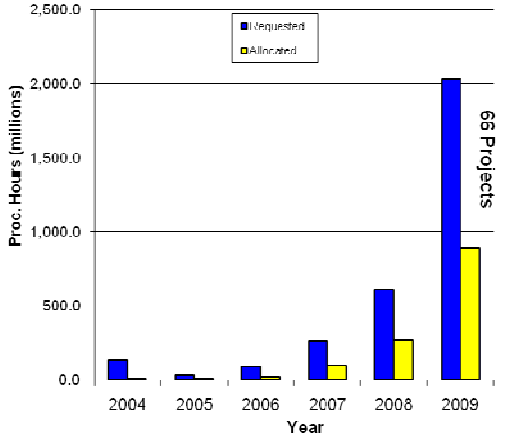

U.S. Department of Energy

Office of Science

Oak Ridge's Cray XT5 Breaks the Petaflop Barrier

Jaguar	Total	XT5	XT4
Peak Performance	1,645	1,382	263
AMD Opteron Cores	**181,504**	**150,176**	**31,328**
System Memory (TB)	**362**	**300**	**62**
Disk Bandwidth (GB/s)	**284**	**240**	**44**
Disk Space (TB)	**10,750**	**10,000**	**750**
Interconnect Bandwidth (TB/s)	**532**	**374**	**157**

More Information: Buddy Bland ASCAC presentation 1:00-1:45 pm today

Argonne's IBM Blue Gene/P – 556 TFs

U.S. Department of Energy

Office of Science

National Energy Research Scientific Computing Center (NERSC)

- Located at Lawrence Berkeley National Lab
  - IBM Power 5 Bassi: 6.7 Tflop/s
  - Linux Opteron Cluster Jacquard: 3.1 Tflop/s
  - Cray XT4 Franklin upgraded to 350 Tflop/s, currently in pre-acceptance
- NERSC-6 Project
  - RFP issued in September 2008
  - Proposals are being reviewed

Franklin

Bassi and Jacquard

U.S. Department of Energy

Office of Science

ESnet
40 Gbps Core

ESnet news

June–November 2008

Princeton Gets a 6,400 Percent Increase in Bandwidth With ESnet Upgrades

ESnet finished improving its Internet connections to several institutions on Princeton University's Forrestal Campus, including the Princeton Plasma Physics Lab (PPPL), the High Energy Physics (HEP) Group within the Physics Department at Princeton University, and the National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory (GFDL).

Now researchers around the globe can access data from these science facilities with increasing speeds and scalability, helping enable international collaborations on bandwidth-intensive applications and experiments.

"This is a great achievement," says Steve Cotter, head of ESnet. "With the availability of cutting-edge instruments and supercomputers, scientists around the world are collaborating to carry out large experiments that produce tremendous amounts of data. This upgrade links Princeton's physics researchers to that data through our robust and reliable network. ESnet's via point-to-point dedicated circuits and IP services at multiple gigabit per second speeds."

The Princeton network upgrade took approximately five months to complete, and involved running fiber optic cabling underground from the Forrestal Campus outside Princeton, New Jersey, along Route 1 to South Brunswick, then to Philadelphia, where it is transported across the ESnet infrastructure to ESnet's main point of presence in McLean, Va.

On the Princeton campus, the PPPL's Internet connection is now operating at 10 gigabit speeds, 10 billion bits per second, significantly continued on page 2

PPPL (shown here) and GFDL are both located on Princeton University's Forrestal Campus.

ESnet4 Provides Critical Link for U.S. Researchers Accessing LHC Data

Approaching the speed of light, millions of protons will collide per second when the Large Hadron Collider (LHC) comes on line next year. The experiment will generate more data than the international scientific community has ever tried to manage. Scientists suspect the outcome of these "subatomic smashups" will provide valuable insights into the origins of matter and dark energy in the universe.

