Preparing for Petascale

NATIONAL CENTER
FOR COMPUTATIONAL SCIENCES

Robert Whitten Jr

May 8, 2008

OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

The Prophecy

"In the woods, as I lay on the ground and looked up into the sky, there came to me a voice as loud and as sharp as thunder. The voice told me to sleep with my head on the ground for 40 nights and I would be shown visions of what the future holds for this land.... And I tell you, Bear Creek Valley someday will be filled with great buildings and factories, and they will help toward winning the greatest war that ever will be. And there will be a city on Black Oak Ridge.... Big engines will dig big ditches, and thousands of people will be running to and fro. They will be building things, and there will be great noise and confusion and the earth will shake." **- John Hendrix, 1903**

Petascale simulations

What if . . .

We understood our own impact on the earth's climate so completely that we could head off global warming?

We understood exactly how stars exploded to provide us with the building blocks of life?

We were able to exploit an unlimited source of energy, free of pollution and greenhouse gases?

We could economically produce wire and other materials that had no electrical resistance?

New ability to model coupled ETG/ITG turbulence in shaped plasma in nonlinear phase

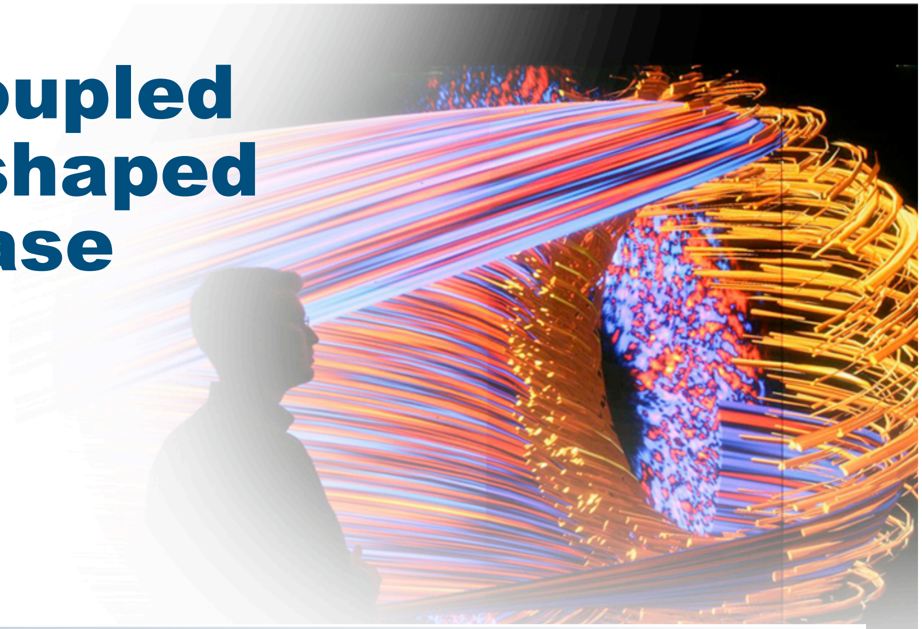

Impact: Modeling and understanding of plasma turbulence is crucial for development of stable and efficient fusion devices

Problem:

- Computational modeling of interaction of turbulence on ion and electron spatial and temporal scales
- Scales differ by orders of magnitude and have traditionally been treated by separate simulations

Results:
Simulations shed new light on how short-wavelength ETG turbulence comes into play as long-wavelength (ITG/TEM) turbulence is suppressed in pedestal (an edge transport barrier)

Ron Waltz, General Atomics

Fusion plasma:
Evolution of nonthermal plasma distributions during ion cyclotron resonance heating

Interaction of radio-frequency (RF) waves with fusion alpha particles must be understood and optimized for successful power production

Increasing current knowledge will help to plan and analyze ITER experiments

Self-consistent simulations of energetic ion evolution in plasma heated by RF waves in the ion cyclotron range of frequencies (ICRF) have been carried out on the LCF systems

Using the AORSA code coupled to the CQL3D Fokker-Planck code, we can now simulate evolution of ion distributions during ICRF heating

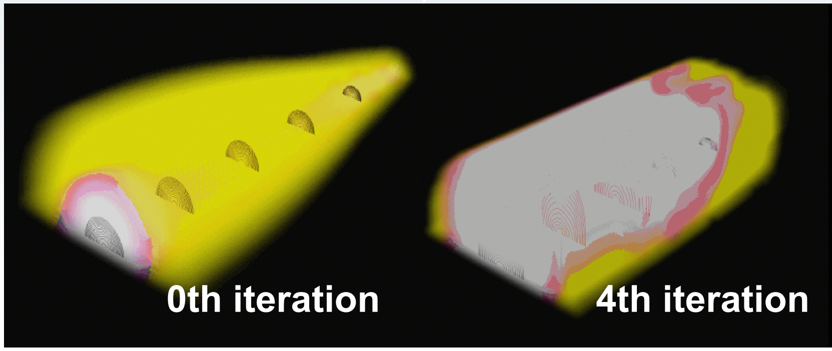

Bounce-averaged minority hydrogen distribution function in Alcator C-MOD shot 1051206002 at f = 80 MHz

Don Batchelor, Oak Ridge National Laboratory

Creation of efficient enzymes for cellulose degradation through protein engineering

- **Renewable energy**: Ethanol production from cellulose
- Developing a detailed understanding of cellulase enzyme mechanisms from multiscale modeling
  - 1- to 100-ns trajectories for systems with more than 800,000 atoms
- Carrying out simulations with different substrates and mutant enzymes

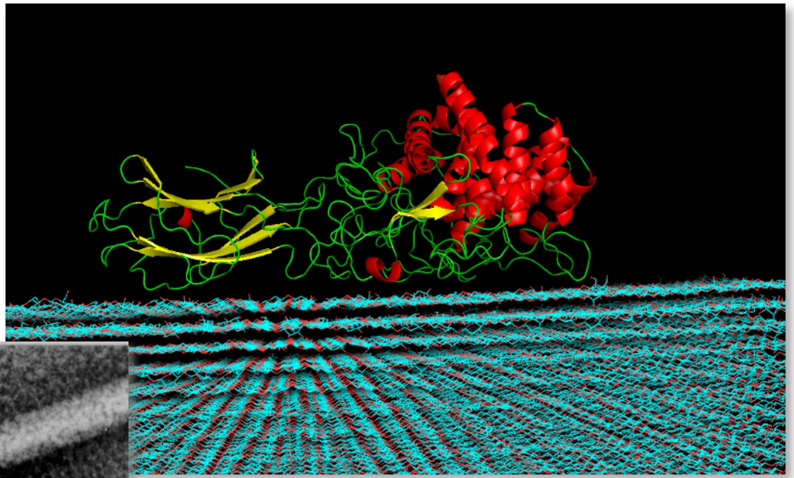

Metabolism of cellulose by exocellulase enzyme

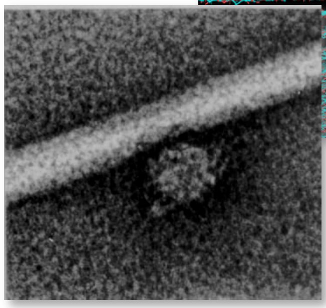

C. thermocellum growing on cellulose

Pratul Agarwal, Oak Ridge National Laboratory

Firsts in global climate modeling

Results from model advances:

- First series of equilibration runs for Carbon Land-Model Intercomparison Project (C-LAMP)
- First-ever CCSM validation with new finite-volume dynamical core (FV dycore): Critical for chemistry and full carbon cycle
- Completed first 300-year run
- Four ensemble runs with natural CO_2 forcing
- Four new ensemble runs started for anthropogenic forcing
- New results from equilibrium runs of CASA′ and CN carbon-cycle models

