Department of Energy
Office of Science

2005 DOE Annual PI Meeting
Brookhaven National Lab
September 27-30, 2005

Thomas Ndousse
Program Manager MICS Division
Office of Science, US Department of Energy

U.S. Department of Energy's Office of Science

OFFICE OF SCIENCE

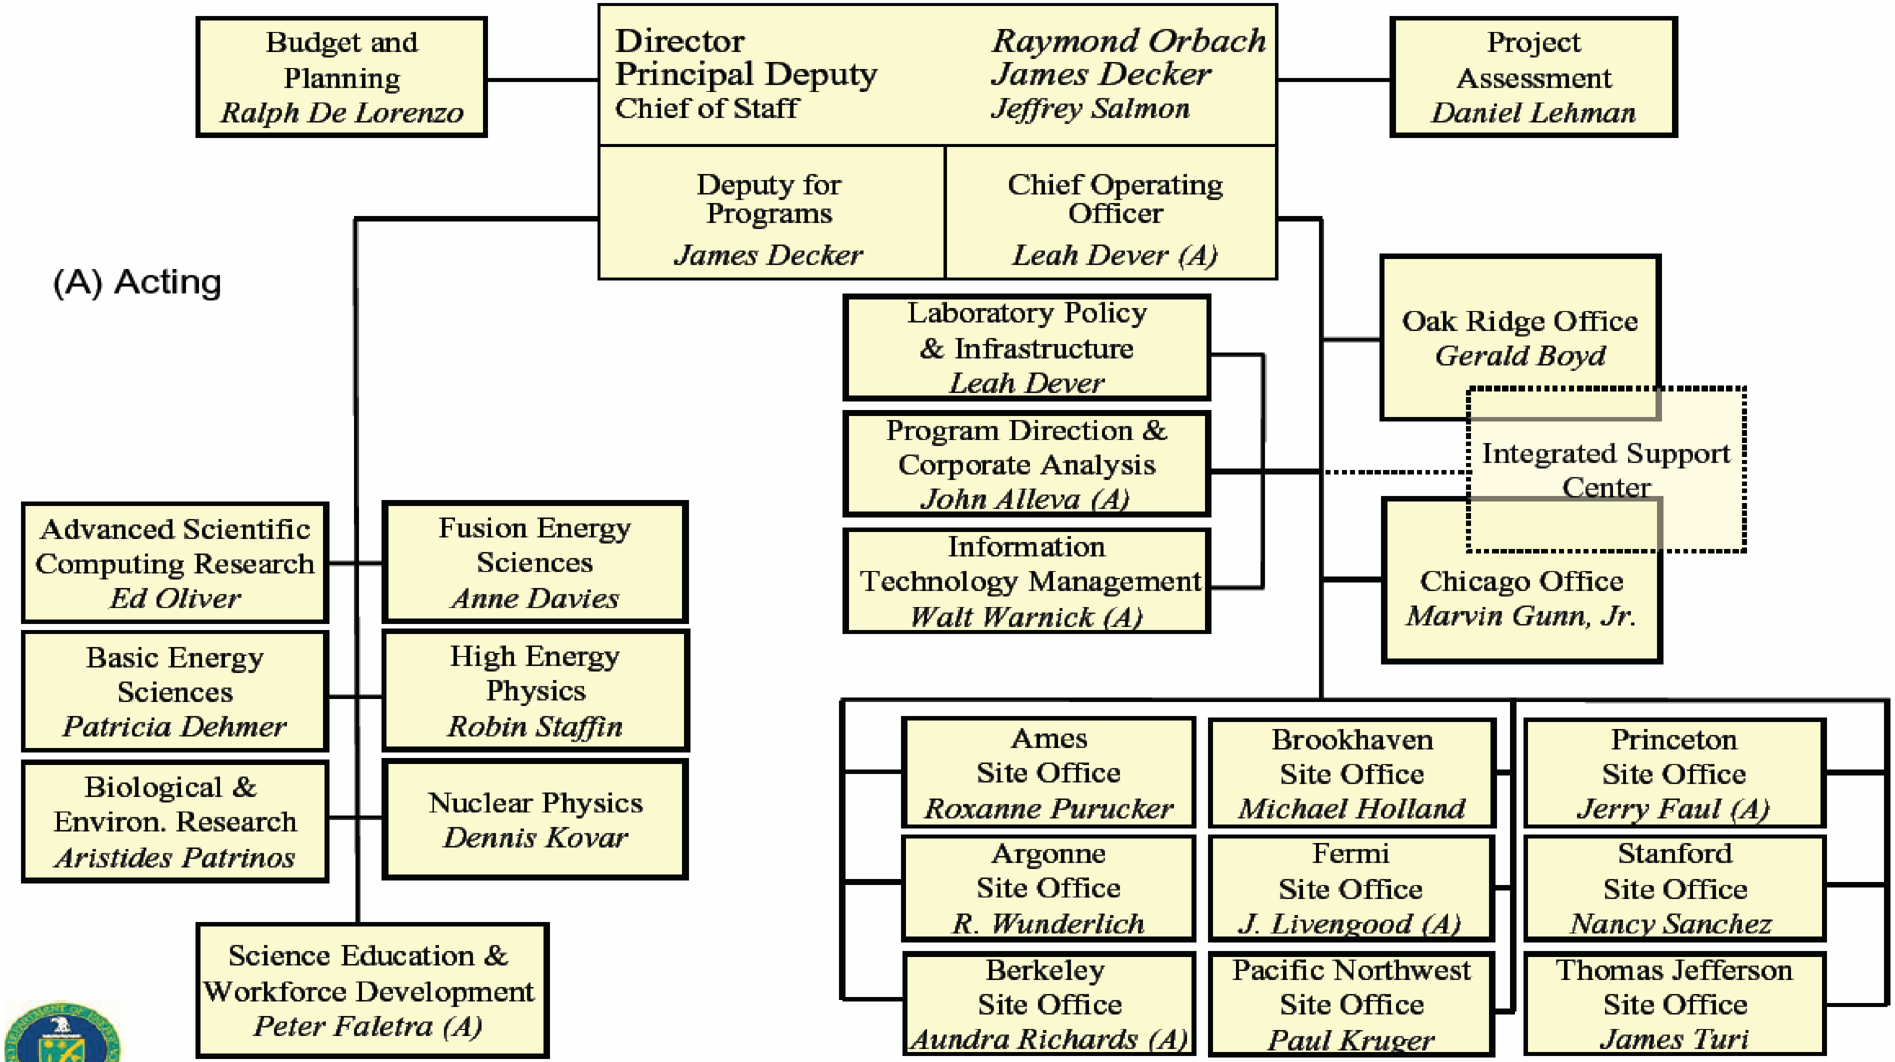

March 18, 2005

U.S. Department of Energy's Office of Science

Deputy for Programs Structure

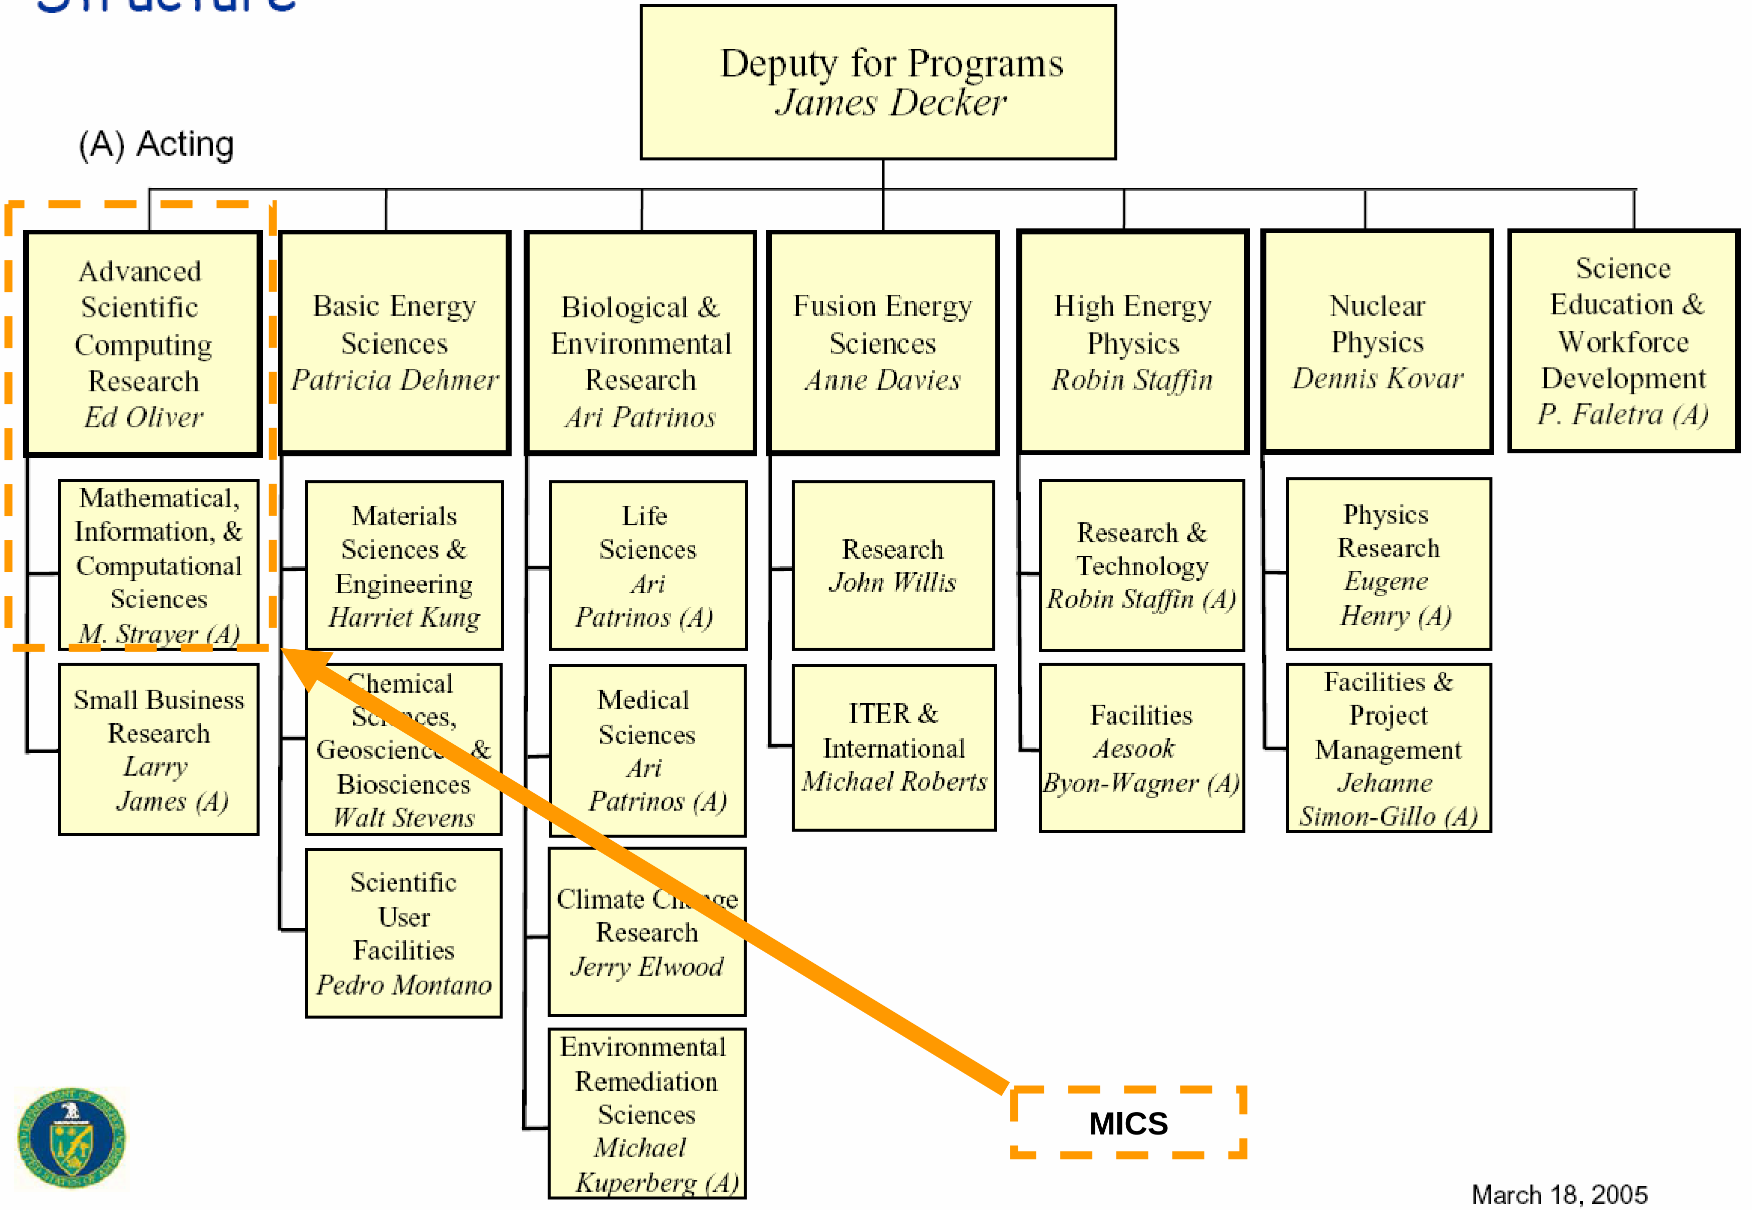

March 18, 2005

Advanced Scientific Computing Research (ASCR)

Mission

Deliver forefront computational and networking capabilities to scientists nationwide that enable them to extend the frontiers of science, answering critical questions that range from the function of living cells to the power of fusion energy

Advanced Scientific Computing Research (ASCR)

Strategies

- **R&D leadership in core technologies and computing sciences**
- **First-class computing and networking facilities**
- **In-house development with vendors and network carriers**
- **Leverage other federal agencies research effort**
- **Partnership with application scientists**

ASCR Program Overview
Budget: $300M

ASCR | **BES, BER, FES, HEP, NP**

Underlying Technologies → Fundamental Research → SciDAC → Partnerships → Science

Underlying Technologies
- Hardware
- Software
- Networking

Enabling Science through Computing
- Applied Mathematics
- Computer Science
- Advanced Computing Software Tools / Integrated Software Infrastructure Centers
- Scientific Application Partnerships

