U.S. Department of Energy

Office of Science

Advanced Scientific Computing Research

An Overview for the Advanced Scientific Computing Advisory Committee

March 13, 2003
Washington, DC

Ed Oliver
Walt Polansky

U.S. Department of Energy

Office of Science

Staff

- Ed Oliver, Associate Director for Advanced Scientific Computing Research
- Dan Hitchcock, Senior Scientific Advisor
- Linda Twenty, Senior Budget & Financial Specialist
- Melea Baker, Senior Administrative Assistant

- Walt Polansky, Acting Director MICS

- Gary Johnson, ACRTs, Genomes to Life
- Fred Johnson, Computer Science
- William (Buff) Miner, NERSC & Scientific Applications
- Thomas Ndousse-Fetter, Network Research
- Chuck Romine, Applied Mathematics
- Mary Anne Scott, Collaboratories
- George Seweryniak, ESnet
- John van Rosendale, Visualization/Data Management, Applied Mathematics- SciDAC
- Jane Hiegel, Administrative Assistant

- Vacancies (3)

Phone- 301-903-5800
Fax- 301-903-7774
http://www.sc.doe.gov/production/octr/mics/index.html

Mission

Discover, develop, and deploy the computational and networking advances that enable researchers in the scientific disciplines to analyze, model, simulate, and predict complex physical, chemical, and biological phenomena important to the Department of Energy (DOE).

support a broad research portfolio in advanced scientific computing – applied mathematics, computer science, networking and collaboratory software

operate supercomputers, a high performance network, and related facilities.

U.S. Department of Energy

Office of Science

Mathematical, Information and Computational Sciences

$ in thousands (actual expenditures)

Base Research	FY2001	FY2002
Univ.	14,352	12,668
Labs	37,372	37,186

SciDAC	FY2001	FY2002
Univ.	17,548	6,413
Labs	19,895	31,030

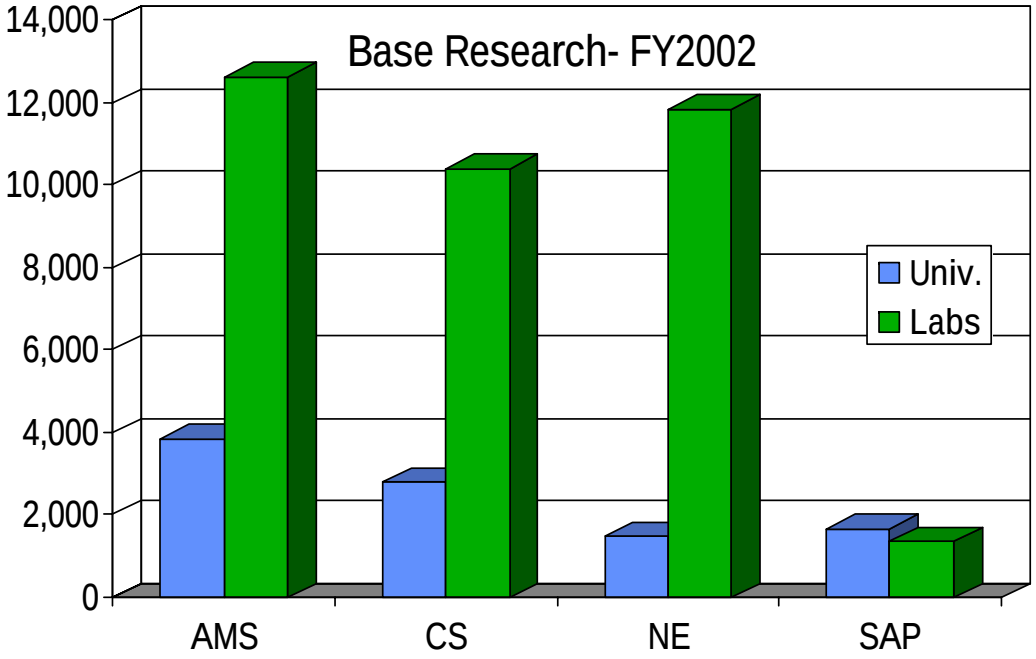

Legend

AMS- Applied Mathematical Sciences
CS- Computer Sciences
NE- Network Environment- (National Collaboratories- Advanced Computing Software Tools, Network Research)
SAP- Scientific Application Projects- (inc. Genomes to Life)
SciDAC- Scientific Discovery through Advanced Computing

