Mathematical, Information and Computational Sciences

An Overview
for the
ESnet Steering Committee

March 18, 2003
Bethesda, MD

Walt Polansky

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

U.S. Department of Energy

Office of Science

Program Execution FY2004

Basic Research

Research to enable…

…simulation of complex systems

…distributed teams, remote access to facilities

…Applications

BES, BER, FES, HEP, NP

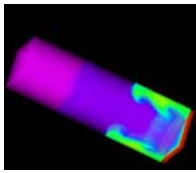
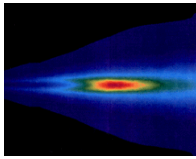
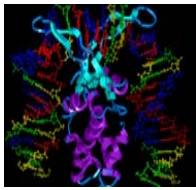
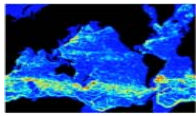
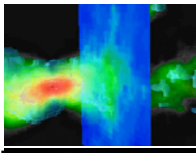
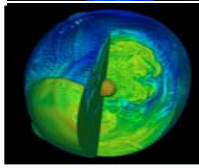
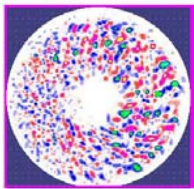

- Applied Mathematics
- Computer Science
- Network Environment
- Scientific Applications
- Genomes to Life

SciDAC!

- Nanoscience
- Grid enabling research
- Integrated Software Infrastructure Centers

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

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

Next Generation Architecture

High Performance Computing and Network Facilities for Science

National Energy Research Scientific Computing Center (NERSC)

Advanced Computing Research Testbeds

Energy Sciences Network (ESnet)

Budgets

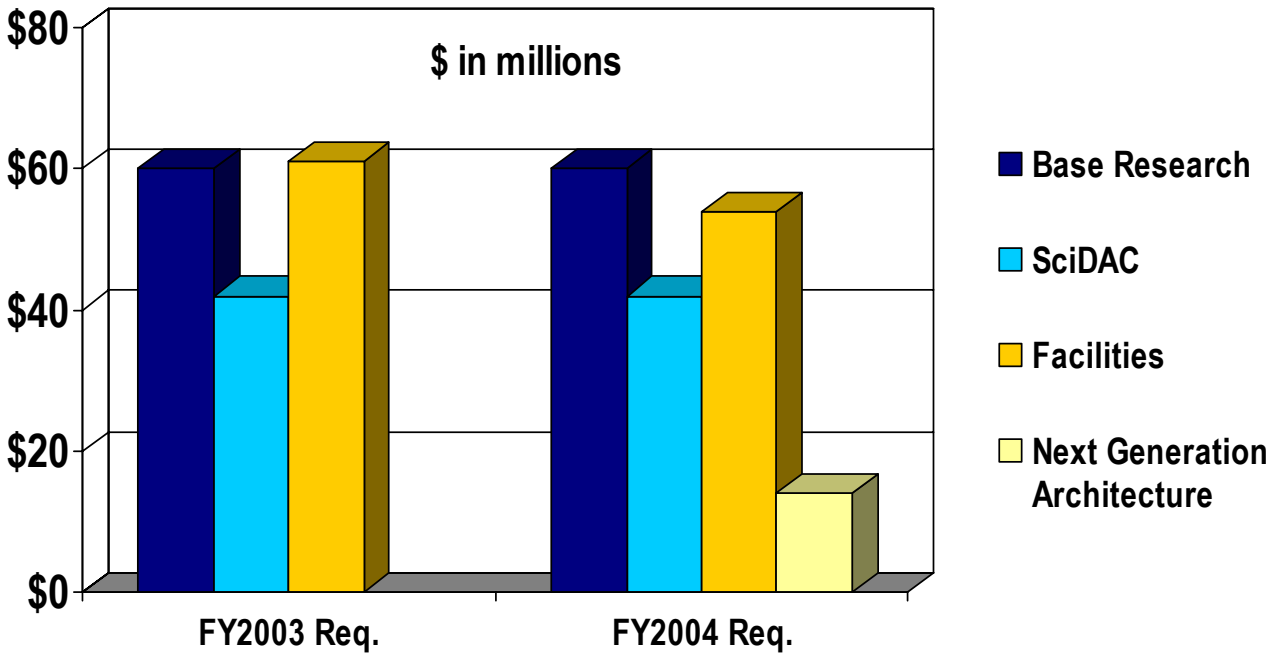

Fiscal Year	Request	Appropriation
2002	$156,170,000	$147,159,000
2003	$163,557,000	$164,480,000*
2004	$170,490,000	TBD

* Following General Reduction & Omnibus Rescission

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.