→ Applications, e.g.
- Materials Sciences
- Chemical Sciences
- Combustion Modeling
- Biology
- High Energy Physics
- Nuclear Physics
- Fusion Energy
- Global Climate

High Performance Network Environment for Science
- **Network Research and Testbeds**
- Collaboratory Tools
- Collaboratory Pilots

→
- Access to Facilities
- Link Researchers

Next Generation Architecture
- **Advanced Computing Research Testbeds**
- Research and Evaluation Prototypes
- **Leadership Class Computing**

→
- Leadership in High End Computing

Core Facilities
- **High Performance Production Computing Capability**
- **Energy Sciences Network (ESnet)**

PI Meeting Goals

- **Opportunity for researchers to learn about DOE network research portfolio**
- **To enable PIs to network, exchange ideas, and develop research collaboration**
- **Provide a forum to discuss technology transfer strategies to advance DOE science mission**
- **Learn about new directions and updates in DOE avanced networking requirements**

PI meeting Participants

- DOE MICS & SciDAC Network Research PIs
- MICS Career Awards (ECPI) network research PIs
- DOE SBIR/STTR Network Research PIs
- Scientific Application Developers PIs
- Scientists with Data and Distributed Supercomputing problems
- DOE HQ Program Managers

Meeting Format

- **Each project has 10 minutes presentation follow by 5 minutes Q&A**
- **General discussion on the new direction DOE network requirements**

High-Performance Network Research

Mission

- Research, develop, test, and deploy high-performance and high-capacity network technologies to support distributed high-end science applications

Goal

- Provide DOE researchers anywhere with leading-edge network capabilities to securely access computing resources and science facilities

Distribution of Funds

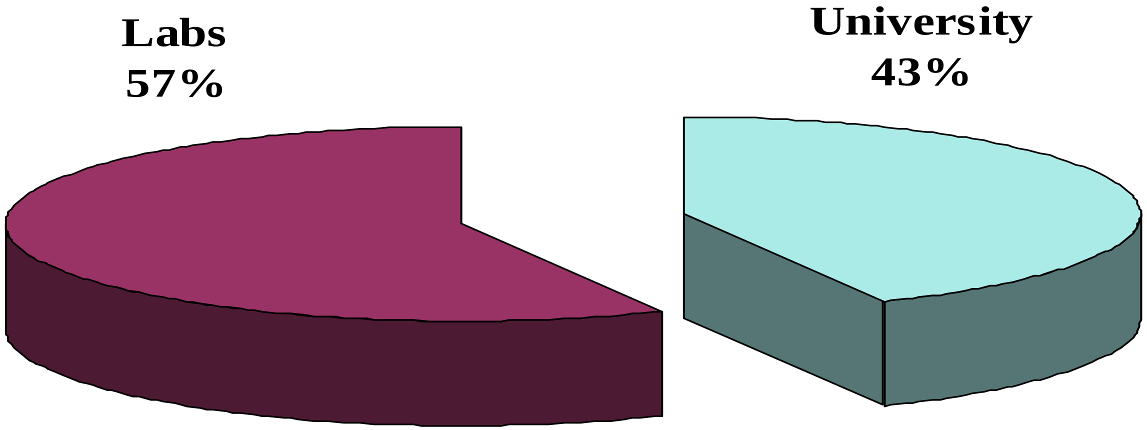

DOE Labs

- Argonne National Lab
- Brookhaven National Lab
- Stanford Linear Accelerator
- Oak Ridge National Lab
- Berkeley National Lab
- Pacific National Lab
- Los Alamos National Lab
- Sandia National Lab
- Livermore National Lab
- Jefferson Lab

University

- Any University with Graduate Program

Network Requirements for Large-scale Science

Science Areas	Today End2End Throughput	5 years End2End Throughput	5-10 Years End2End Throughput	Remarks: Basic research, testing and deployment
High Energy Physics	0.5 Gbps E2E	100 Gbps E2e	1 Tbps	high throughput
Climate Data & Computations	0.5 Gbps E2E	160-200 Gbps	Tbps	high throughput
SNS NanoScience	does not exist	1Gbps steady state	Tbps & control channels	remote control & high throughput
Fusion Energy	500MB/min (Burst)	500MB/20sec (burst)	Tbps	time critical transport
Astrophysics	1TB/week	N*N multicast	1TB+ & stable streams	computational steering & collaborations
Genomics Data & Computations	1TB/day	100s users	Tbps & control channels	high throughput & steering