As thousands of researchers across the globe anxiously await the results of the experiment, getting the massive amounts of data to them is no insignificant task. Fortunately, network engineers at the U.S. Department of Energy's (DOE) Energy Sciences Network (ESnet) foresaw this data challenge years ago and developed ESnet4, a new large-scale science data transport network with enough bandwidth to transport multiple streams of 10 gigabits of information per second — the equivalent of transmitting 500 hours of digital music per second for each 10 gigabit line.

The LHC, which straddles the Swiss and French borders on the outskirts of Geneva, will be the first experiment to fully utilize the advanced capabilities of this network, which connects DOE national laboratories to researchers across the country and collaborators worldwide.

"ESnet4 is one of the most robust scientific data networks in existence," says Steve Cotter, Department Head for ESnet. "The science environment of today is very different from that of a few years ago. ESnet4 provides the high-speed, extremely reliable connectivity between labs and U.S. and international research institutions required to support the inherently collaborative, global nature of modern large-scale science." continued on page 3

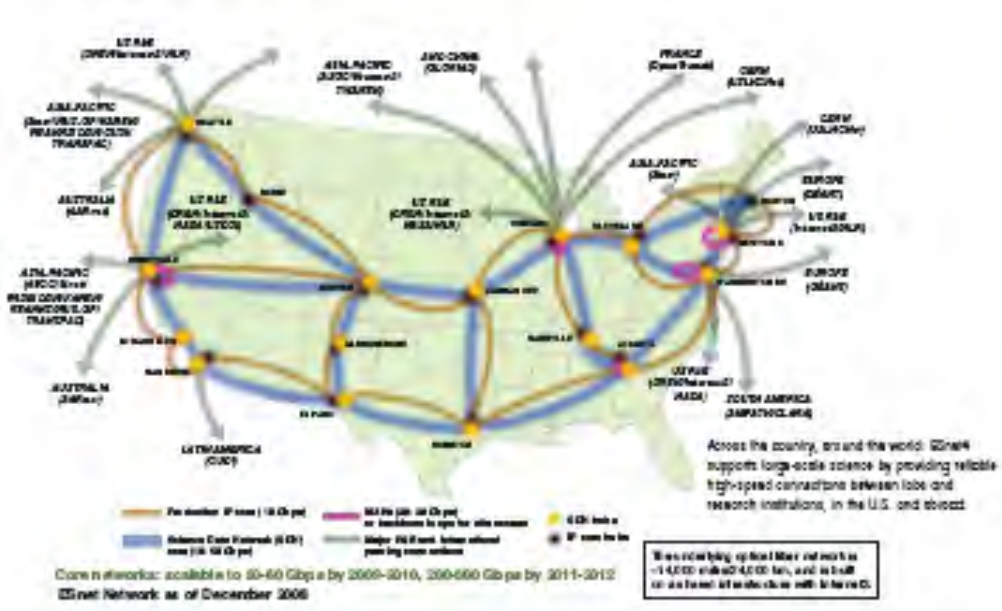

Leader in Networks for Science

- OSCARS
- PerfSONAR
- DanteInternet2CanarieESnet

More Information: Steve Cotter ASCAC presentation 10:15-11:00 am Wednesday

U.S. Department of Energy

Office of Science

ASCR
Budget Details

	FY 2008 Approp.	FY 2009 Request	FY 2009 HEWD	FY 2009 SEWD
Applied Mathematics	36,885	43,164	**46,164**	43,164
Computer Science	27,226	34,618	34,618	34,618
Computational Partnerships	53,767	52,064	**54,064**	52,064
Next Gen. Networking for Science	12,017	17,221	17,221	17,221
High Performance Production Computing	54,200	54,790	54,790	54,790
Leadership Computing Facilities	110,158	115,000	**120,000**	115,000
High Performance Network Facilities & Testbeds	23,936	25,000	25,000	25,000
Research and Evaluation Prototypes	23,585	17,000	17,000	17,000
Subtotal, ASCR	*341,774*	*358,857*	*368,857*	*358,857*
All other (GPP, GPE, SBIR/STTR)	9,399	9,963	9,963	9,963
Total, ASCR	351,173	368,820	**378,820**	368,820

U.S. Department of Energy

Office of Science

ASCR Staffing

- New Hires:
  - Karen Pao, February, 2009 as an Applied Math Program Manager in the Research Division
- In Progess
  - Offer made to candidate for Data and Viz Program Manager in Computer Science
- Positions to be posted soon
  - Computer Scientist Program Manager to fill F. Johnson's position
  - Collaboratories Program Manager

Maintaining “Best in Class”

What computing will be needed to enable the grand challenges in Science?