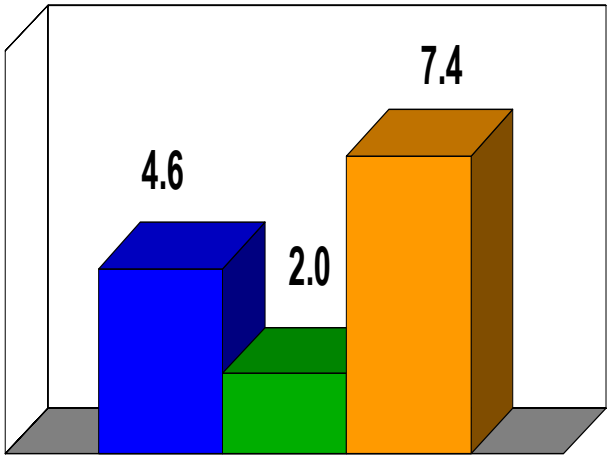

Research- Operating Systems/Runtimes
Research- Application Performance/Scaling
Evaluation- Cray X1

U.S. Department of Energy

Office of Science

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

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

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

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…

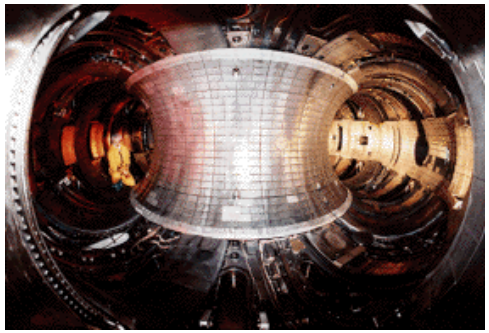

Design Refinement

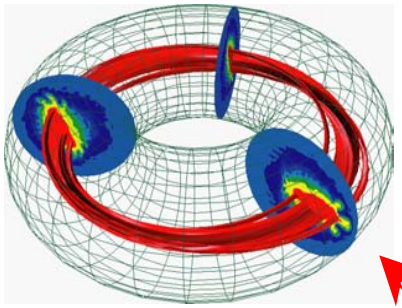

Optimization- advanced optimization theory, efficient run ensembles.
Predictability- error estimation, parameter estimation, model reduction

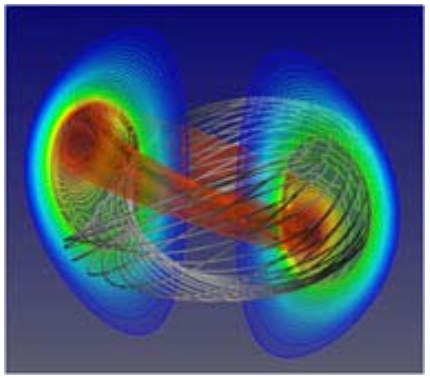

Mathematical description- scientifically accurate across multiple scales, with proper boundary conditions.

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

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

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

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

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

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

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

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

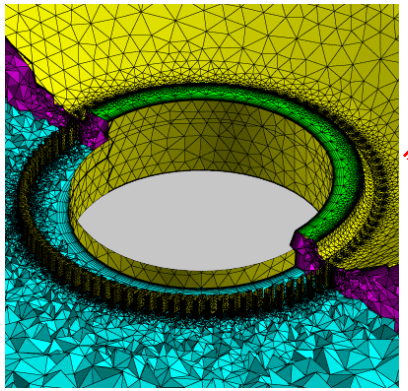

Computational solution- high-performance computing; accurate, efficient, scalable, tunable, robust, modular, and fast numerical algorithms

Network Environment

- Establish themes for Office of Science network environment “story”
- Compare and contrast ESnet with similar networks.
- Document present utilization of ESnet.
- Identify major bottlenecks in networks.
- Sketch a 5-yr. vision for networking environment.