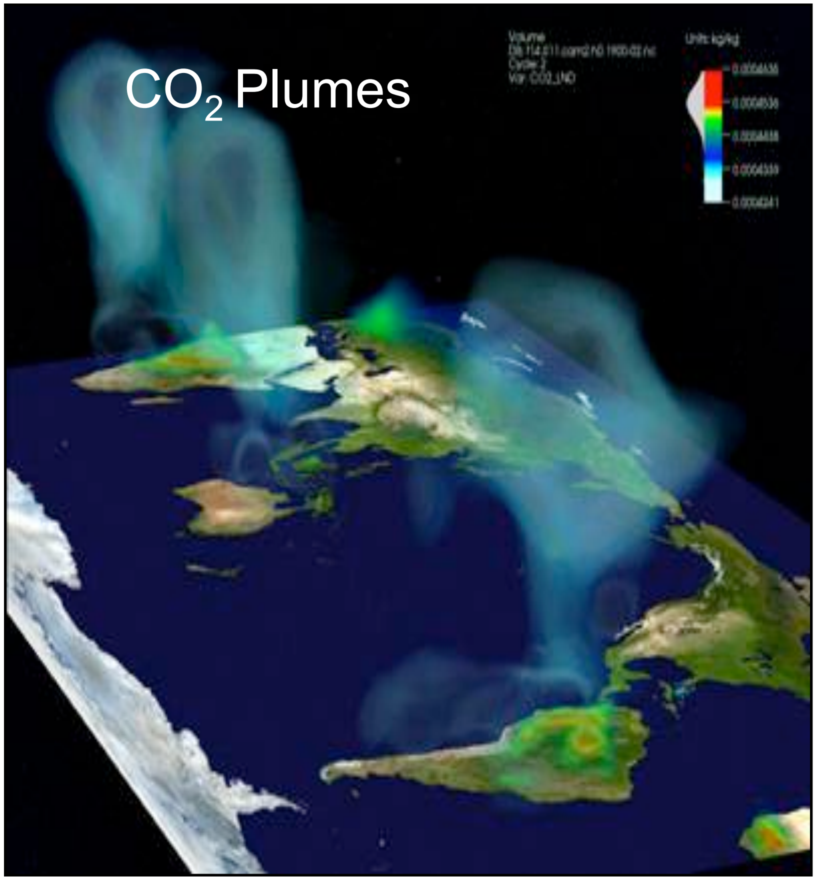

Warren Washington, National Center for Atmospheric Research

Insight into next-generation recording media:
Magnetic properties of FePt nanoparticles

FePt magnetic nanoparticles	Highly optimized calculations on LCF	Summary of results
• Material identified by industry for next-generation magnetic recording (>1 Tb psi) 	• DFT calculations using codes optimized with Cray Center of Excellence • Largest, most complex calculations of this type to date • ~50% of peak on 512 Jaguar processors (1 TF) for ~800 atoms • 2000 atoms planned	• Revealed, for first time, strong influence of nanoscale on magnetic structure of FePt nanoparticles • Sensitivity of magnetic structure to small changes demonstrates importance of calculations for materials design and optimization 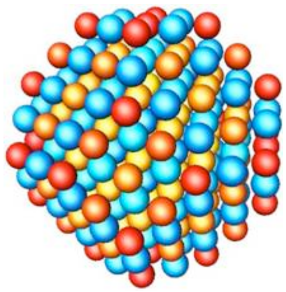

Thomas Schulthess, Oak Ridge National Laboratory

Evolution of supernovae

Supernova models must incorporate all known physics (in spite of computational demands) to capture phenomena

- Explosions were obtained for 11 and 15 solar mass progenitors
- Explosions seem to be contingent on simulating **all** of:
  - Multidimensional hydro
  - Good transport (MGFLD)
  - Nuclear burning
  - Long physical times (equivalent to long run times)

- Researchers were surprised to find that nuclear burning has an important dynamic effect
- New result builds on earlier SASI findings
  - Longer time scales required to observe explosion

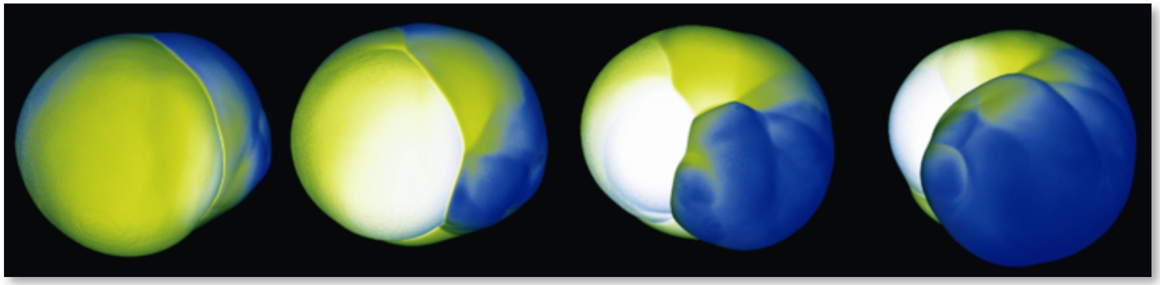

Tony Mezzacappa, Oak Ridge National Laboratory

Energy security through combustion insight

Impact: Increase thermal efficiency and decrease emissions in new land-based natural gas turbine designs

Problem:

- Understand the flamelet and thin-reaction zones where lean premixed combustion occurs
- Characterized by strong turbulence and chemistry interactions

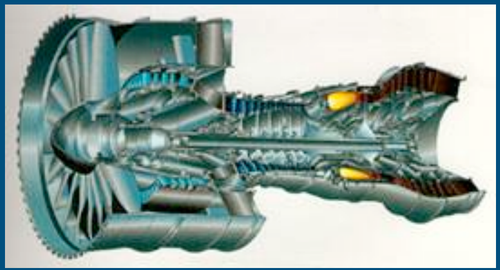

Challenges:

- Combustion at the lean flammability limit is hard
  - Prone to extinction
  - Unburned hydrocarbon emission
  - Large-amplitude pressure oscillations
  - Emission of toxic carbon monoxide

Results:

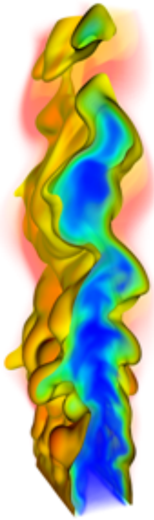

- Flame structure penetrated by small-scale eddies, but mean reaction rates still resemble a strained laminar flame
- New insights into source terms that influence flame thickness