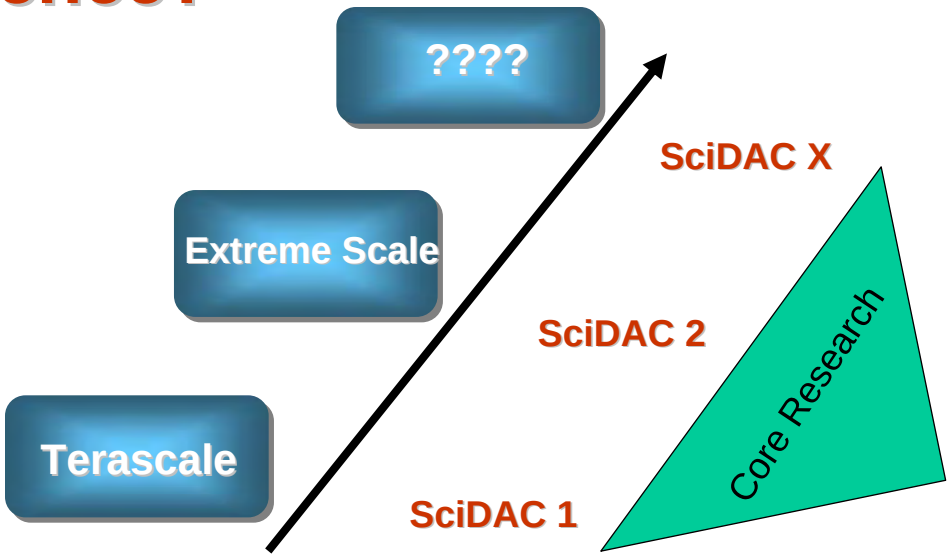

I climb the "Hill of Science,"
I "view the landscape o'er;"
Such transcendental prospect,
I ne'er beheld before!

Emily Dickinson

U.S. Department of Energy

Office of Science

ASCR's Vision

- **Deliver Petascale Science Today**
  - Continue to make the Leadership Computing Facilities available to the very best science through Innovative and Novel Computational Impact on Theory and Experiment (INCITE).
  - Continue to work with Pioneer Applications to deliver scientific results from day one.
- **Build the Intellectual Foundation for the Future**
  - Continue to nurture –
    - World class mathematics and computer science research efforts
    - Applications critical to DOE missions through Scientific Discovery through Advanced Computing (SciDAC).
  - Provide direct support for "bleeding-edge" research groups willing to take on the risk of working with emerging languages and operating systems.
  - Foster innovative research at the ever blurring boundary between Applied Mathematics and Computer Science.
- **Realize the Promise of Extreme Scale**
  - Work with key science applications to identify opportunities for new research areas only possible through extreme scale computing.
  - Support innovative research on advanced architectures and algorithms that accelerates the development of hardware and software that is well suited to extreme scale computational science.

U.S. Department of Energy

Office of Science

Identify the Barriers

Gathering Community Input

Experimental Particle Physics

Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale

December 9-11, 2008 · Menlo Park, CA

Simulation and Algorithms

- Full exploration of the physics frontier is currently limited by our ability to reach statistically meaningful levels of sensitivity & precision in simulation & analysis
  - In order to make new physics discoveries, need to simulate:
    - tails of the detector response
    - tails of Standard Model distributions.
  - Algorithm development to optimize performance
    - Detector simulation and Physics generator level
    - Sophisticated analysis in a large parameter space