Global Infrastructure for Science

Supercomputing and Large-Scale Storage

experiment

theory

simulation

High Speed Networks

Computing and Storage of Scientific Groups

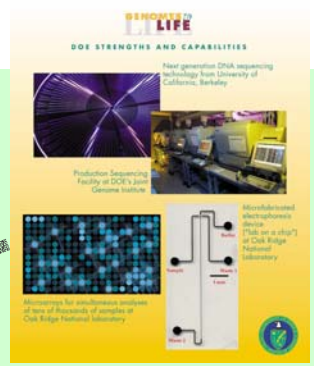

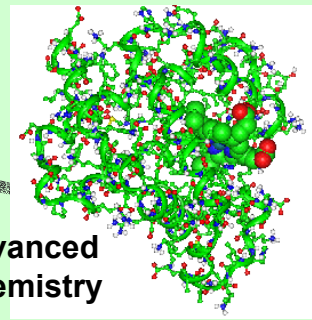

Advanced Chemistry

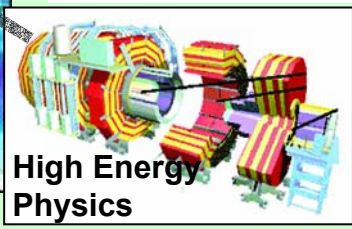

High Energy Physics

Advanced Engine Design

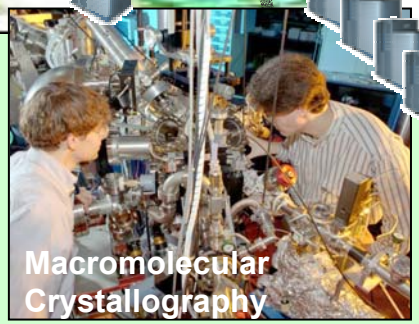

Macromolecular Crystallography

Advanced Photon Source

Spallation Neutron Source

Office of Science
Distributed Resources, Distributed Expertise

Supercomputing and Large-Scale Storage

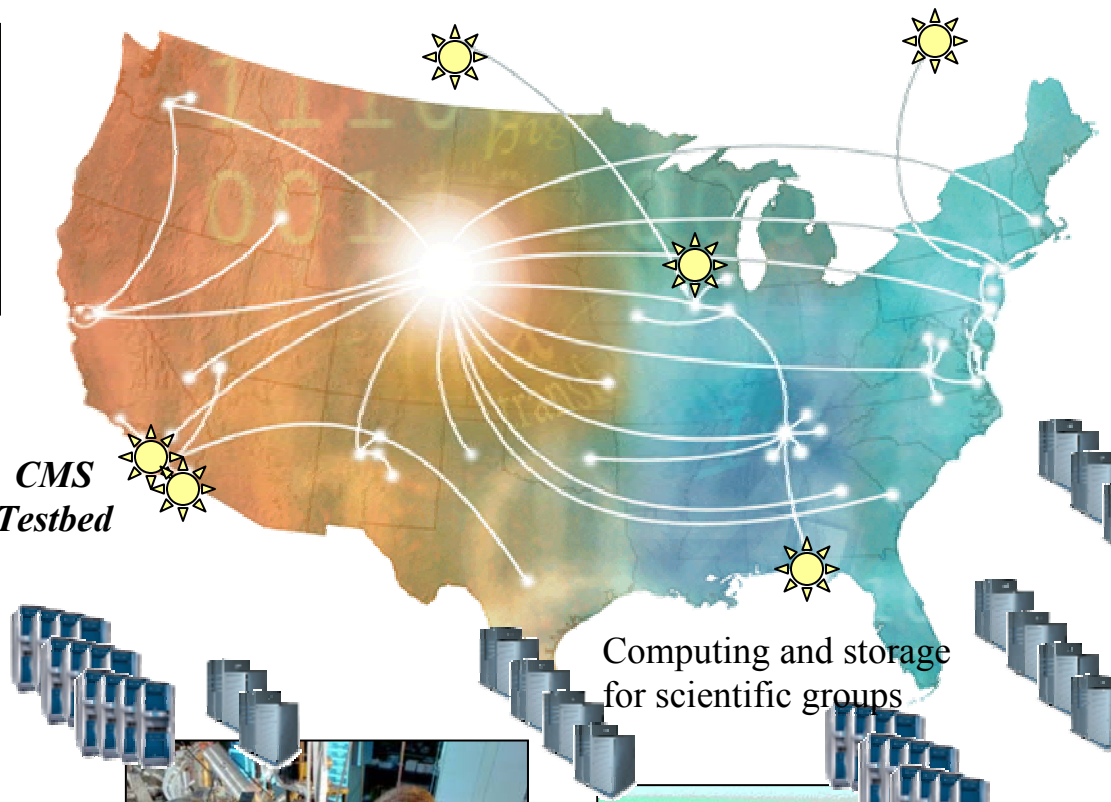

CMS Testbed

Computing and storage for scientific groups