U.S. Department of Energy

Office of Science

MICS Activities
reported at October, 2002 ASCAC Meeting

FY2002 Activity	Status
Conducted a workshop and 8 Town Meetings to evaluate Earth Simulator impact	Science case- 15 working documents; 11 final release documents www.ultrasim.info
Launched Early Career Principal Investigator activity to strengthen core research program	17 awards in FY2002; FY2003 Call for Proposals closed February 20, 2003
Convened ASCAC-BESAC sponsored workshop on Computational Nanoscience	Basis for Call for Proposals www.science.doe.gov/grants/Fr03-17.html
Conducted workshop on networking requirements for future of science	Discussions continuing
Conducted Genomes to Life workshops on applied mathematics and computer science	Collaboration continuing
Initiated ESnet backbone upgrade from 622 Mbs (OC12) to 10 Gbs (OC192) to service increased networking requirements for science	Northern route at OC192 as of March 5, 2003; southern route at OC48 en route to OC192

Budgets

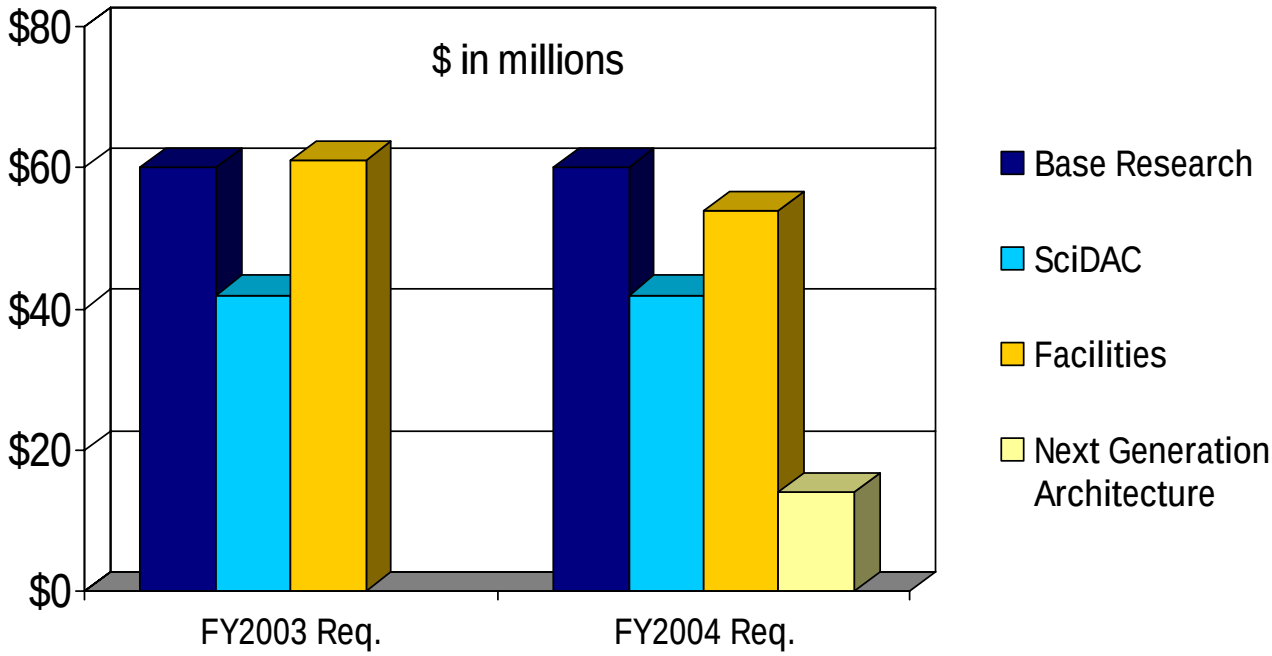

Fiscal Year	Request	Appropriation
2002	$156,170,000	$147,159,000
2003	$163,557,000	$164,480,000*
2004	$170,490,000	TBD

* Following General Reduction & Omnibus Rescission

MICS Activities
reported at October, 2002 ASCAC Meeting

Continued

FY2003 Plans	Status
Initiate reviews of applied mathematics and collaboratory pilot research activities	Review of AMS lab activities held each October; collaboratories- March/April, 2003
Initiate review of SciDAC portfolio	Computer science ISICs- Under way Math ISICs- May, 2003
Continue workshops and Town Meetings to assess UltraScale Simulation needs	No recent activity