More Information: Paul Messina ASCAC presentation 8:30-9:15 am Wednesday

U.S. Department of Energy

Office of Science

Identify the Needs
Community Input

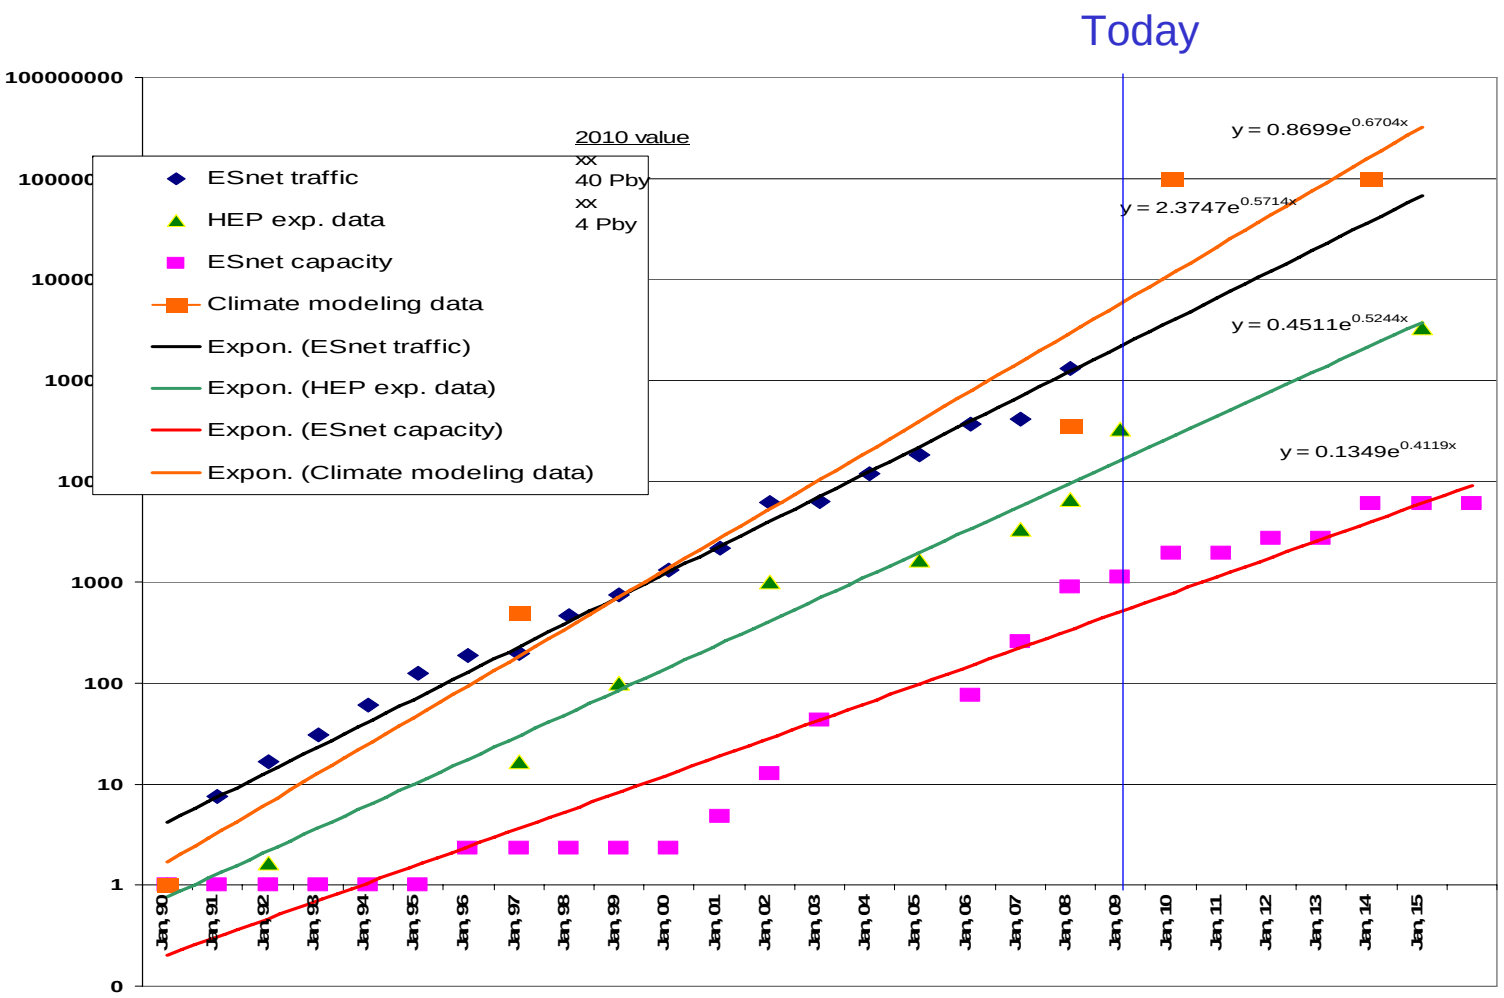

Data gathered from ESnet Requirements Gathering Workshops

U.S. Department of Energy

Office of Science

Identify the Barriers
Multicore

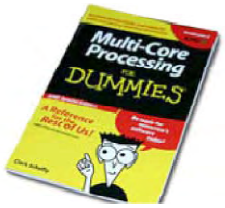

- Multiple parallel general-purpose processors (GPPs)
- Multiple application-specific processors (ASPs)

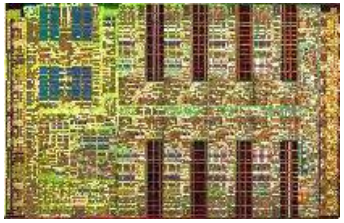