Network Requirements for Large-scale Science - II

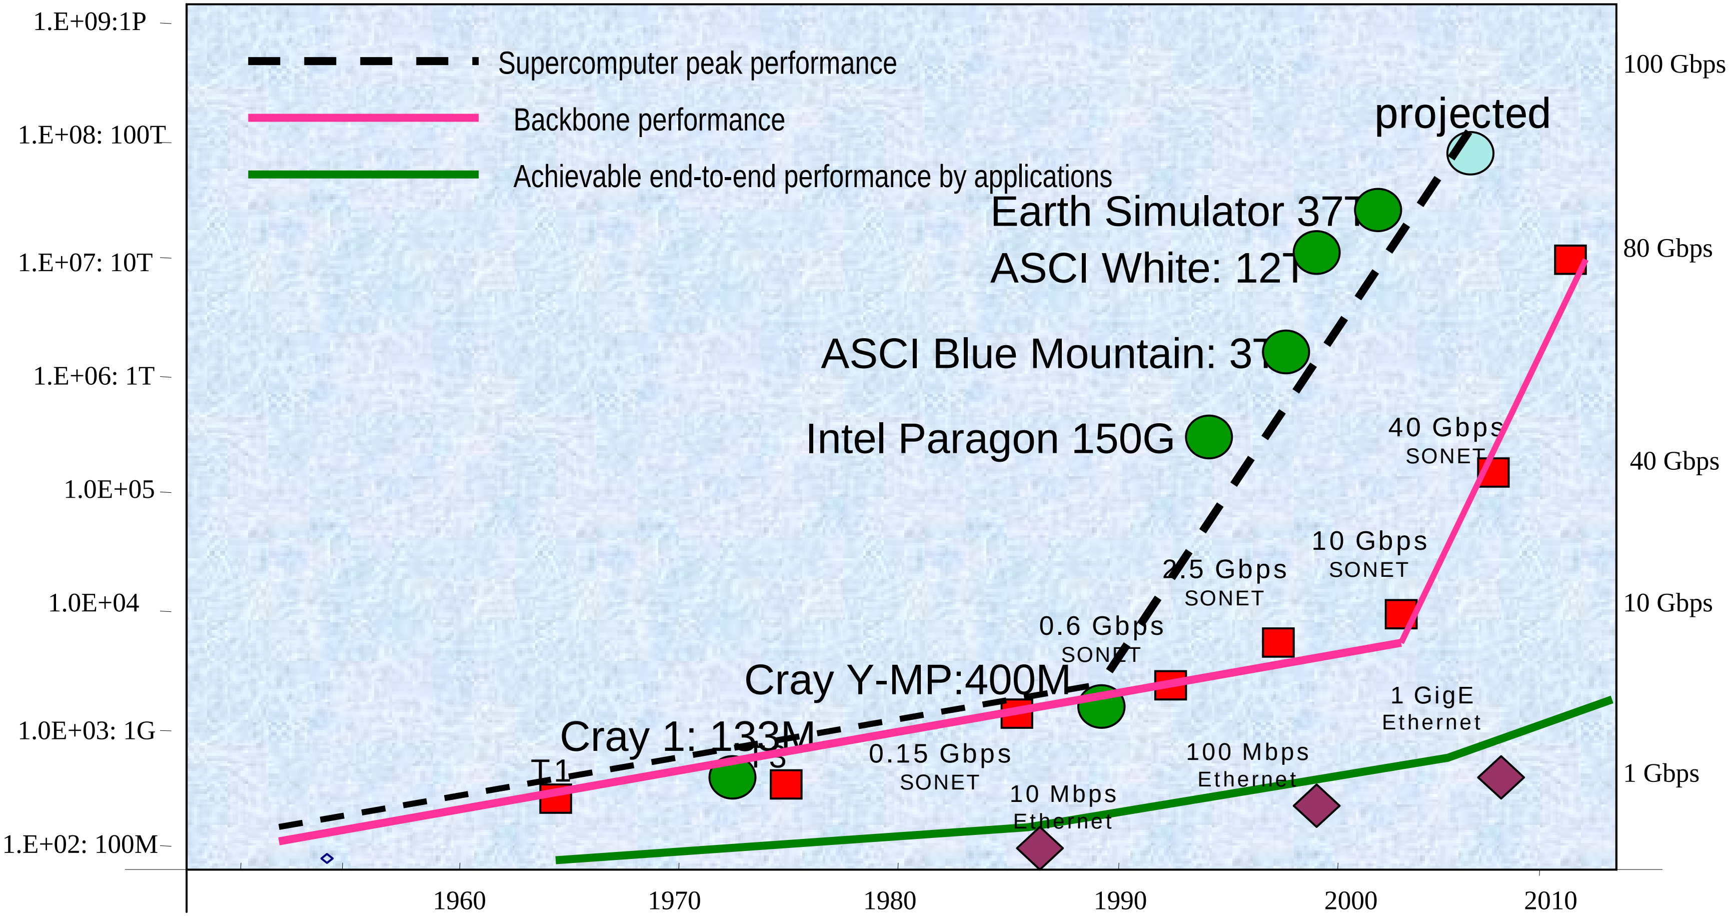

<u>Rule of thumb</u>: The bandwidth must be adequate to transfer Petabyte/day

~ 200Gbps - <u>NOT</u> on the evolutionary path of backbone, much less application throughput