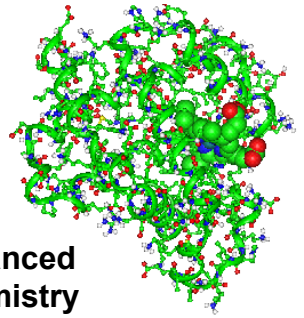

Advanced Chemistry

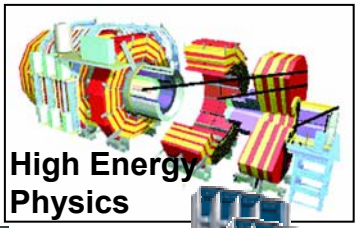

High Energy Physics

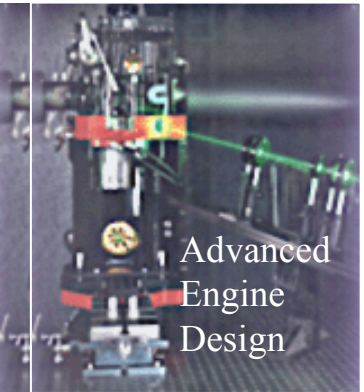

Advanced Engine Design

Macromolecular Crystallography

Advanced Photon Source

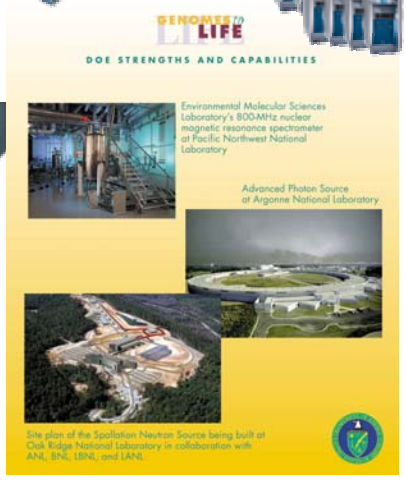

Science Applications- Bit Rate Budget
End-to-end

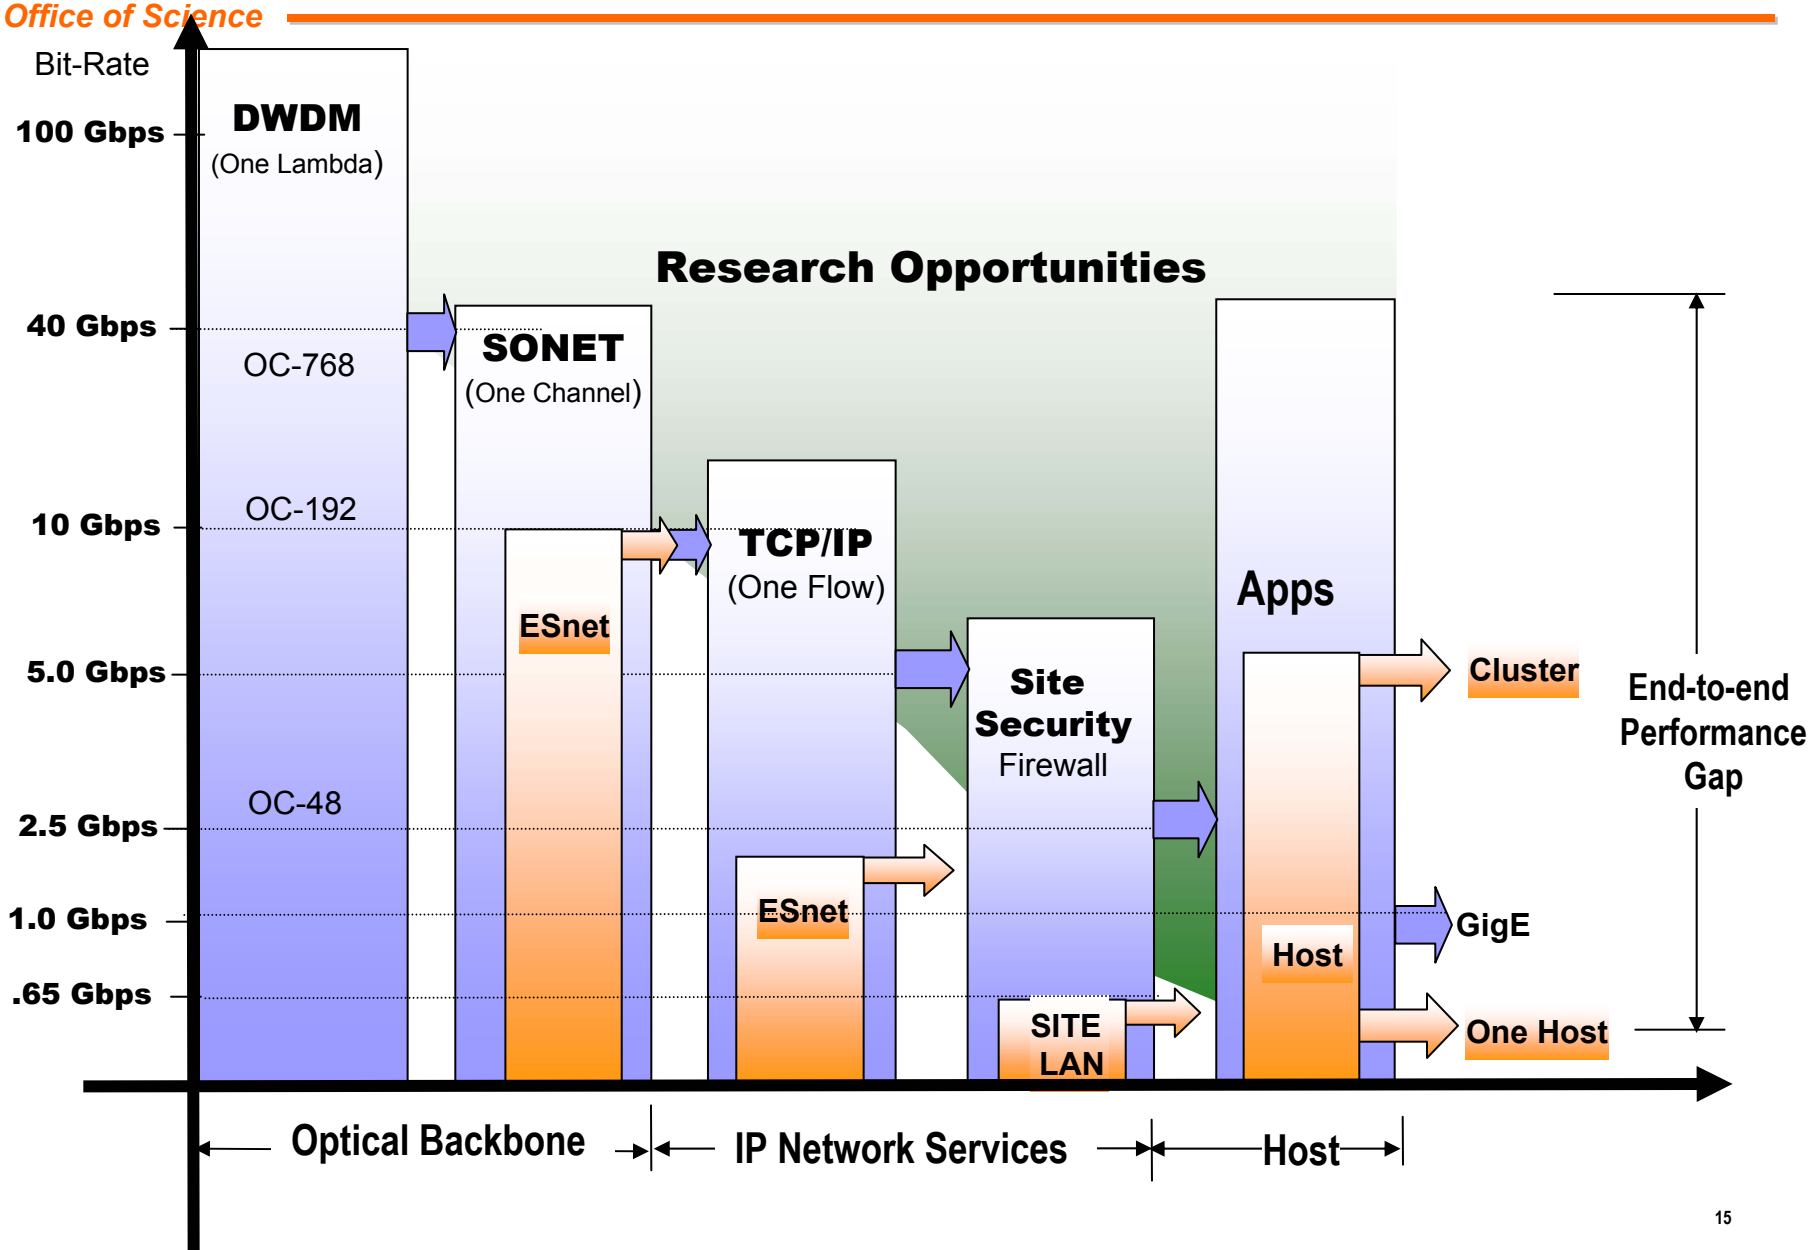

Transmission Control Protocol Internet Protocol (TCP/IP)

Features

- Reliable end-to-end transmission control protocol for the Internet
- Error control algorithm, CRC-32, optimized for 1.5KB packet size
- TCP stack resides in host
- "Fair" to all users; Prohibits aggressive use of the network
- Constraints on throughput