Jackie Chen,
Sandia National Laboratories

NCCS systems, 1991–2008, paving the way to Petascale

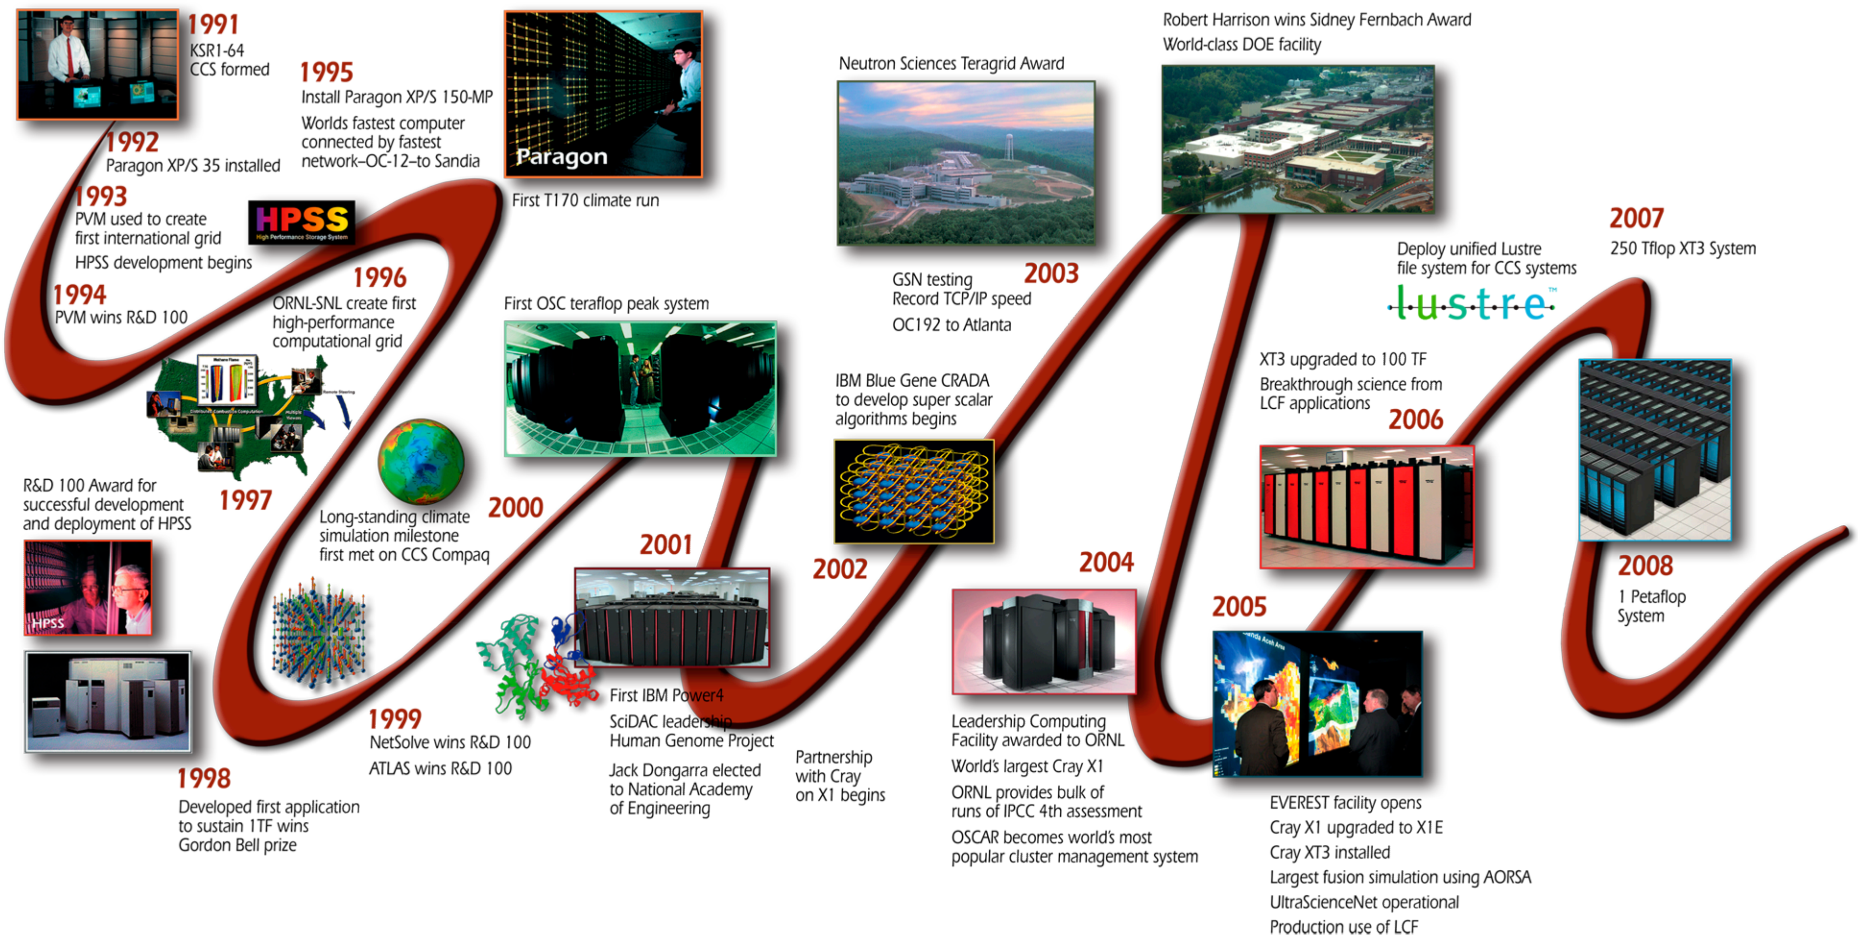

NCCS systems roadmap
Deliver 1-PF system by FY 2009
Deliver 250 TF by FY 2008

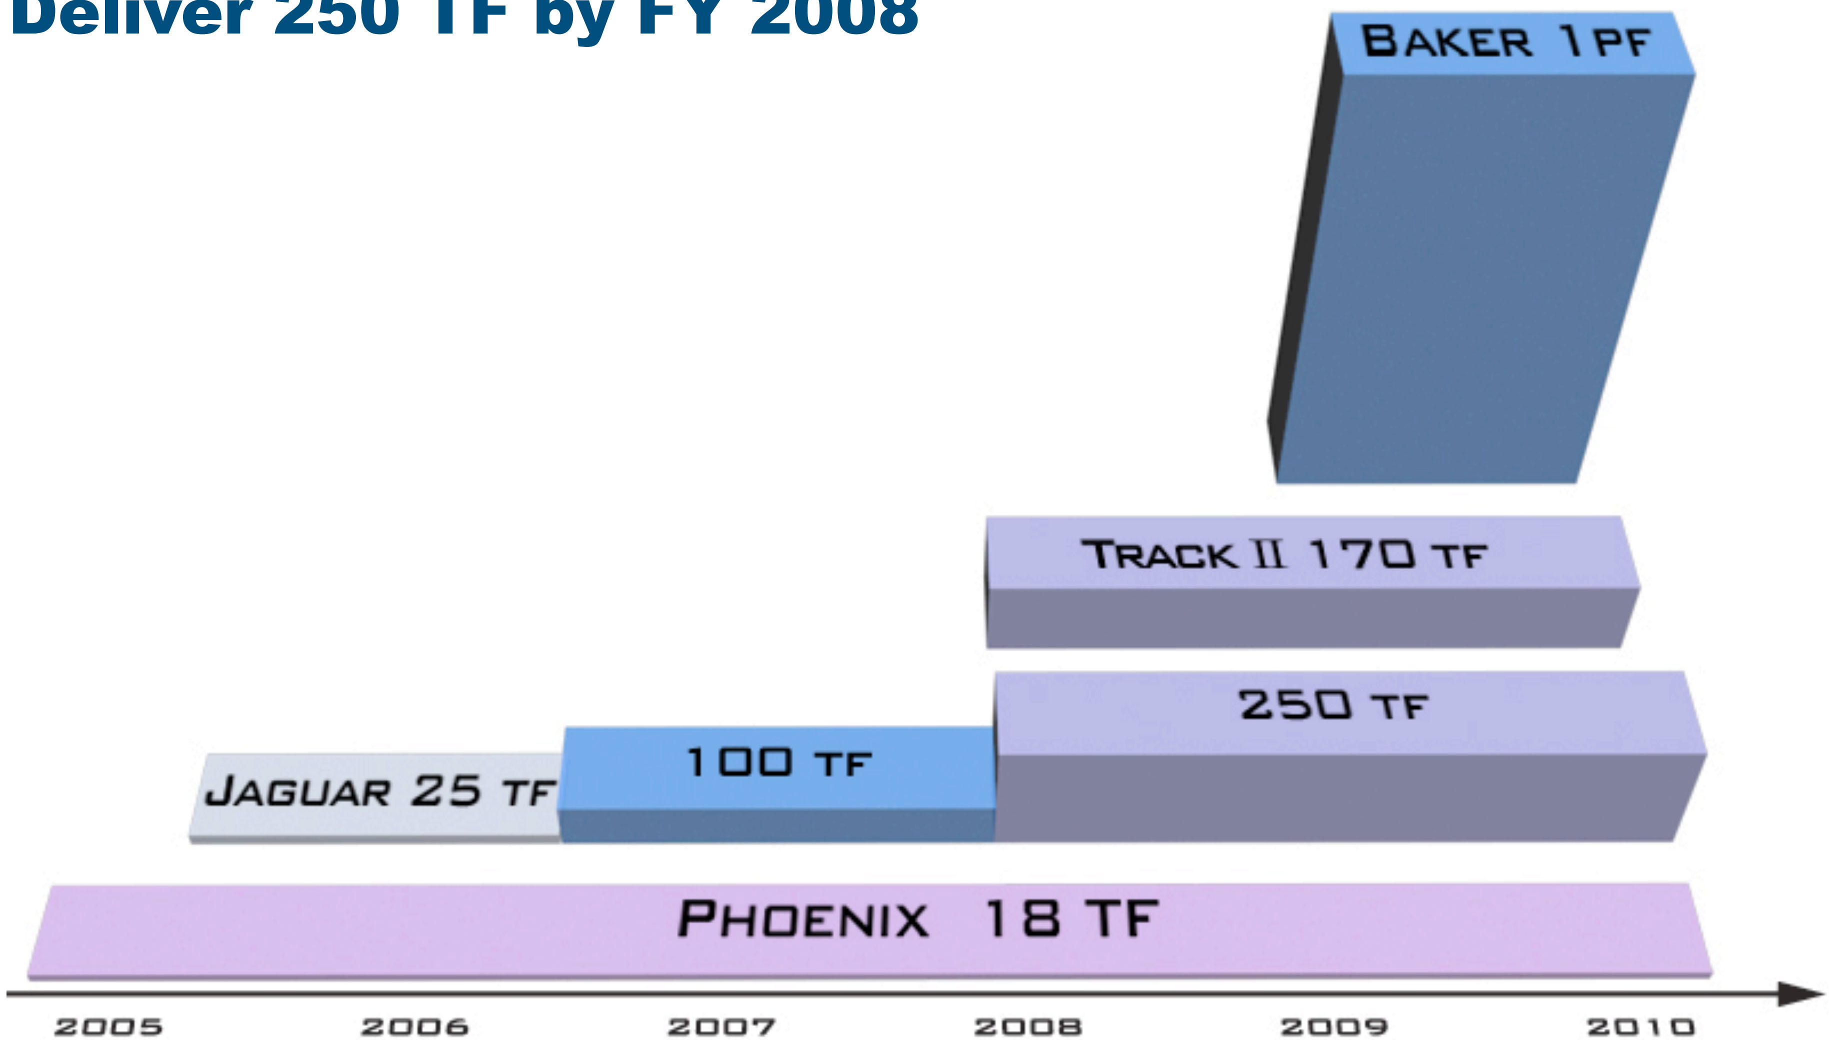

NCCS resources

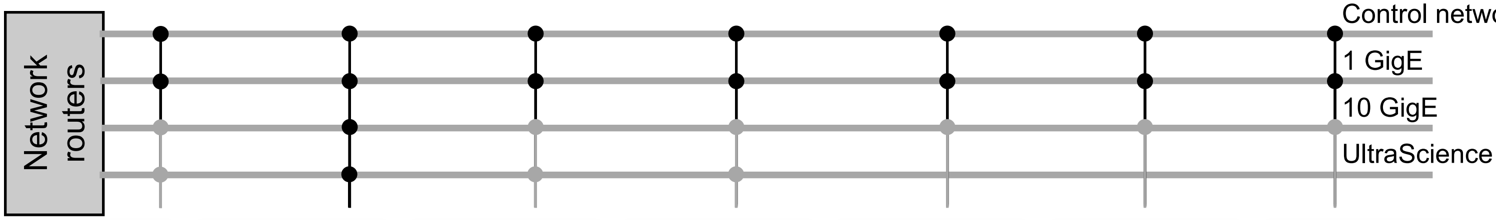

October 2007 summary

7 Systems

Supercomputers

51,110 CPUs
87 TB Memory

479 TFlops

CRAY XT4 JAGUAR	CRAY XIE PHOENIX	CRAY XT4 Kraken	IBM BLUE/GENE P	IBM LINUX NSTG	VISUALIZATION CLUSTER	IBM HPSS
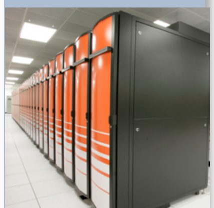		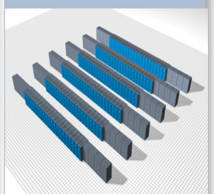			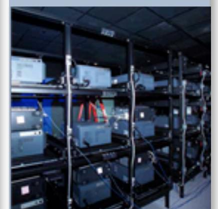	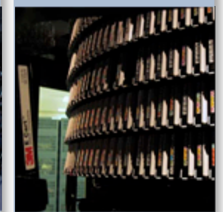
(31,238) 2.1GHz 62 TB Memory	(1,024) 0.5 GHz 2 TB Memory	(4512) 2.3 GHz 18 TB Memory	(2048) 850 MHz 4 TB Memory	(56) 3 GHz 76 GB	(128) 2.2 GHz 128 GB	Many storage devices supported
900 TB	44 TB	500 TB	60 TB	4.5 TB	9 TB	5 TB

Total local disk
1.5 PB

Evaluation platforms

- 144-processor Cray XD1 with FPGAs
- SRC Mapstation
- Clearspeed

Test systems

- 96-processor Cray XT3
- 320-Processor Development Cluster

Scientific visualization lab

35 megapixels Powerwall

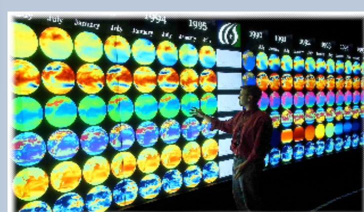

Backup storage

10 PB

10 PB
Archival storage

Jaguar - Cray XT4

- 7,832 Nodes
- 2.1Ghz quad core AMD
- 263TF
- 62TB Main Memory
- 600TB Local Disk
- Lustre filesystem
- Blueprint for Petascale

