Office of Science
Science for America's Future

Mathematical, Information and Computational Sciences

An Overview of Computational and Network Resources

For the

Advanced Scientific Computing Advisory Committee

Washington, DC
May 2-3, 2001

Walter M. Polansky
Acting Director, MICS

Mission

Discover, develop, and deploy the computational and networking tools 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).

foster and support fundamental research in advanced scientific computing – applied mathematics, computer science, and networking

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

-- Expenditures --
Research vs. Facilities

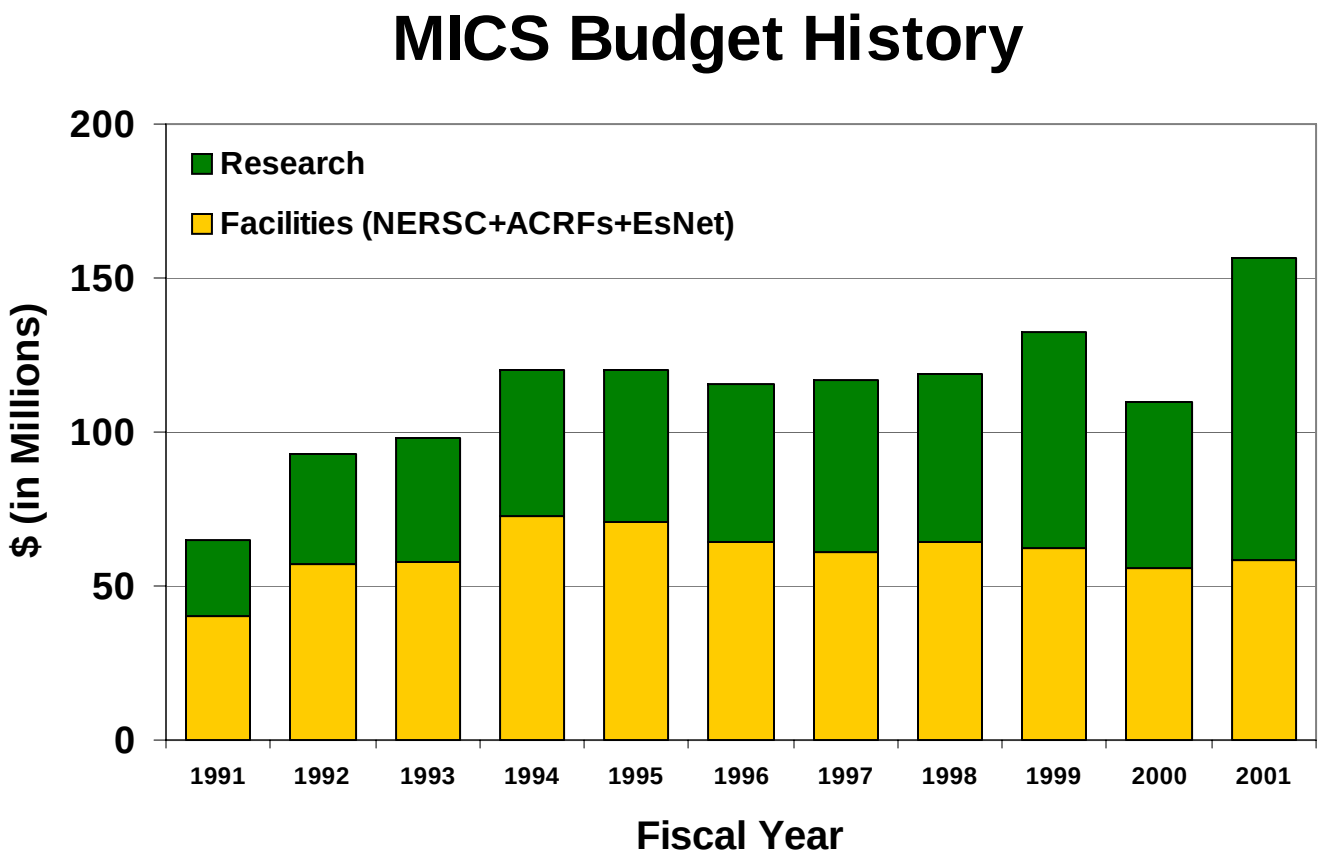

Computing Resources and Networking

Budget History

($ in millions)