IBM Cell
1 GPP (2 threads)
8 ASPs

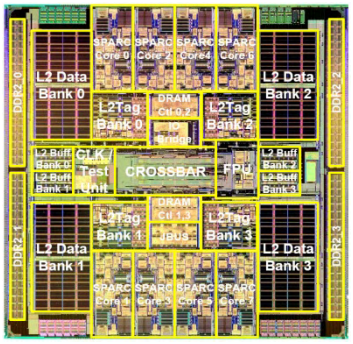

Sun Niagara
8 GPP cores (32 threads)

Intel Network Processor
1 GPP Core
16 ASPs (128 threads)

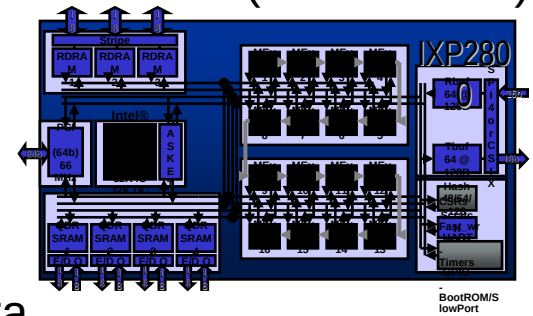

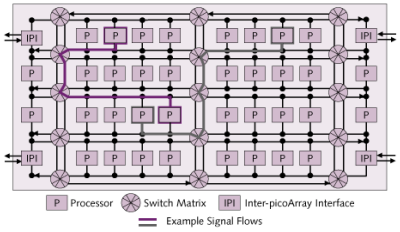

Picochip DSP
1 GPP core
248 ASPs

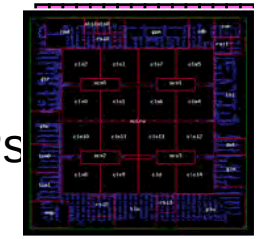

Cisco CRS-1
188 Tensilica GPPs

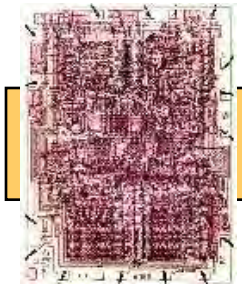

Intel 4004 (1971):
4-bit processor,
2312 transistors,
~100 KIPS,
10 micron PMOS,
11 mm² chip