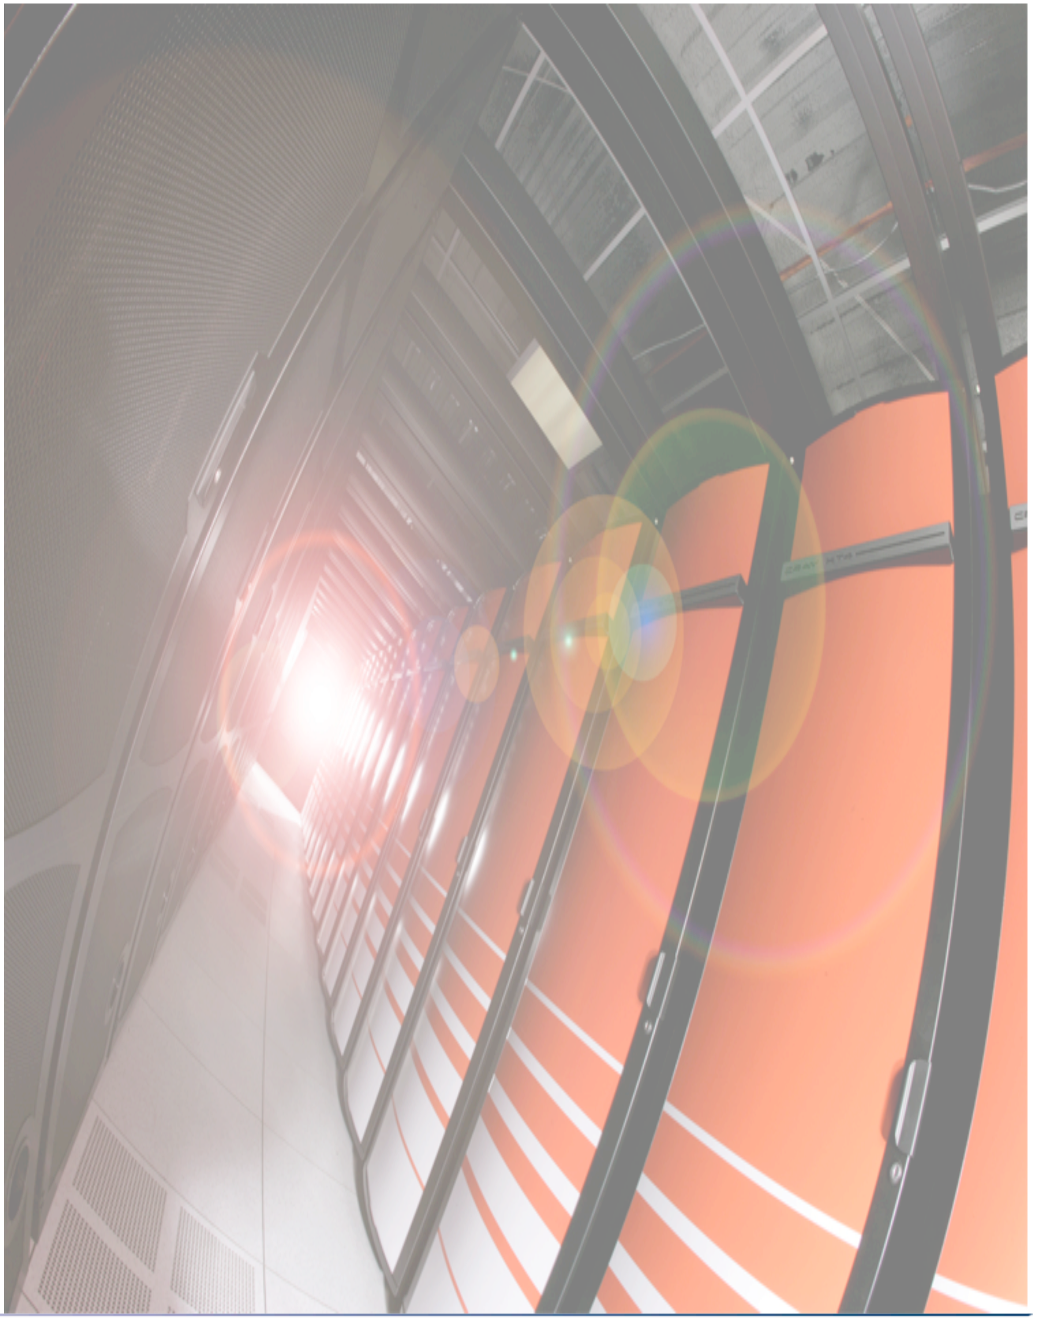

Phoenix - Cray X1E

- 1024 Multi Streaming Processors (MSP)
- 2TB Memory
- 18.5TF
- Largest X1E in world
- Liquid cooling system
- Glimpse of future?

Smoky

- 320 nodes
  - 2.2Ghz quad core AMD processors
- Dedicated development resource
- Scaling test bed

Lens - Analysis & Visualization

- **32 node Linux cluster**
  - **128 Quad-core AMD processors**
- **27 panel PowerWall**
- **35 million pixel resolution**

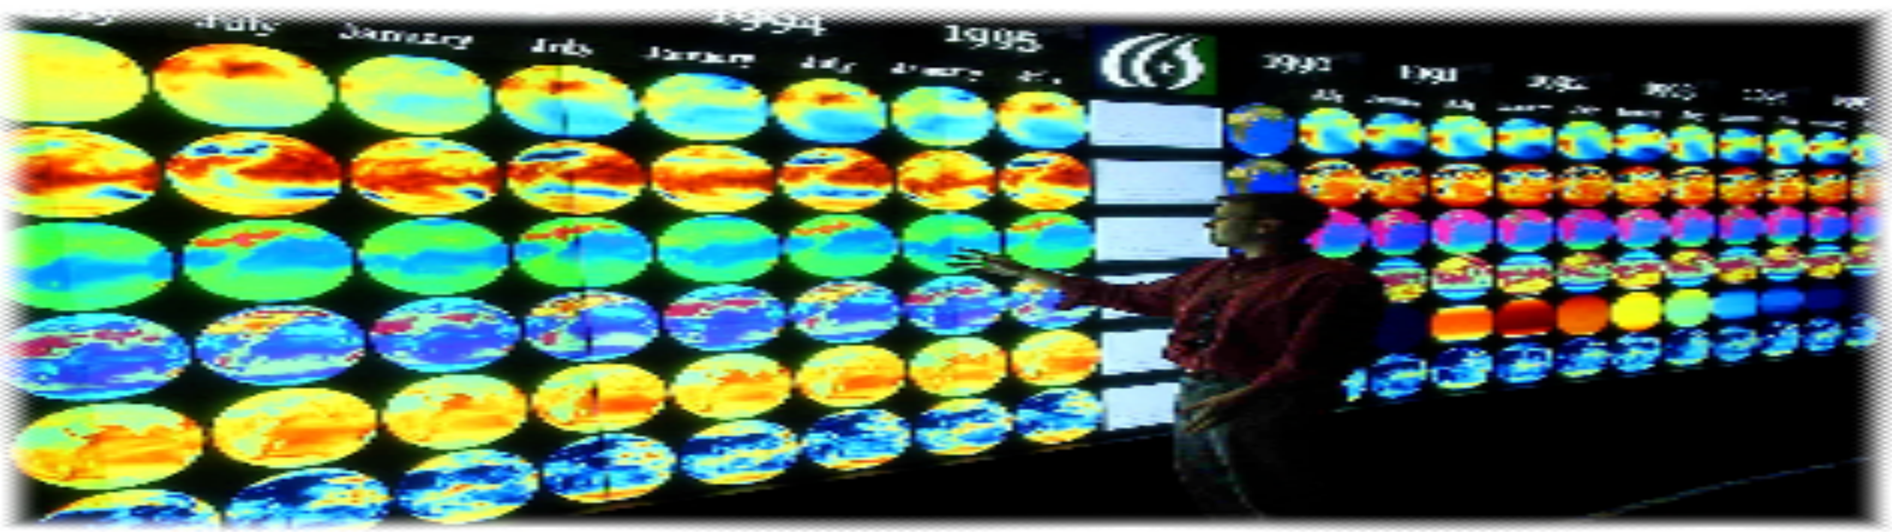

HPSS – Mass Storage

- Tape library and disk storage
- 2PB stored data
- Expandable to 10PB
- Recently expanded in preparation for petascale computing