  - [1] Congested network: 0-15% of link capacity
  - [2] Uncongested network: up to 30% of link capacity
- Used for FTP

Issues

- Maximum throughput ~13 Gbps
- Maximum number of packets 2^{32}
- Designing effective congestion mechanism require router modification
- Increasing packet size requires new error control mechanism
- Operating system bypass needed to reduce host congestion

Firewalls

Features

- Monitor, process and filter ports and packets, in real time
- Compute intensive
- Utilized for cyber protection

Issues

- DOE Cyber Security Program Plan Guide, DOE G205.1X (Draft), 2.3a. Perimeter Protection- "…describe approaches used to secure the perimeter (e.g. firewalls and other intrusion detection systems)".
- Commercially available at OC-12; OC-48 firewalls are rare
- OC-192 not in the pipe line

Network Needs Vary Widely…

…a partial list

Community	Goal	Network services needed	Performance requirements
HENP, Climate	move a petabyte in 24 hours	Secure file movement	160-200 Gbs, "best effort" service
Supernovae, Cell Biology	shared immersive environments	multicast, latency and bandwidth guarantees	2.4 Gbs with strong QoS guarantees
Remote data access, remote computing	real time exploration of remote datasets	secure remote connectivity	1-10 Gbs to desktop, modest QoS needs
Grid and network research	experimental, non-production use	remote connectivity, access to hardware	unused lambdas for experimentation