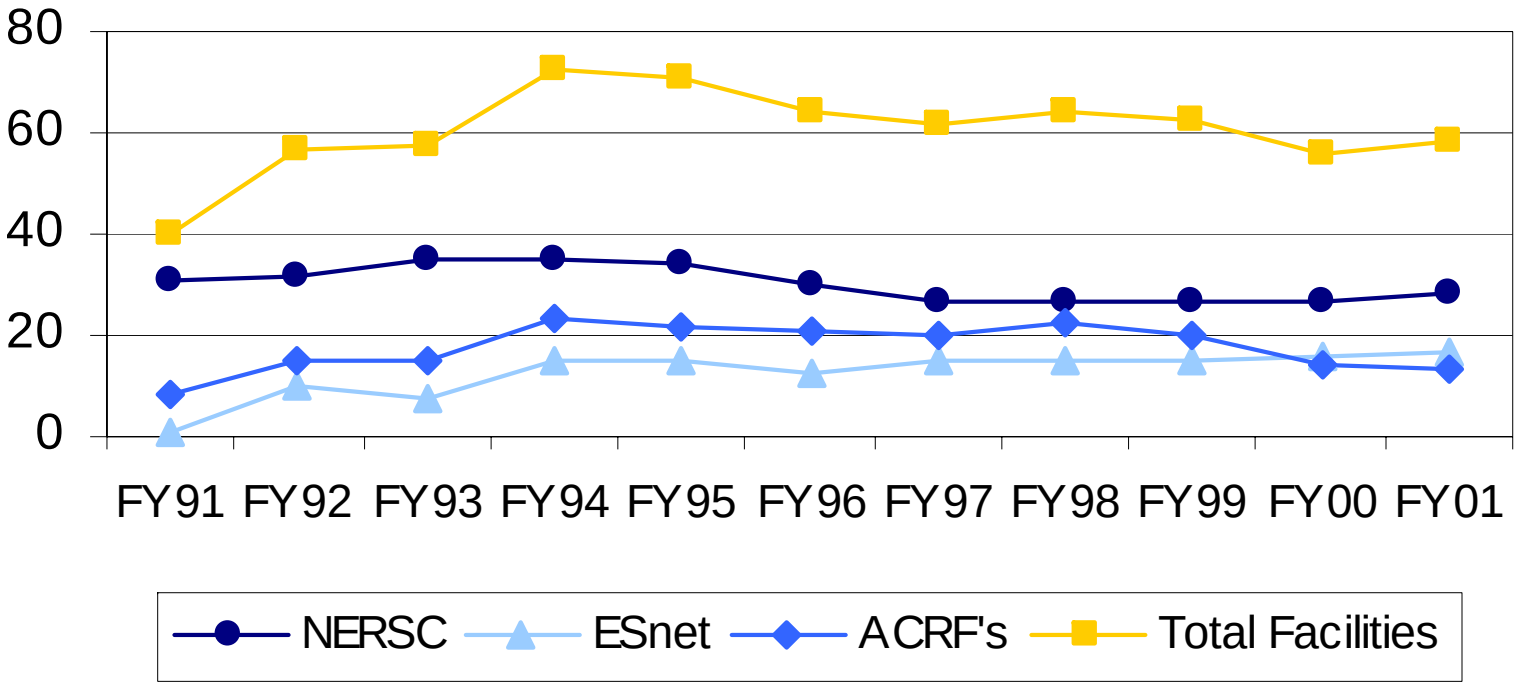

National Energy Research Scientific Computing Center

NERSC - Provides capability resources and professional user friendly services to computational scientists on projects within the missions of the Department of Energy

- Began in 1974 at LLNL as computing resource for magnetic fusion researchers
- Transferred to LBNL in 1996; moved to Oakland in 2000
- Provides open computing environment for nearly 2,400 users
- Offers several compute platforms (PVP & MPP); Terasacale; Technology refreshed approximately every 3 years
- Allocates compute resources competitively
  - 60% Office of Science program managers
  - 40% Committee appointed by NERSC director

Accomplishments Enabled by NERSC

- Turbulent Transport Reduction by Zonal Flows
  - 400 million particles; 5,000 time steps
  - Turbulence in tokomak plasmas can reduce turbulent transport
- Computational Accelerator Design
  - 500 million particles; predicted beam pipe aperature for SNS linac
- Kaon Decay Amplitudes
  - Reproduced observed effect; first time successfully computed
- Parallel climate model
  - Combines state-of-art models for atmosphere, land, ocean and sea ice
- Mechanisms of Enzyme Catalysis
  - Combined density functional theory & molecular mechanics to study catalysis pathways
- Modeling Molecular Processes in the Environment

Recognition of NERSC Users

- ACS Award for Computers in Chem. & Pharm. Res.- 2001
- APS Rahman Prize for Computational Physics- 2001
- Feynman Prize in Nanotechnology- 2000
- Presidential Early Career Award- 2000
- National Weather Service Modernization Award- 1999
- AIChE Award- 1998
- Gordon Bell Prize- 1998
- Hirschfelder Theoretical Chemistry Prize- 1998
- IEEE Computer Society's Sidney Fernbach Award- 1998
- Royal Society of Chemistry Spiers Medal- 1998
- Science Magazine's Breakthrough of the Year- 1998
- APS Prize for Achievement in Accelerator Physics and Technology- 1997

Advanced Computing Research Facilities

Objectives

Provide pioneer capability computing for scientific applications relevant to the Office of Science mission.

Provide Testbeds to examine critical CS issues

Accomplishments

Provided state-of-the-art computational resources for Grand Challenge calculations.