1000s of processor cores per die

"The Processor is the new Transistor" [Rowen]

U.S. Department of Energy

Office of Science

Identify the Barriers

Cybersecurity

Current State	Opportunities for Mathematical Research			Future State
▪ "cybersecurity today is more of a craft than a science." ▪ Policies and practices are derived from trial and error	Predictive modeling of large-scale networks ▪ Simulate realistic networks for testing network defense algorithm ▪ Identification of multiple network attributes	Cyber threats discovery ▪ Rare-event detection methods ▪ Data analysis and modeling in presence of missing values and uncertainty	Network dynamics ▪ Adapt game theory and dynamical systems methods to anticipate threats and alter interactions between attackers and defenders ▪ Automated model checking and formal reasoning methods to reduce software and hardware errors	▪ Predictive Awareness for Secure Systems ▪ Self-Protective Data and Software ▪ Trustworthy systems for untrusted Components

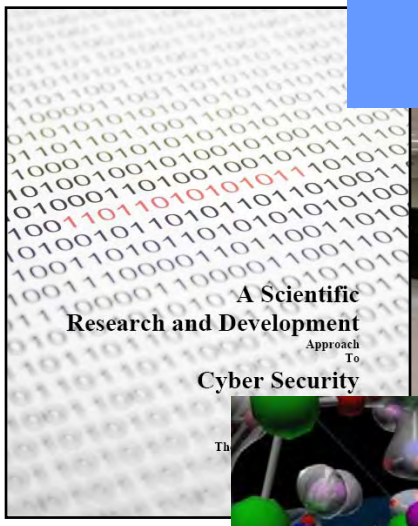

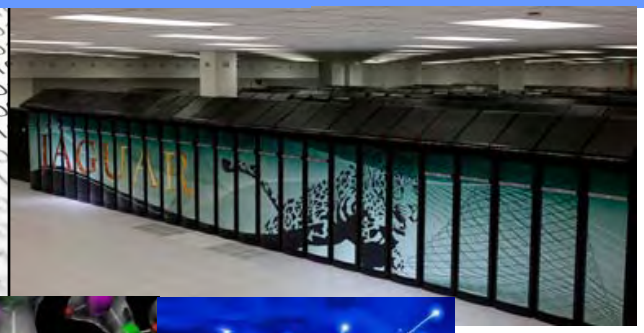

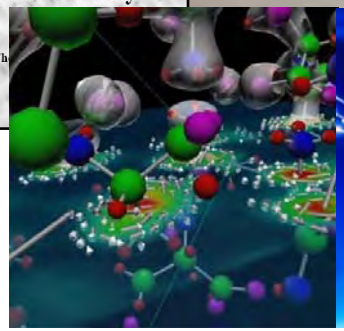

DOE Cyber challenges

- Open Science:
  - Remote access to high-performance computing facilities; international scientific collaborations; collaboratories; grid computing; emergence of 'clouds'
  - Ensuring the integrity of large scientific data files.
- Other DOE Programs
  - Securing high-value IT assets and data

