1 Degree POP on Various Platforms

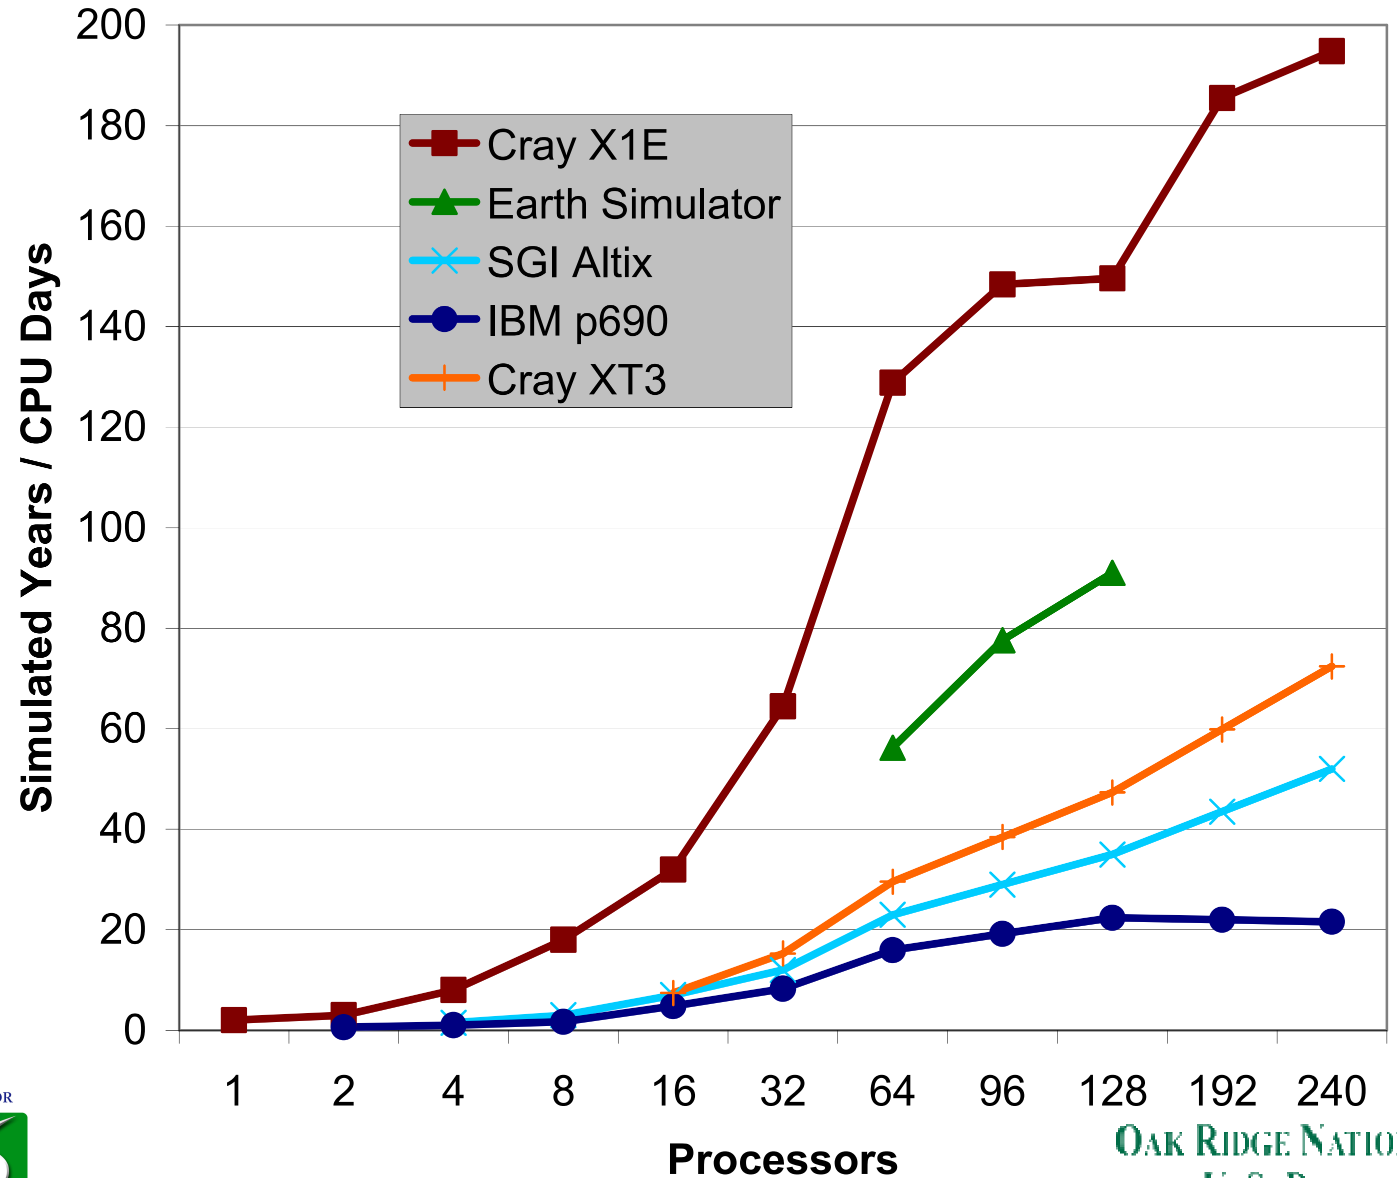

National Center for Computational Sciences

Oak Ridge National Laboratory
U. S. Department of Energy

1/10th Degree POP Performance