U.S. Department of Energy

Office of Science

Major R&D Network Characteristics

Network	Abilene	DTF/ETF	ESnet	DFN	GEANT
Mission	Production; U.S. academic community	"Backplane" testbed	Production; DOE science community	Production; German science community	Production; pan-European
Transport	OC192; 1,000 TB/month	4x OC192	OC192 (N); OC48(S) 100-200 TB/month	OC48 1000 TB/month	OC192 1,000 ++ TB/month
Funding Model	Site-connect fees, commercial donations	NSF	DOE-SC; no site connect fees	Site connect fees; some govt. funding	Site-connect fees; some govt. funding
External Connectivity	International; No commercial traffic	Member sites only	Commercial and international	Commercial and international	Commercial and international
Services	24x7 NOC.	None.	24x7 NOC, video, data, audio, PKI, directory.	24x7 NOC	24x7 NOC
Security Services	Some.	No.	Some, expanding	No available information	"CERT-like"
Performance data	Extensive	None, yet.	Limited.	Limited.	Limited.
Connections	~ 50 Gigapops directly connecting 200+ universities	~6-12 high-capability systems	~35 sites; DOE National labs and collaborators	~550 Universities and laboratories	~ 30 European National networks

Network Environment Vision

Features	2003	(2008)	Actions
End-to-End throughput	~ 0.1- 0.5 Gbps to OC-48	~ 35 Gbps to OC-192	
Protocol	TCP/IP	Base users-TCP/IP High-impact- ???	New strategy; Research investments
Cybersecurity (Firewalls, IDS, authentication, access control etc.)	OC-12	OC-192; some OC-768	Technology advances; Policy review (?)
Middleware services	Ad hoc; available for limited resources	Ubiquitous grid services	Research investments
QoS, on-demand Services	None	Wavelength on demand	Research investments
Network Monitoring	Ad-hoc; limited number of sites	Systematic; all sites	Establish & implement new requirement
Backbone Network	Static Provisioning	Dynamic Provisioning	New business model

Candidate Options - Business Model

- Growing recognition that TCP/IP-based networks only meets some needs
- Advent of wavelength services (DWDM)
- Networking slump; so industry may be eager to explore alternatives

Vendor supplies bandwidth	Current model; simple and complete	Limted options for QoS and network/grid reseach; expensive?
Lease lambdas from vendor	Flexible; allows differentiated service	Undertake responsibility for routing; cost-effective?
Purchase dark fiber	Most comprehensive	Stresses DOE-system; too much capability?

U.S. Department of Energy

Office of Science

An ESnet "slice" to address MICS program needs

ESnet Backbone Traffic (1994-2002)