National Institute for Computational Science (NICS)

- NSF Track II award
- UT/ORNL Partnership
- 170TF system in 2008
- 1PF system in 2009
- Housed at NCCS
- Managed by UT

Kraken - Cray XT4

- NICS Resource
- 170TF
- 2.2Ghz quad core AMD
- Online Summer '08

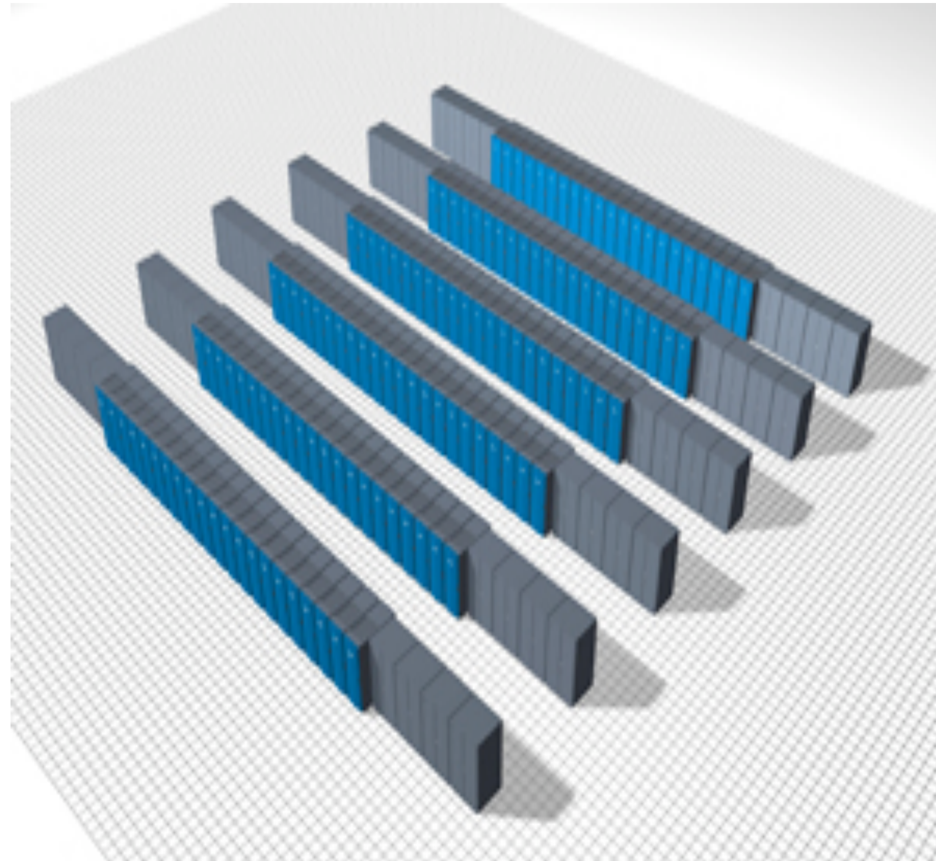

Operations Council
M. Dobbs, Facility Mgmt.
P. Lindsey, Facility Mgmt.
D. Ridz, Recruiting
D. Vasil, Cyber Security
W. McCrosky, Finance Officer
R. Toedte, Safety & Health
H. Abbott, Org. Mgmt. Specialist
N. Wright, HR Manager

National Center for Computational Sciences
J. Hack, Director
A. Bland, LCF Project Director
L. Gregg, Division Secretary
A. Brannon, ORHS Student Intern[1]

Advisory Committee
J. Dongarra
T. Dunning
K. Droegemeier
S. Karin
D. Reed
S. Tichenor

Director of Science
D. Kothe

Director of Operations
J. Rogers

LCF System Architect
S. Poole

Chief Technology Officer
A. Geist

Deputy Project Director
K. Boudwin

Cray Supercomputing Center for Excellence
J. Levesque
S. Allen

L. DeRose[4]
D.Kiefer[4]
J. Larkin[4]
N. Wichmann[4]

[1]Student
[2]Post Graduate
[3]JICS
[4]Cray, Inc.
[5]Matrixed
[6]Subcontract
[7]Interim

Scientific Computing
R. Kendall
L. Rael

S. Ahern[##]
S. Alam[5]
E. Apra[5]
R. Hartman-Baker
D. Banks[3]
R. Barreto[2]
R. Barrett[5]
J. Daniel
M. Eisenbach
M. Fahey
J. Gergel[5]
S. Hodson
C. Jin
S. Klasky[#]
J. Kuehn[5]
V. Lynch[5]
B. Messer
R. Mills[5]
G.Ostrouchov[5]
N. Podhorszki
D. Pugmire
R. Sankaran
S. Thornton[5]
A. Tharrington
R. Toedte
T. White

[#]End-to-End Solutions Lead
[##]Viz Task Lead

User Assistance And Outreach
J. White
L. Rael

T. Anderson[6]
D. Beach[6]
D. Frederick
C. Fuson
B. Gajus[6,5]
C. Halloy[3]
S. Hempfling
M. Henley
J. Hines
S. Jones[6]
D. Levy[5]
B. Renaud
B. Whitten
L. Williams[5]
K. Wong[3]

High-Performance Computing Operations
A. Baker
S. Allen

M. Bast
J. Becklehimer[4]
J. Breazeale[6]
J. Brown[6]
M. Disney
A. Enger[4]
C. England
J. Evanko[4]
A. Funk[4]
D. Giles[5]
M. Griffith
V. Hazelwood[5]
J. Hill
C. Leach[6]
D. Londo[4]
J. Lothian
D. Maxwell[#]
M. McNamara[4]
J. Miller[5]
G. Phipps, Jr.[6]
S. Shpanskiy
D. Vasil
S. White
C. Willis[4]
T. Wilson[6]

[#]Technical Coordinator

Technology Integration
S. Canon
S. Allen

T. Barron
D. Dillow
R. Graham
K. Matney
S. Oral
D. Steinert
G. Shipman
F. Wang
V. White
W. Yu[5]

Project
Site Preparation
K. Dempsey

Hardware Acquisition
J. Rogers

Test & Acceptance Development
S. Canon

Commissioning
A. Baker

Project Management
D. Hudson[5]

Project R & D
A. Geist

Cray Proj. Director
T. Watts[4]

NCCS facility upgrades

- **3.3 MW power upgrade completed in May 2006**
  - Provides power through the 250-TF system upgrade
- **7 MW power upgrade completed in 2007**
- **15 MW power upgrade in 2008**
- **Upgrade of second-floor for 100-TF system**
  - Increased raised floor height to 36 inches
  - Installed 25 air handlers
  - Replaced 4-inch pipe with 12-inch pipe
  - Installed 11 power distribution units

Utility upgrades for 1-PF system

- Upgrade building power from 7.3 to ≥25 MW
- Tie building cooling into ORNL cooling plant via 36-inch pipes
- Upgrade UPS and generator capability

Scaling other services to meet user needs

	100 TF	250 TF	1000 TF
Local-area network	60 GB/s	140 GB/s	240 GB/s
Wide-area network	3 GB/s	4 GB/s	5 GB/s
Archival storage	4 PB, 4 GB/s	10 PB, 10 GB/s	18 PB, 19 GB/s
Center-wide file system: Spider	250 TB, 10 GB/s	1 PB, 60 GB/s	10 PB, 240 GB/s
Data analysis and visualization	Cluster based	Lens Cluster	Use 250-TF Jaguar