FY2003 Program Attributes

- Advanced Architectures
  - $3M in FY2003, per Congressional intent
  - Lease payments on Cray X1; early start on evaluation
- Base research and SciDAC at FY2002 levels
- Computational nanoscience
  - 64 preproposals submitted
  - FY2003 req.- $3M (similar amount in BES request)
- Genomes to Life at FY2002 levels
- Early Career Principal Investigators
  - 65 grant applications submitted; under review

U.S. Department of Energy

Office of Science

Next Generation Architecture
$ in millions

Objective- To identify and to address major architectural bottlenecks (e.g. internal data movement in very large systems) in the performance of existing and planned DOE science applications

FY2003 Request- $0
FY2004 Request- $14

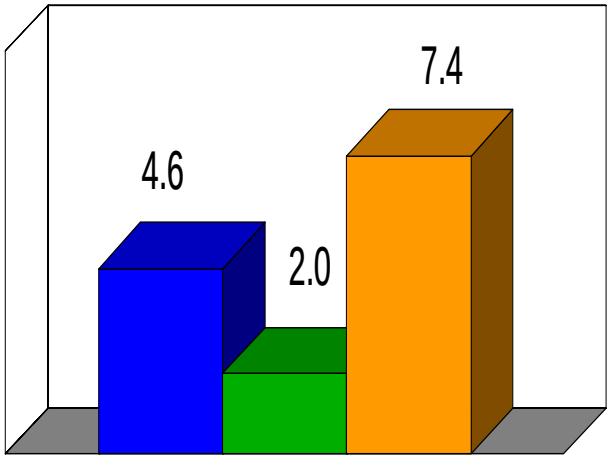

Activities

- Impact of alternative computer architectures on application performance
- Improve application performance and system reliability through software development
- Evaluate hardware testbeds of sufficient size to understand key issues.

Research- Operating Systems/Runtimes
Research- Application Performance/Scaling
Evaluation- Cray X1

U.S. Department of Energy

Office of Science

Program Execution FY2004

Research to enable...

Basic Research

...simulation of complex systems

...distributed teams, remote access to facilities

...Applications

BES, BER, FES, HEP, NP

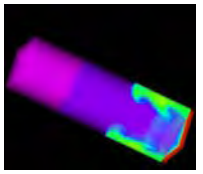
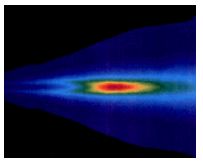
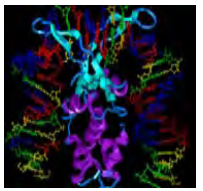
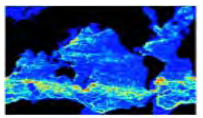
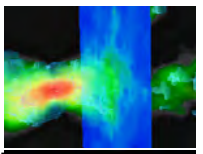
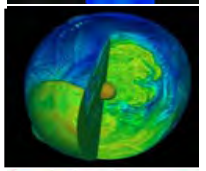


- Applied Mathematics
- Computer Science
- Network Environment
- Scientific Applications
- Genomes to Life

SciDAC!

- Nanoscience
- Grid enabling research
- Integrated Software Infrastructure Centers

(Mathematicians, computer scientists, application scientists, and software engineers)

Next Generation Architecture

- Nanoscience
- Materials
- Chemistry
- Combustion
- Accelerator
- High energy Physics
- Nuclear physics
- Fusion
- Climate
- Astrophysics
- Biology

High Performance Computing and Network Facilities for Science

National Energy Research Scientific Computing Center (NERSC)

Advanced Computing Research Testbeds

Energy Sciences Network (ESnet)

U.S. Department of Energy

Office of Science

Simulation Capability Needs
FY2005 Timeframe - Partial List

Application	Simulation Need	Sustained Computational Capability Needed (Tflops)	Significance
Climate Science	Calculate chemical balances in atmosphere, including clouds, rivers, and vegetation effects.	> 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.

U.S. Department of Energy

Office of Science

Applied Mathematics Research...