More Information: Charlie Catlett ASCAC presentation 3:30-4:15 pm today

U.S. Department of Energy

Office of Science

Outreach
DOE Community

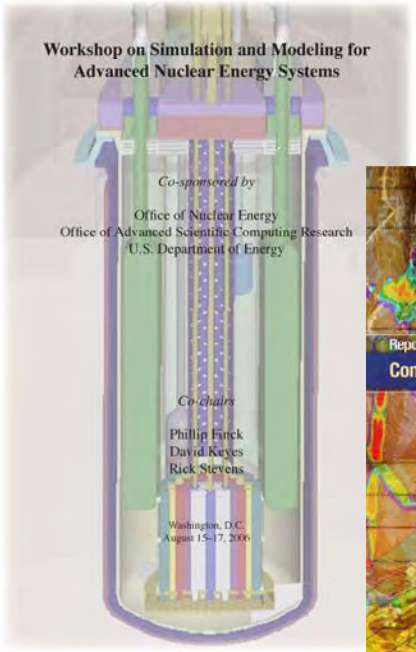

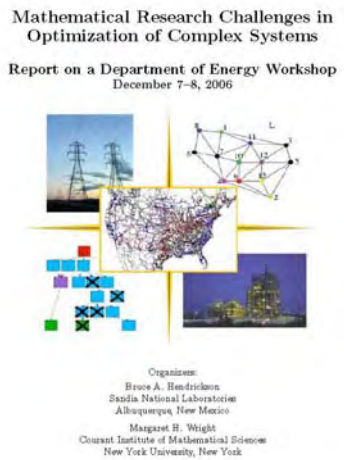

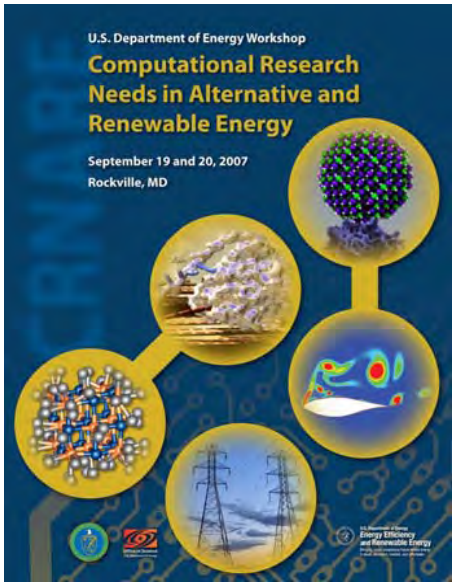

More Information: Patricia Hoffman, Acting Assistant Secretary, OE, ASCAC presentation 11:00-11:45 am Wednesday

U.S. Department of Energy

Office of Science

Outreach
Around the World

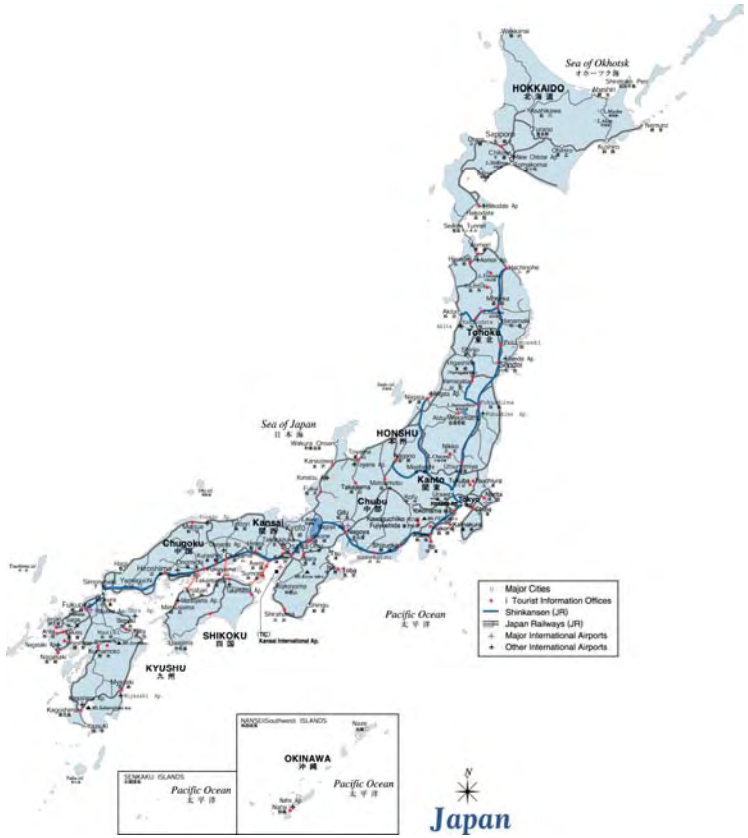

More Information: Pete Beckman, ASCAC presentation 11:00-11:45 am today

U.S. Department of Energy

Office of Science

ASCR Research Strategy

- **Break new ground in science:**
  - **SciDAC:** Deliver computational tools and techniques to advance DOE-science through modeling and simulation
  - **Multiscale mathematics:** Discover methods and algorithms to fully-describe understanding of nature over vast scales of time and space.