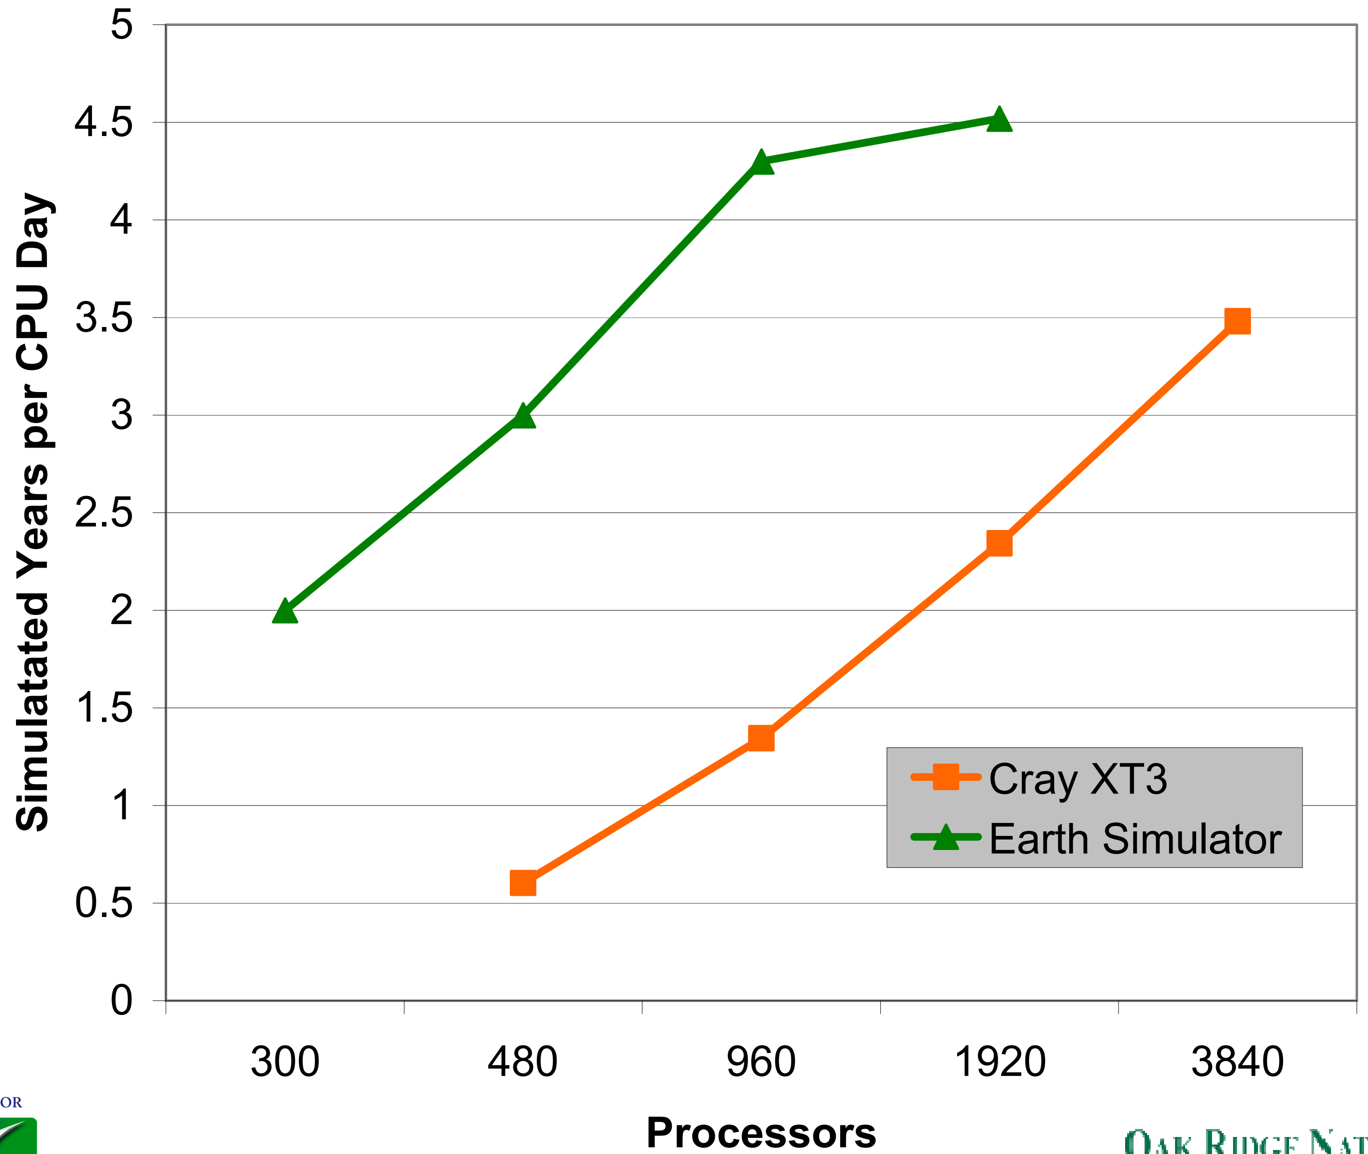

National Center for Computational Sciences

Oak Ridge National Laboratory
U. S. Department of Energy

Amber for computational biology –