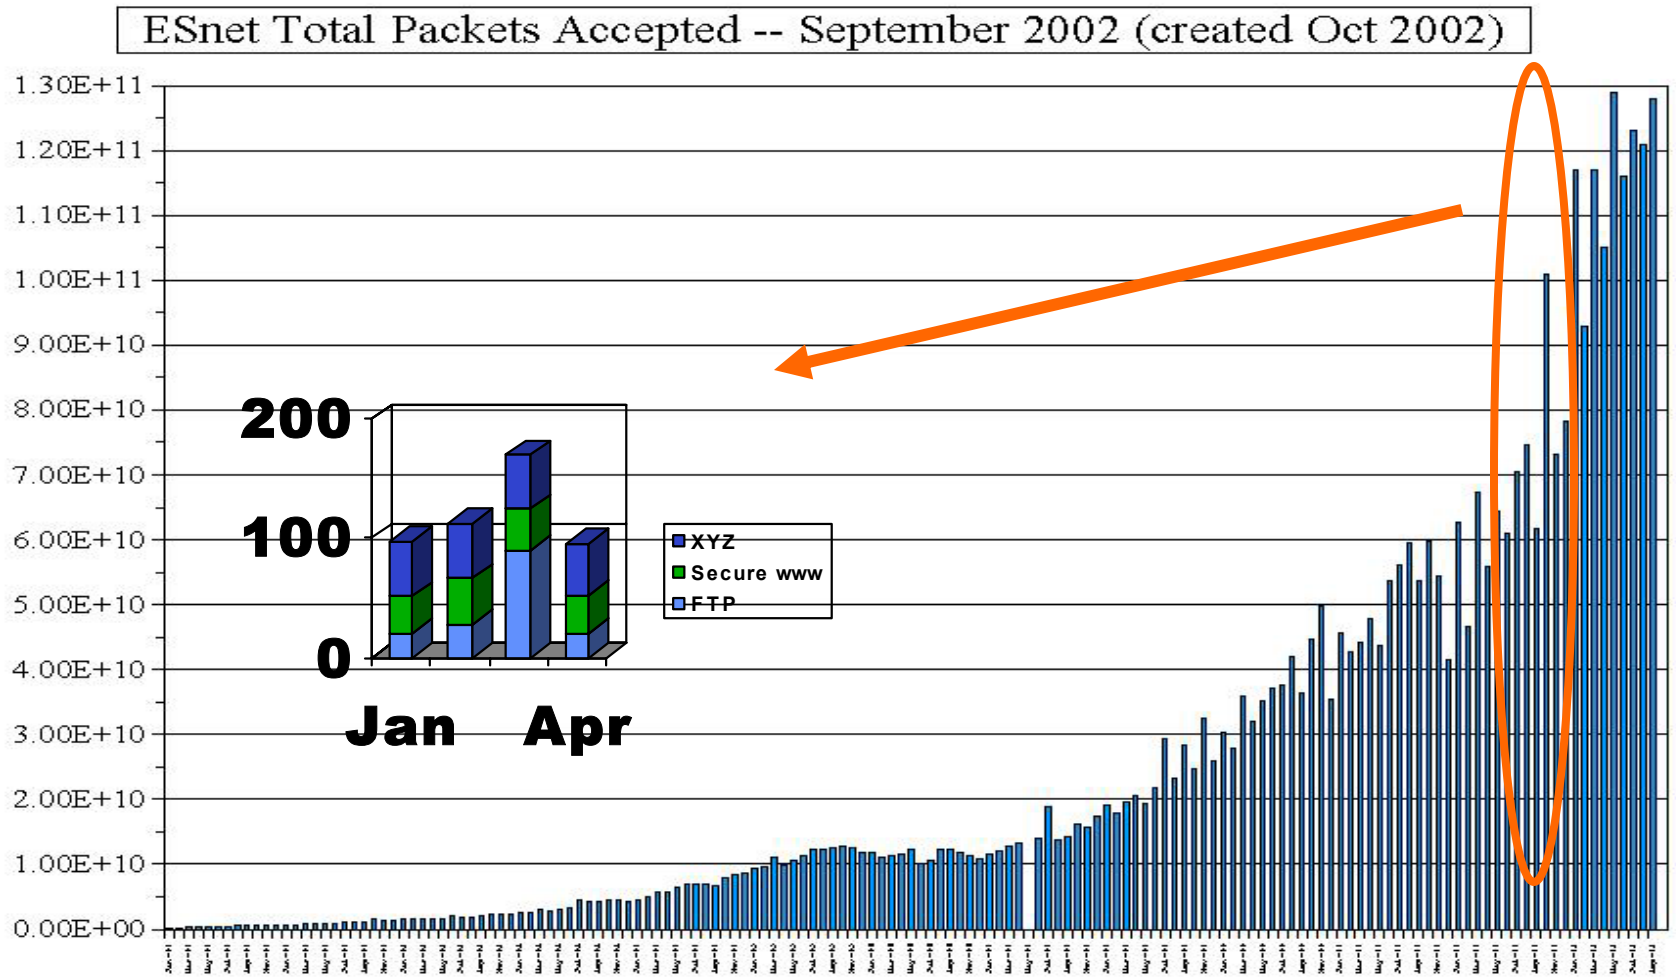

Network Transit Times

Another "look" to address MICS program needs

Transfer A - B	Network	Result	Objective
1 bit	Empty	Calculated	Establish minimum time
1 packet	Empty	Calculated	Establish minimum TCP overhead
Files (benchmarks)	Actual conditions	Table of experimental results	Quantify performance

Summary

- Network environment in 2008, dramatically different from the present.
- Current description of the network environment fails to capture special nature of Office of Science research endeavors
- Data on utilization/performance of ESnet
  - System: TB/mo., bandwidths, configurations
  - Local: limited number of sites, data not quantified
  - Does not meet MICS-program needs (What about users needs?)
    - Type of traffic (FTP, video, etc.);
    - Origin of traffic (BER, BES, FES, HENP)
- Detailed comments (including revised ppt slides) are needed and are welcome.
- First opportunity to tell the story- Congressional Hearings March, 2003