Network Capabilities for Science applications

- Advanced Networks to support leadership class supercomputers (100T)
- Multi-Gigabits/sec to distributed science applications
  - High-speed data transfer
  - Distributed remote visualization
  - Tele-instrumentation
  - Distributed supercomputing
- Non-intrusive cyber security for open science environment
- Effective and agile utilization of network bandwidth
- Guaranteed end-to-end performance to scientific applications

Challenges

- How far can IP go?
- Is TCP and UDP up to the task for terabits networks?
- Packet-switching, circuit-switching, burst switching, or combination?
- How to harness the abundant optical bandwidth
- Will Existing network components scale efficiently at Terabits/sec
- The case for optical buffers and processors
- High-performance middleware

Strategy:
Composable protocol Stack

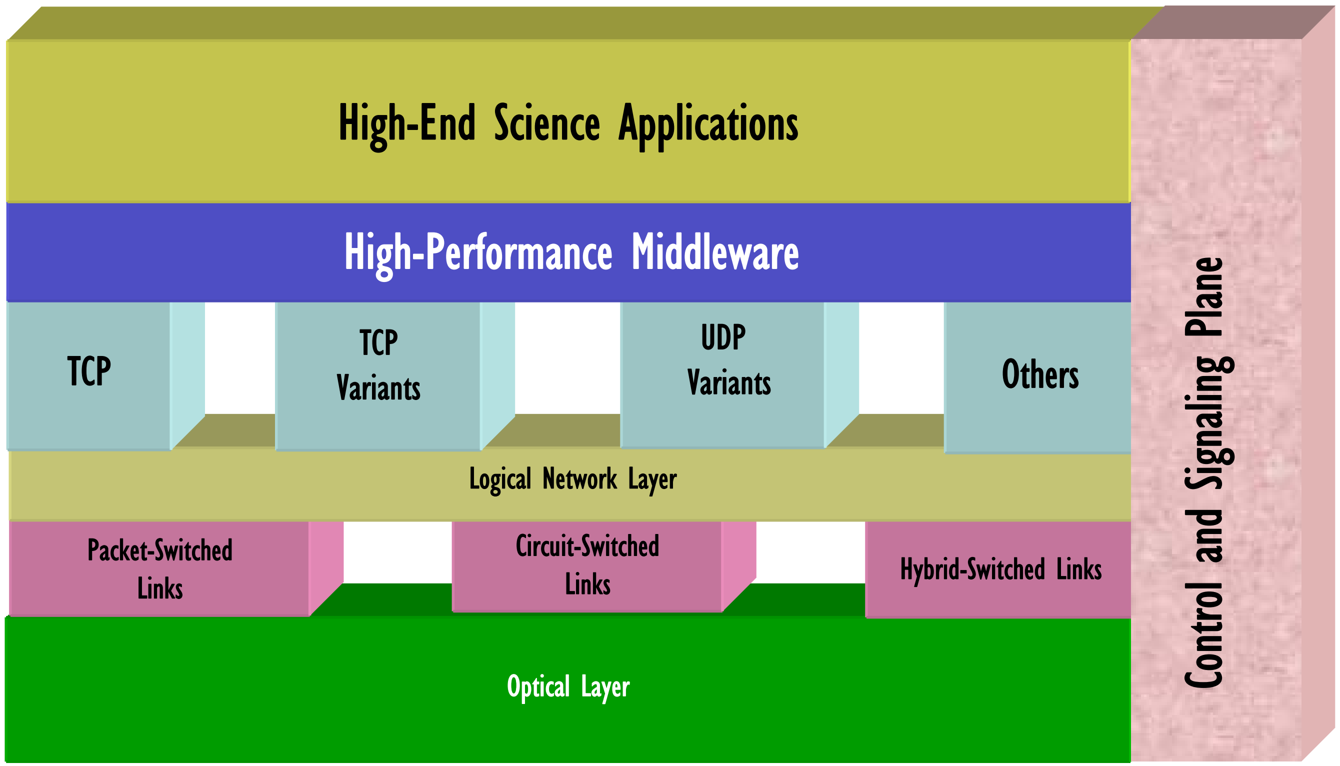

Focus Areas: Research and Development

End-to-end

- Innovative end-to-end optical network architectures
- Ultra high-speed transport protocols
- Ultra high-speed data transfer protocols
- On-demand bandwidth/scheduled bandwidth
- Control and signaling plane technologies
- High-performance middleware
- High-performance host protocol stack
- Ultra high-speed cyber security systems
- End-to-end network monitoring