Approaching one nanosecond / day throughput

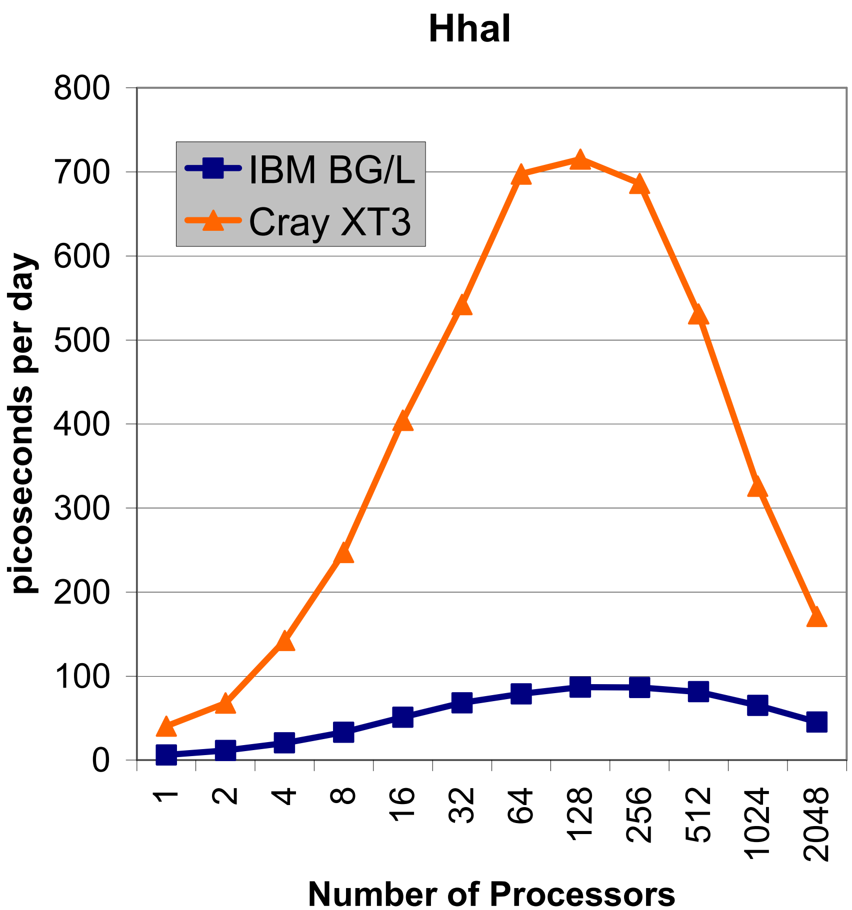

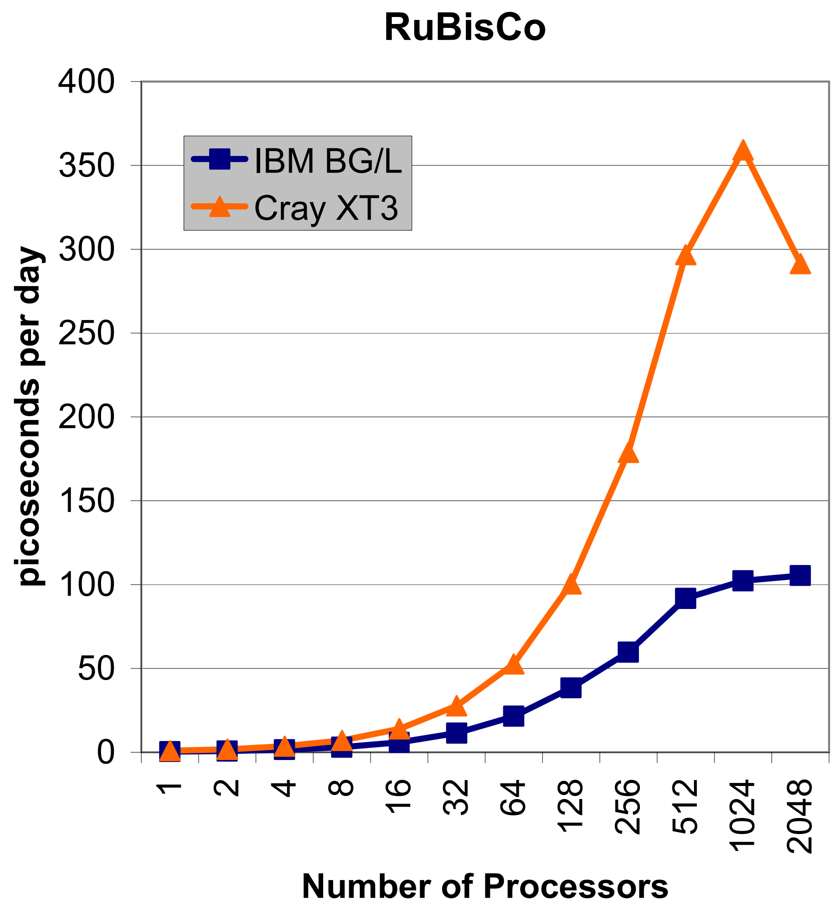

Early results show very positive XT3 results

Cray XT3 on the HPC Challenge Benchmarks

• G-HPL	1st place	20.527 TF/s
• G-PTRANS	1st place	874.899 GB/s
• G-Random Access	1st place	0.533072 Gup/s
• G-FFTE	1st place	905.569 GF/s
• G-Stream Triad	1st place	26,020.8 GB/s