Evaluated feasibility of innovative computer architectures (IBM SP, SGI Origin 2000, Paragon, CM-5, Kendall Sq.) to meet SC computational needs.

Ongoing Projects

CHIBA City, a 512 CPU Linux cluster (ANL).

Falcon & Colt - Compaq AlphaServer SCs (ORNL).

Nirvana, 2048 Processor SGI (LANL)

“Probe” HPSS Testbed (LBNL/ORNL)

TERA Evaluation- (UCSD)

Plans

Explore novel architectures & testbeds

Evaluate for topical applications

Nurture/expand technical & vendor bases for future hardware purchases.

ACRF Evolution

- High Performance Computing Research Centers
  - FY1992- Established High Performance Computing Research Centers at LANL & ORNL (global climate/materials)
  - FY1993- Initiated Grand Challenge computational research program
  - FY1995- Established an HPCRC at ANL (applications testing/computer science)
- Advanced Computing Research Facilities FY95-00
  - Upgraded hardware at LANL
  - Focused ORNL and ANL efforts
  - Coupled Grand Challenges to specific ACRFs
  - Allocated portion of NERSC for Grand Challenges
- FY 2000- Completed Grand Challenges

ACRF Accomplishments

- Developed GUSTO (Globus Ubiquitous Supercomputing Testbed)- a prototype for next-generation networking
  - High Performance Computing Challenge Award- SC98
  - Global Information Infrastructure Award- 1998
- Introduced novel approaches for Parallel I/O in portable system ROMIO
  - Implemented on all ASCI machines
  - Incorporated into HP and Silicon Graphics products
- Combined high-speed parallel computing with 3D virtual reality graphics for pollution control applications
  - DOE Young Scientist Award- 1998
- Developed/Enhanced the High Performance Storage System
  - R&D 100 Award- 1997
  - Production system operational throughout DOE-sites

ESnet Overview

Objective-

Provide highly capable and reliable communications infrastructure and leading-edge network services that support Office of Science and the Department's missions

Program Elements-

- Nation-wide high-performance research network
- Advanced network services to support science in DOE
- Cooperation between DOE, ESnet Mgmt & End Sites
- Extensive structure of domestic and international interconnects
- Advanced Technology and Research program

ESnet - History

- 1986 – 1991: Pre- HPCC
  - ESnet becomes official network for Energy Research (now Office of Science)
  - Backbone - a combination of 9.6K and 56K satellite and terrestrial circuits, DECnet & TCP/IP
  - 22 sites connected to T1 Backbone in 1991
- 1992-1994
  - Competitive RFP launched for “Fast Packet ESnet Services”
  - Contract Awarded to Sprint in 1994: T3 and beyond, ATM
- 1995 - 1999
  - ESnet Operations recompeted in 1995; operation moved to LBNL
  - Traffic increases from 2.4 to over 8.7 Tbytes/month
  - T3 and OC3c ATM based interconnects established, over 50 sites connected
- 2000 ESnet: new contract with Qwest

Office of Science Mission Requirements Drive Computer Resource & Network Investments

- Base SC Research Programs
  - Distributed terabyte files; rapidly increasing in size
  - Growing needs for high-performance computer & network resources

- Scientific Discovery through Advanced Computing- FY01- $60 M ($37M in MICS)
  - Office of Science initiative- scientific challenge codes; computing systems & mathematical software; collaboratory software infrastructure; scientific computing hardware infrastructure
  - ACRF and NERSC resources will be utilized at early phase

- Biotechnology
  - Microbial Cell (NFR 01-20)
  - Computational Biology (NFR 01-21)
  - Genomes to Life

Scientific Computing Infrastructure Hardware

- Strategy for high-performance computing hardware
- SciDAC Program Element
  - Flagship facility : general purpose, production (NERSC)
  - Topical Center(s) : architecture tailored to specific types of simulations (TBD)
  - Experimental Computing Facilities (TBD)
  - High speed communications network (ESnet)
- Flexible and robust- could be adapted for compute needs of biotechnology

Opportunities in Biotechnology

- Biological systems directly convert energy into useable forms
- Biological systems employ low energy reactions
- Biological systems adapt to their environment; they organize, replicate, & repair themselves
- Behavior of biological systems is not described from first principles

Computational Needs- Biology & Bioinformatics

Problem Component	Computing Speed	Storage
Genome Assembly	>10 TeraFlops sustained to keep up with expected sequencing rates	300 TB of trace files per genome
Protein Structure Prediction	>100 TeraFlops per protein set in one microbial genome	Petabytes
Classical Molecular Dynamics	100 TeraFlops per DNA-protein interaction	10s of Petabytes
First Principles Molecular Dynamics	1 PetaFlops per reaction in enzyme active site	100s of Petabytes
Simulations of Biological Networks	>1 TeraFlops for simple correlation analyses of small biological networks	1000s of Petabytes