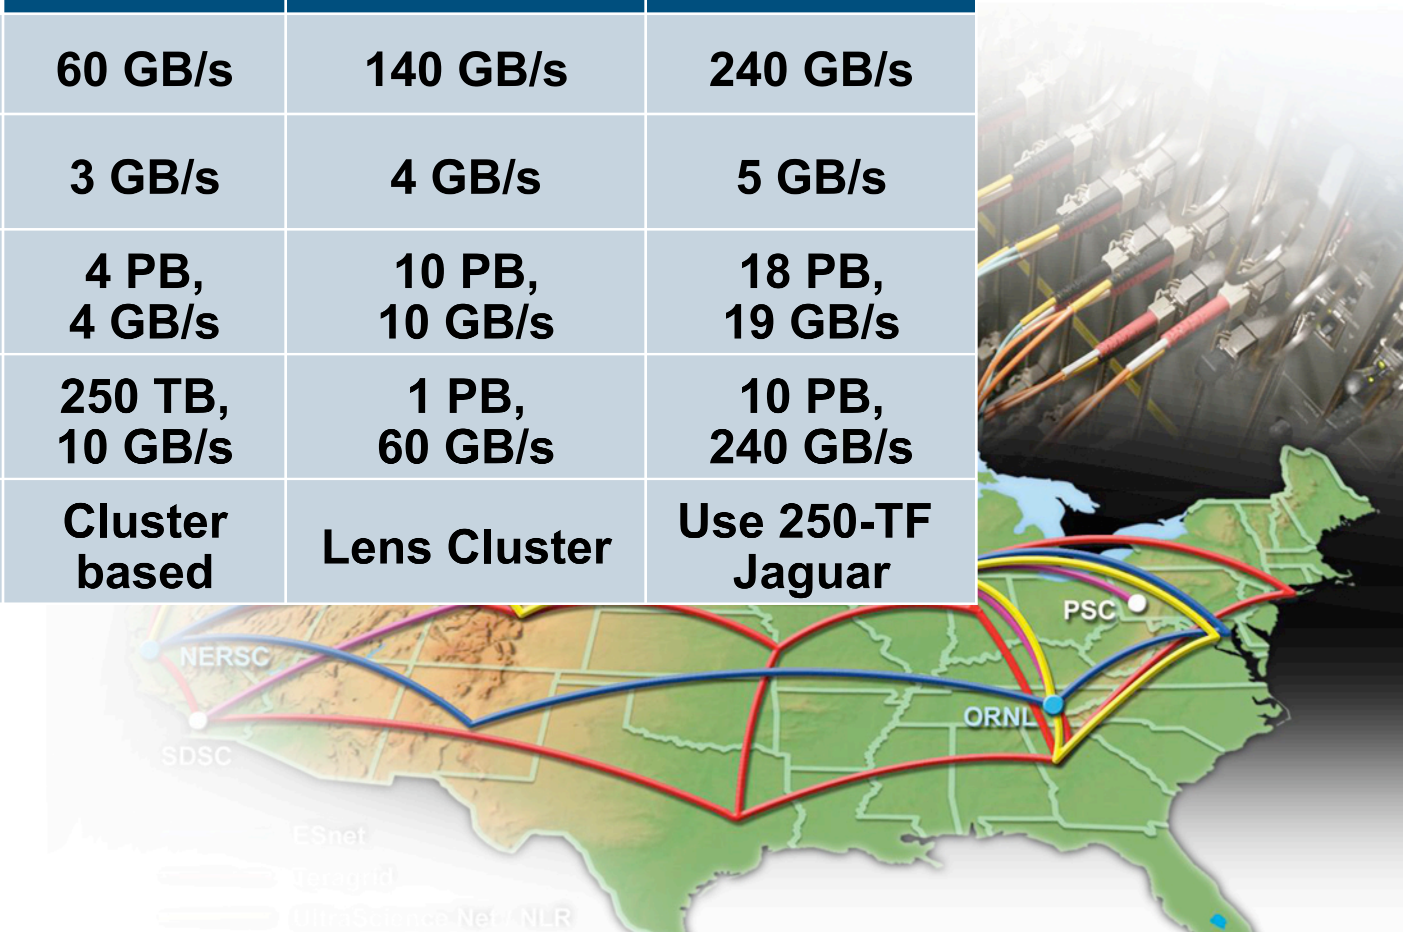

Software and science: Overview

Building the petascale system lays the foundation for science and education at the petascale:

Scientists using a production petascale system with petascale application development software

- Establish major new fully integrated petascale computing environment
- Develop petascale software infrastructure to enable productive utilization and system management
- Empower scientific and engineering progress and allied educational activities using petascale system
- Develop and educate next-generation computational scientists for use of petascale system

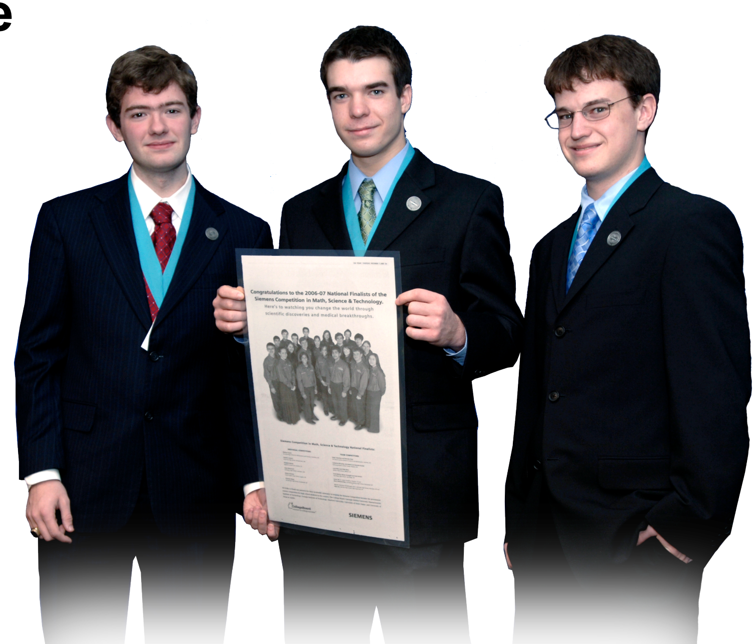

Steps for ensuring application readiness

- Early identification of codes for acceptance test and early access
- Scientific Computing Group: Liaisons to application teams
- Cray Center of Excellence: Scaling of libraries and codes
- Lustre Center of Excellence: Help tune application parallel I/O
- Councils: Identification of application requirements