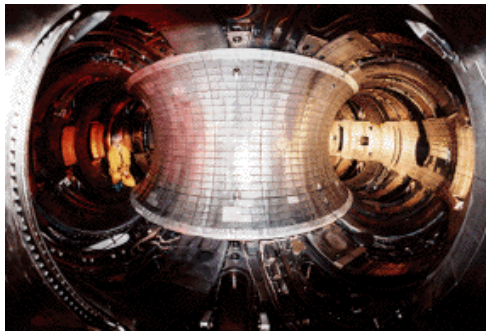

Design Refinement

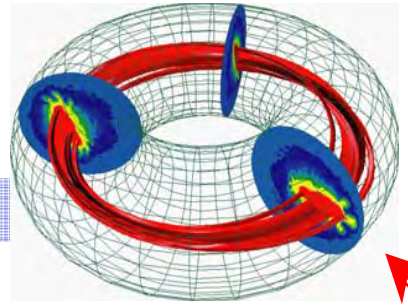

Optimization- advanced optimization theory, efficient run ensembles.
Predictability- error estimation, parameter estimation, model reduction

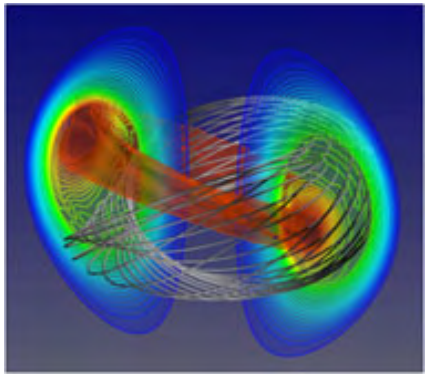

...and its contribution to Fusion Energy Sciences (an example)

Mathematical description- scientifically accurate across multiple scales, with proper boundary conditions.

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} + \kappa_{divb} \nabla\nabla \cdot \mathbf{B}$$

$$\mathbf{E} = -\mathbf{V} \times \mathbf{B} + \eta \mathbf{J}$$

$$\mu_0 \mathbf{J} = \nabla \times \mathbf{B}$$

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\mathbf{V}) = \nabla \cdot D \nabla n$$

$$\rho\left(\frac{\partial \mathbf{V}}{\partial t} + \mathbf{V} \cdot \nabla \mathbf{V}\right) = \mathbf{J} \times \mathbf{B} - \nabla p + \nabla \cdot \nu\rho \nabla \mathbf{V}$$

$$\frac{n}{\gamma - 1}\left(\frac{\partial T}{\partial t} + \mathbf{V} \cdot \nabla T\right) = -p\nabla \cdot \mathbf{V} + \nabla \cdot n\left[(\chi_{\parallel} - \chi_{\perp})\hat{\mathbf{b}}\hat{\mathbf{b}} + \chi_{\perp}\mathbf{I}\right] \cdot \nabla T + Q$$

Discretization- mesh technology, functional analysis, robustness, efficient computability, proper treatment of boundaries.

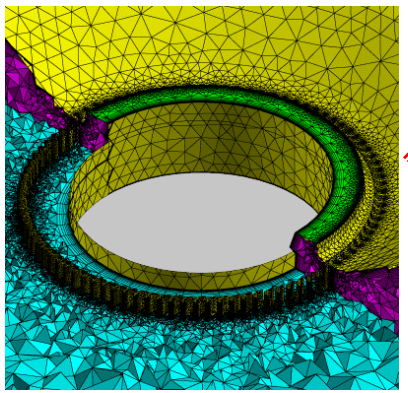

Computational solution- high-performance computing; accurate, efficient, scalable, tunable, robust, modular, and fast numerical algorithms

Research Challenges Architectures and Systems

- Memory management
  - Intelligent, dynamic, adaptive, under programmer's control through new language/compiler techniques
- Legacy issue
  - SMP clusters have driven the development of system software, libraries, and applications for the past 10 years
- Systems- MTBF will be measured in hours, or less
  - Faults common, system must isolate/recover/maintain scalability
- Applications- understanding what is going on
  - Scalability of debuggers, I/O and file systems, visualization