U.S. Department of Energy

Office of Science

Focus Areas: Testbeds and Application Prototypes -II

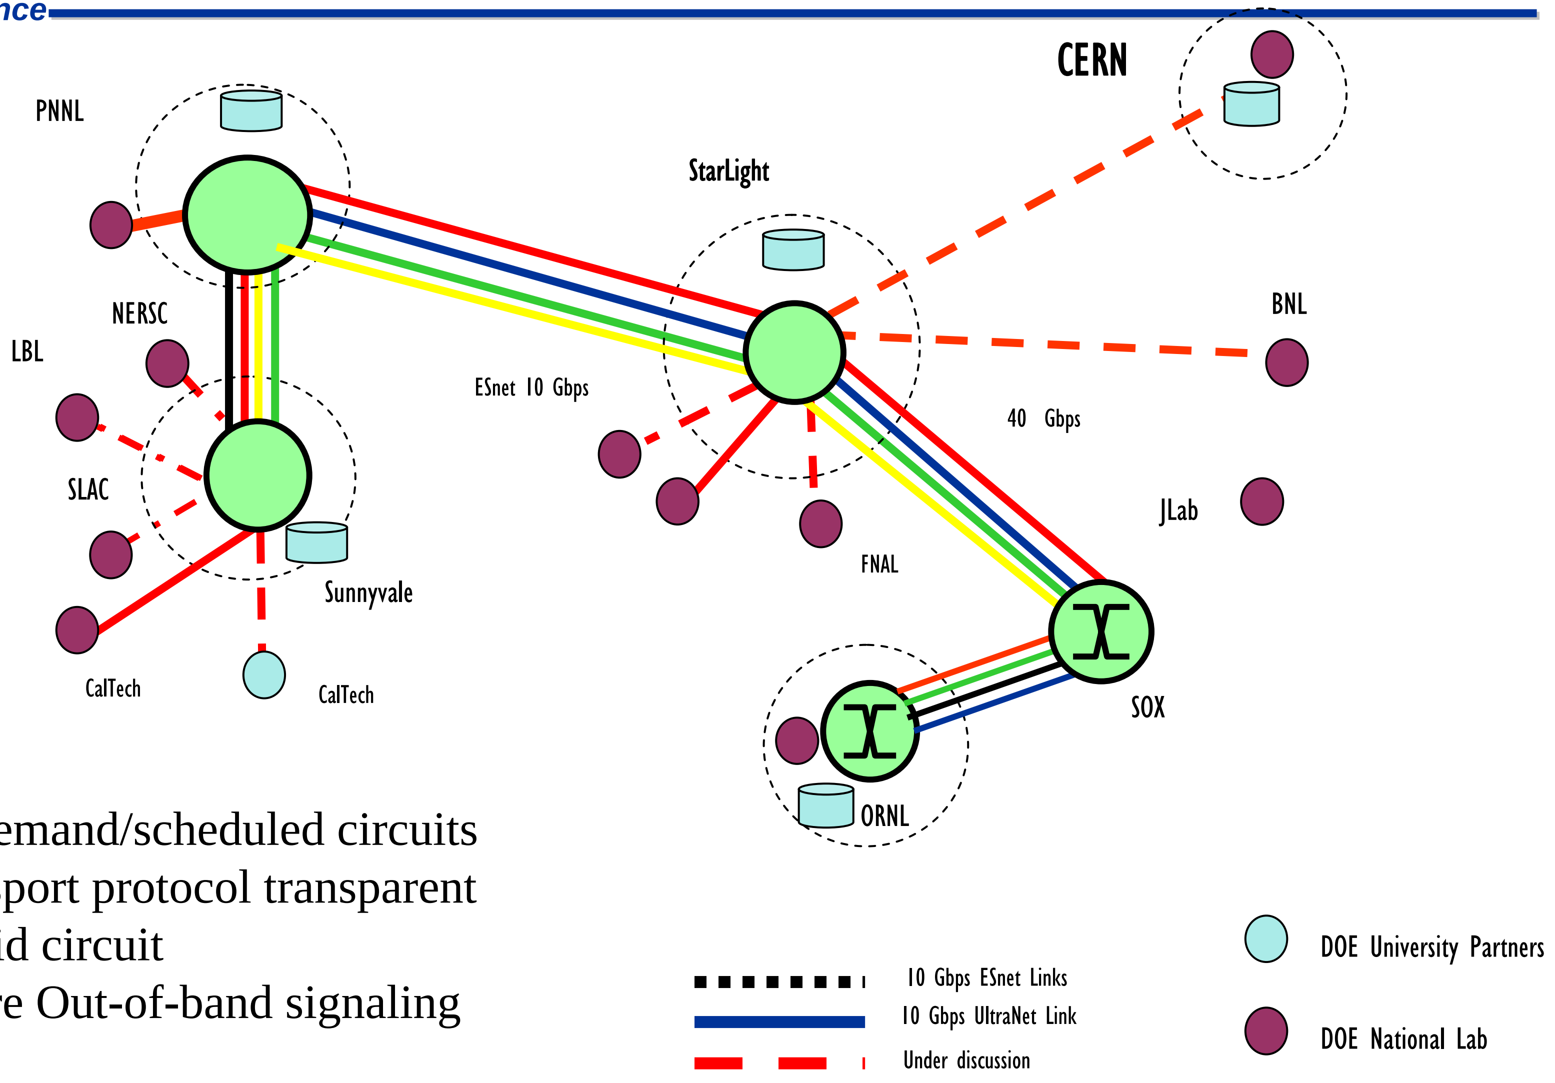

- On-demand/scheduled circuits
- Transport protocol transparent
- Hybrid circuit
- Secure Out-of-band signaling