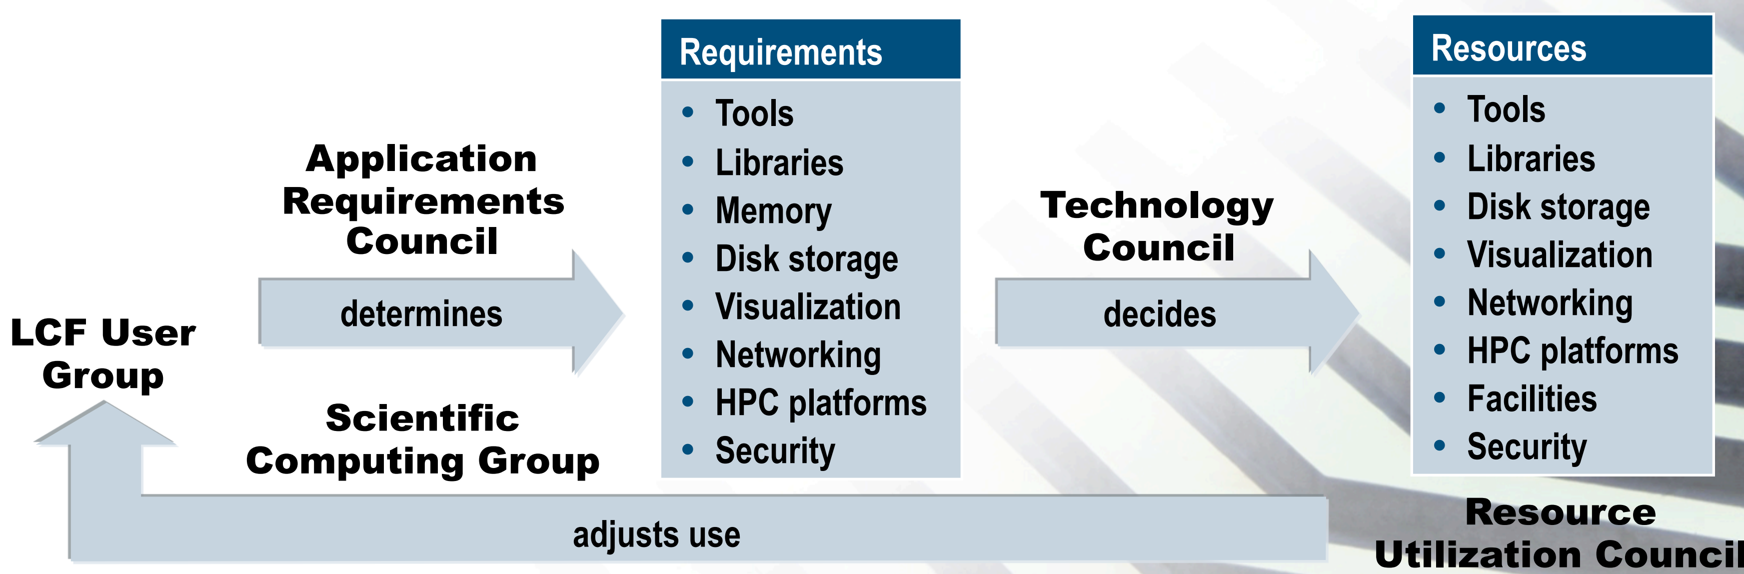

Consulting at Petascale

- Consulting at Petascale means education
  - For consultants
  - For researchers
- Where do you start?
  - PetaApps
    - Learn how they do it
      - Models
      - Methods
      - Tools / frameworks
- Best Practices
  - Learn and share

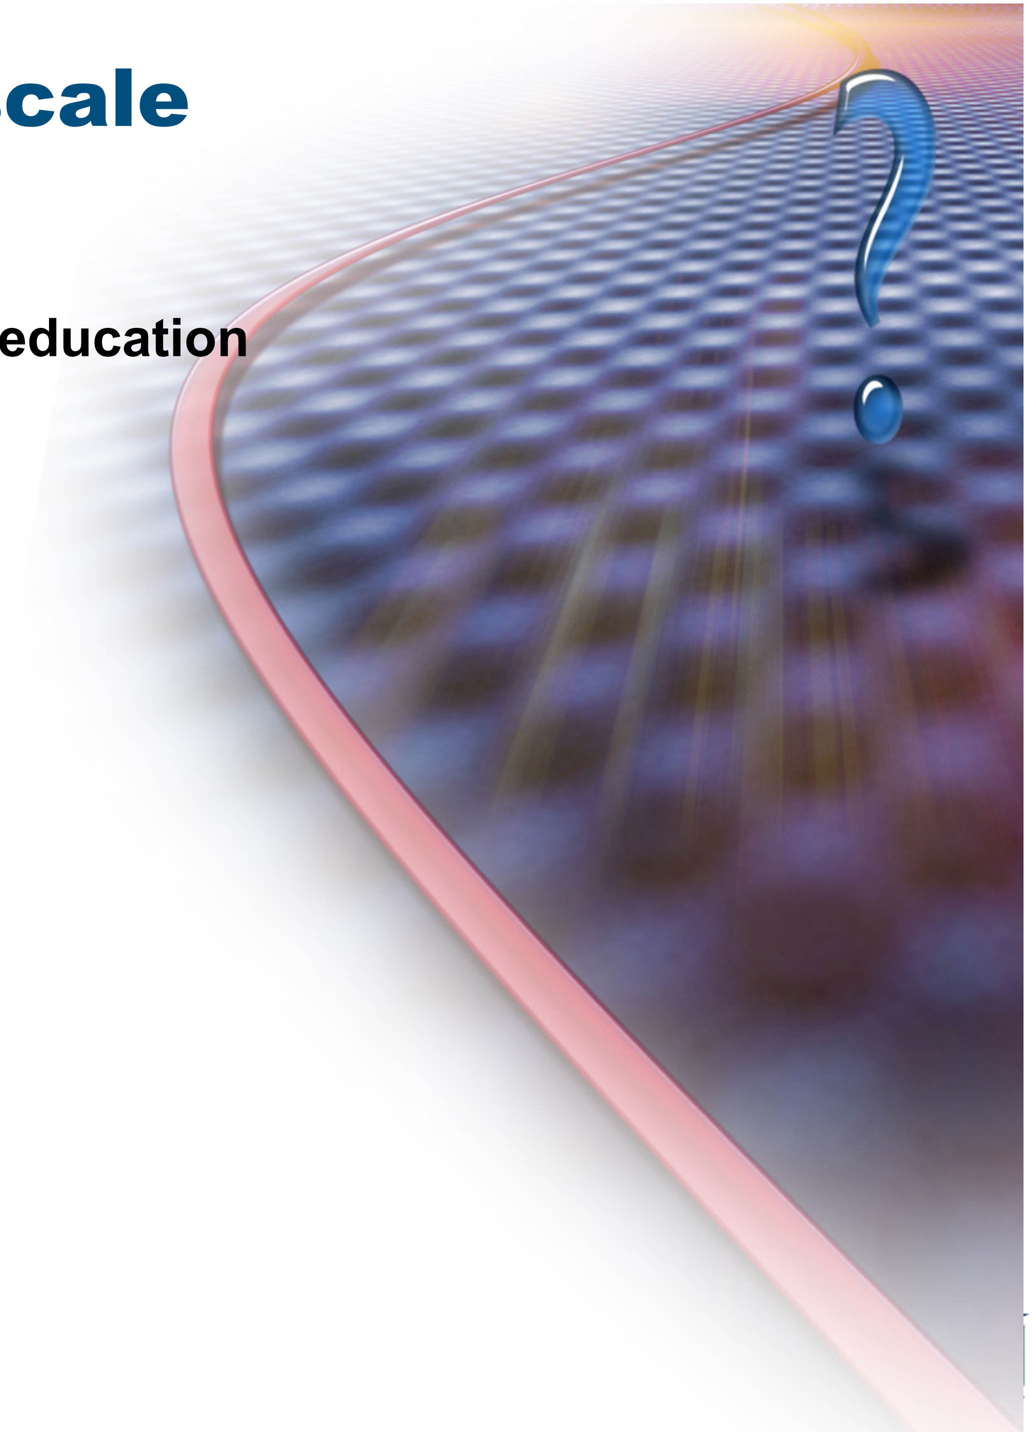

Questions?

http://www.nccs.gov