U.S. Department of Energy

Office of Science

Science Applications- Bit Rate Budget
End-to-end

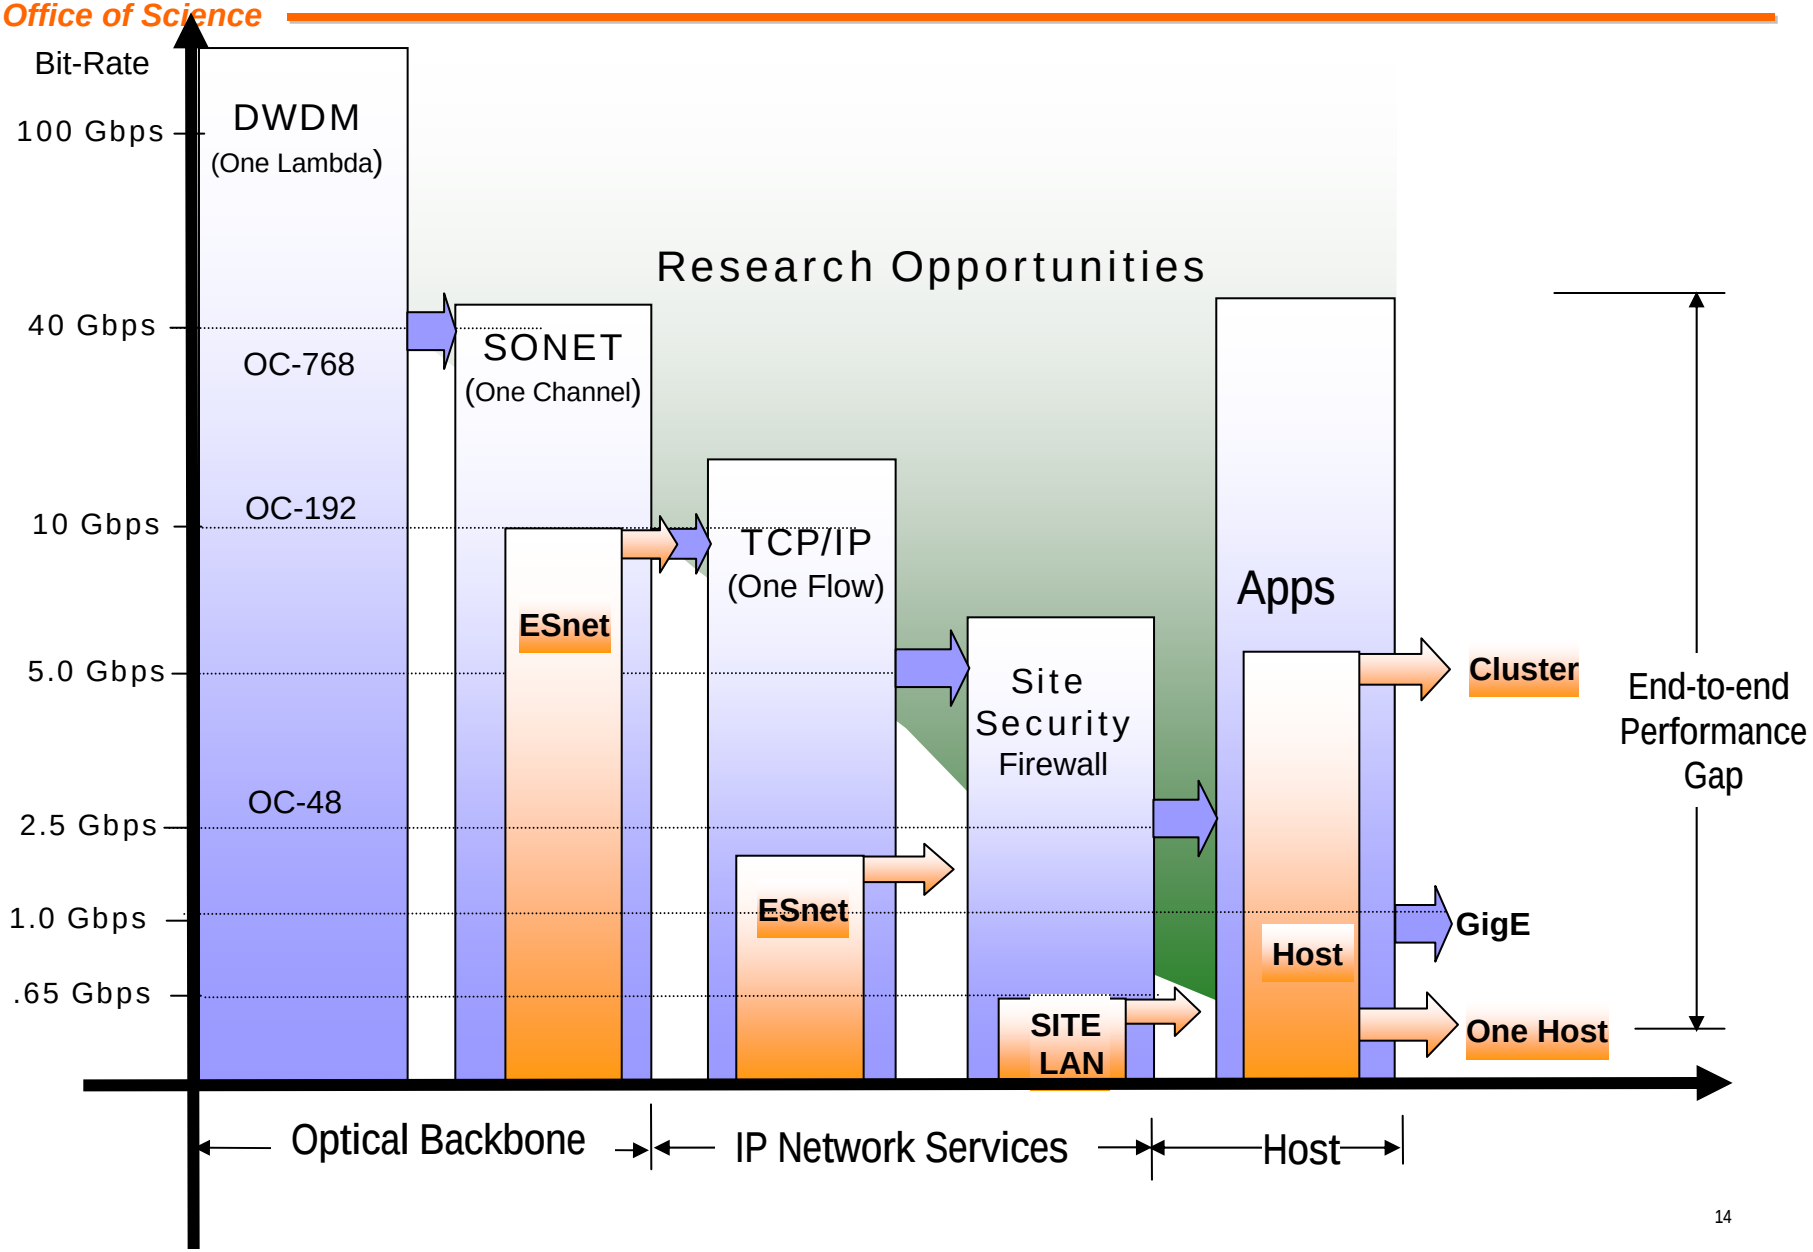

Opportunities

- Establish/strengthen strategic partnerships with other Office of Science programs
  - Genomes to Life, nanoscience, fusion energy
- Embark on a sustainable path to provide high-performance computers for science
  - industry partnerships
  - architecture research
- Provide scientific foundations in applied mathematics and computer science in areas that are barriers to world-leadership in computational science, e.g.
  - multi-scale mathematics
  - operating systems and programming environments
- Restore base research vitality to FY1992 levels
- Build on SciDAC success

Links

Mathematical, Information and Computational Sciences Program
http://www.sc.doe.gov/ascr/mics/index.html
Genomes to Life
http://www.doegenomestolife.org/
Nanoscale Science, Engineering, and Technology Research
http://www.sc.doe.gov/production/bes/NNI.htm
http://www.science.doe.gov/bes/Theory_and_Modeling_in_Nanoscience.pdf
UltraScale Simulation Planning
http://www.ultrasim.info/
Earth Simulator Home Page
http://www.es.jamstec.go.jp/esc/eng/
Other Federal high-performance computing programs

ASCI :	http://www.llnl.gov/asci/platforms/platforms.html
DARPA HPCS :	http://www.darpa.mil/ipto/research/hpcs/index.html
NSF PACI :	http://www.paci.org/
DOD Mod Office :	http://www.hpcmo.hpc.mil/

FY2004 Budget Request

Federal Government::	http://www.whitehouse.gov/omb/budget/fy2004/pdf/spec.pdf (see pg. 181)
DOE :	http://www.cfo.doe.gov/budget/04budget/index.htm
ASCR:	http://www.cfo.doe.gov/budget/04budget/content/science/science4.pdf

More Links

The International Technology Roadmap for Semiconductors (ITRS), 2002 Update. Semiconductor Industry Association. International SEMATECH:Austin, TX, 2002

http://public.itrs.net/Files/2002Update/Home.pdf