Focus Areas: Testbeds and Application Prototypes -II

- End-to-end traffic engineering
- Traffic engineering extension and integration: QoS, MPLS, GMPLS
- Traffic engineering extension for QoS, MPLS, and GMPLS
- Control and signaling plane internetworking
- Security issues of control and signaling technologies

Call for Proposals for FY06

FY-06 SciDAC - $300M/5 years

1. Computing Science Areas
   - Applied Math
   - Computing Sciences
   - Middleware and network research
   - Application Partnerships
2. Application Science Areas
   - Fusion energy
   - High energy physics
   - Nuclear energy physics
   - Biological sciences
   - Environmental sciences

- **SBIR/STTR 2005 Solicitation**
  - ~ $6 M
  - $100k – (Phase I)
  - $800k – 900k (Phase II)

Scientific Discovery through Advanced Computing
SciDAC

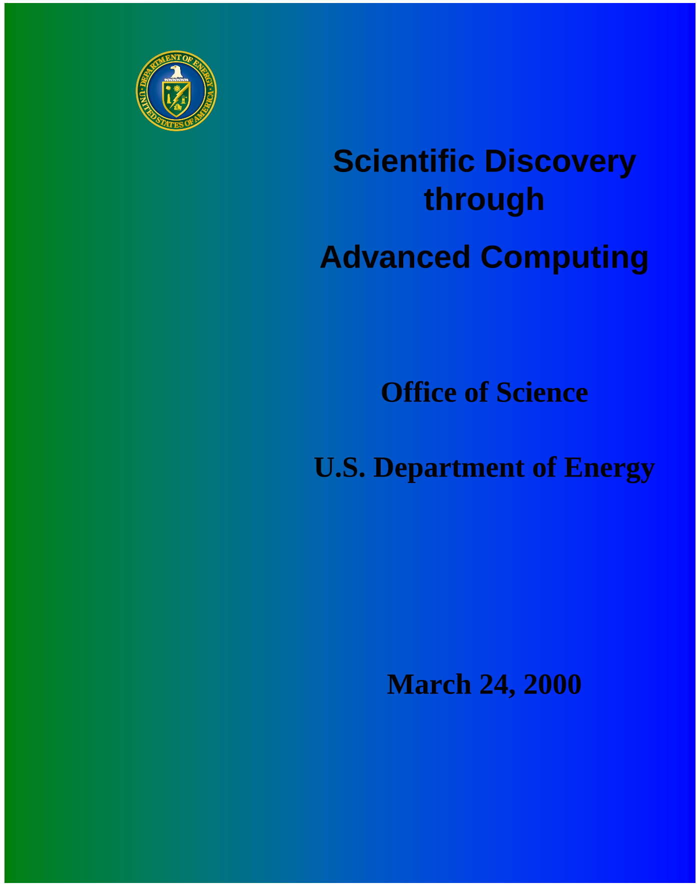
Scientific Discovery through
Advanced Computing

Office of Science

U.S. Department of Energy

March 24, 2000

An integrated program to:

(1) Create a new generation of Scientific Simulation Codes that take full advantage of the extraordinary computing capabilities of terascale computers.

(2) Create the Mathematical and Computing Systems Software to enable the Scientific Simulation Codes to effectively and efficiently use terascale computers.

(3) Create a Collaboratory Software Environment to enable geographically separated scientists to effectively work together as a team and to facilitate remote access to both facilities and data.

SciDAC – I (2001 – 2006) - $300M
SciDAC – II (2006 - 2011) - $300M

Peer-Review Process

Peer-Review Process

- Office of Science peer-review Guidelines
- Office of Science competitive awards process

Review Criteria

- Technical content
- Appropriateness of research approach
- Relevance to DOE Mission
- Budget

Reviewers and Performers Affiliation

- National labs
- University
- Industry

R&D Award Types

Short-Term Horizon Projects

- 2-3 years duration
- $150K-$300K
- Advanced deployment of new technologies
- Application prototypes and pilots

Long-Term Horizon Projects

- 3-6 years
- $300K-$1,000K
- Persistent testbeds
- Basic research projects

Future Workshop

- **Supercomputing 2005 (SC-2005), Seattle Washington**

  http://sc05.supercomp.org/home.htm

  - High-performance network BOF

- **SciDAC Annual PI meeting**

  http://www.csm.ornl.gov/workshops/DOE_SciDAC/

- **HIGH-SPEED NETWORKING WORKSHOP: THE TERABITS CHALLENGE**

  In Conjunction With IEEE INFOCOM 2006)

  Barcelona, Spain

  April 24th, 2006