- **Provide knowledge and foundational tools:**
  - **Applied Mathematics:** Develop algorithms for solving complex, mission-relevant science problems. Understand the mathematics that underlies prediction.
  - **Computer Science:** Facilitate the use of emerging Leadership-scale computing resources. Understand the implications of new computing architectures. Develop tools to extract meaningful information from peta-byte data sets
  - **Next Generation Networking for Science:** Understand the management and performance of federated 100 gbps networks. Enable geographically distributed research teams to collaborate- share data, assess results, plan and conduct experiments.

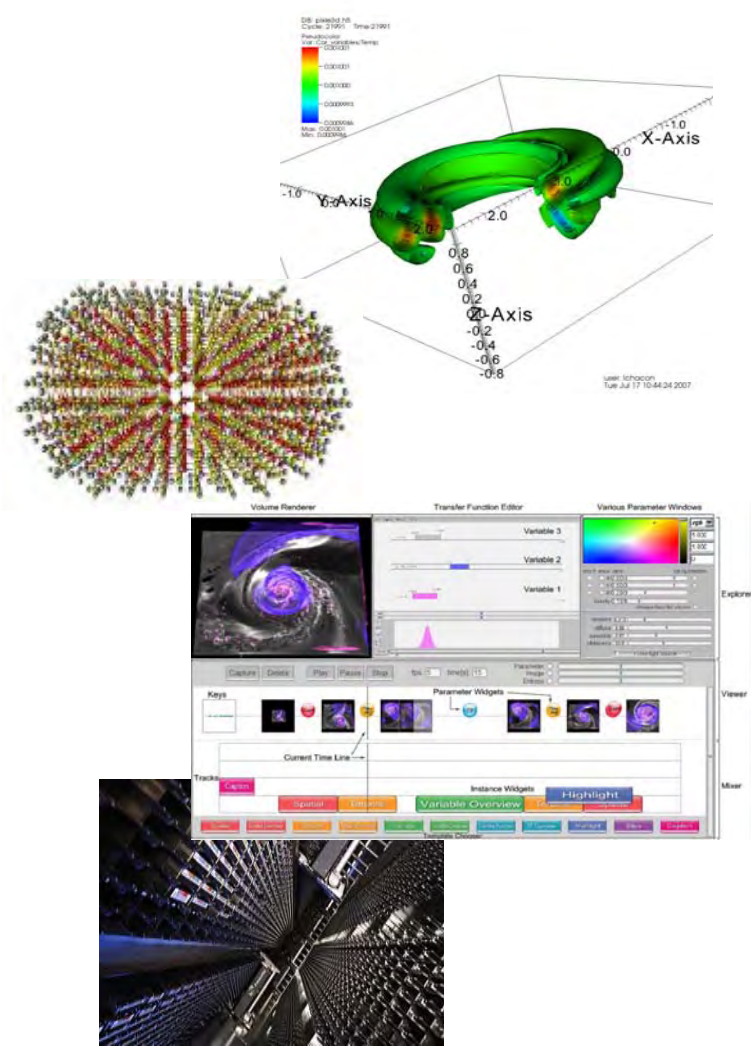

U.S. Department of Energy

Office of Science

ASCR
Facilities Strategy

- **Providing the Tools** – High-End Computing
  - High-Performance Production Computing - National Energy Research Scientific Computing Center **(NERSC)** at Lawrence Berkeley National Laboratory
    - Delivers high-end capacity computing to entire DOE SC research community
  - Leadership-Class Computing – **Leadership Computing Centers at Argonne National Laboratory and Oak Ridge National Laboratory**
    - Delivers highest computational capability to national and international researchers through peer-reviewed **Innovative and Novel Computational Impact on Theory and Computation (INCITE)** program (80% of resources)
- **Investing in the Future** - Research and Evaluation Prototypes
- **Linking it all together** – Energy Sciences Network (ESnet)