Reported results as of Sept. 20, 2005

Outline

- Who are we: NLCF at ORNL
- Science Motivators (first results)
- Allocations and Support for Break-Through Science
- 2006-2010 Roadmap
- Wrapup

Facility plus hardware, software, science teams contribute to Science breakthroughs

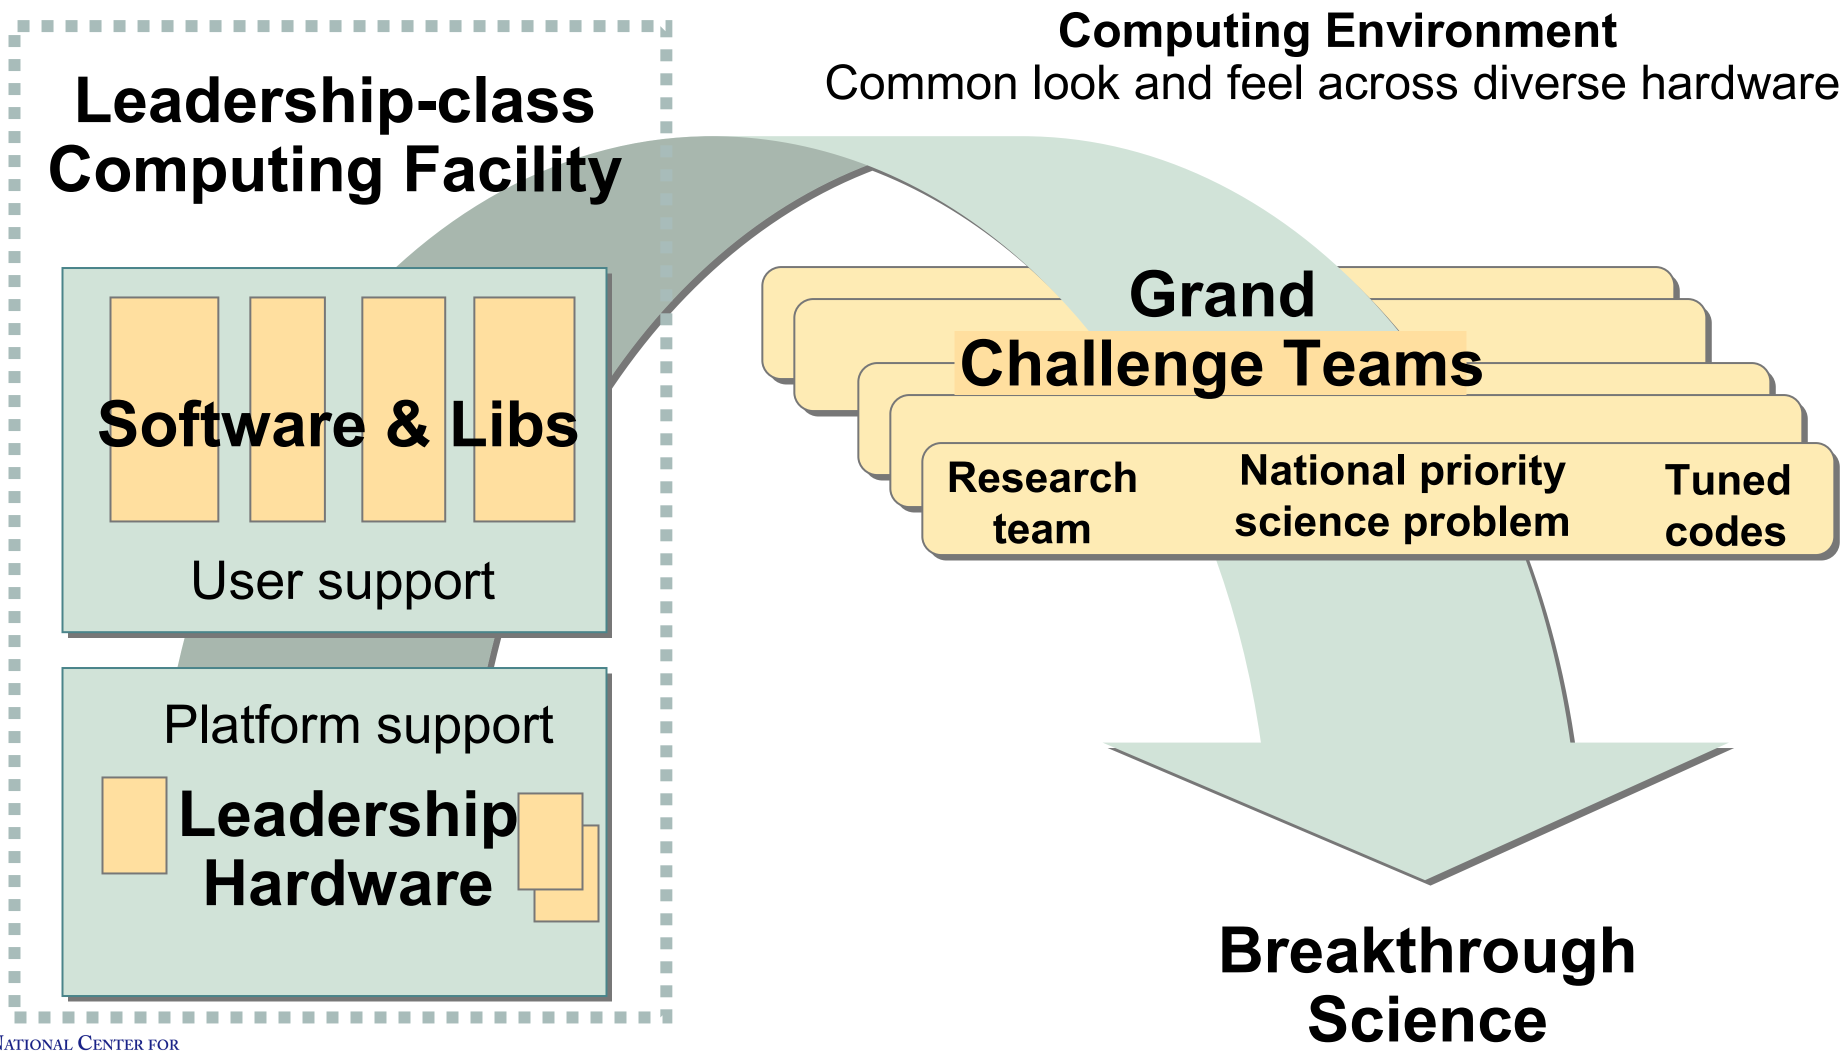

National Leadership Computing Facility
2006 Call for Proposals

- Projects must be funded by the DOE Office of Science or support the mission of the Office of Science
- Principal investigator, multi-principal investigator teams, multi-institutional teams and end station proposal teams may apply for LCF resources
- Multi-year proposals will be considered subject to annual review.
- Expectation that Leadership systems will enable U.S. to be "first to market" with important scientific and technological capabilities, ideas, and software
- Limited set of scientific applications selected and given substantial access to leadership system

Office of Science
U.S. DEPARTMENT OF ENERGY

National Center for Computational Sciences

Access to NLCF

Call for Proposals

- LCF and INCITE calls yearly
- Pilot Project calls biannually

National Center for Computational Sciences

Review

- Technical readiness
- Scalability

Allocations

- Grand Challenges
- End Stations
- Pilot Projects

Project types

Grand Challenge

- Scientific problems that may only be addressed through access to LCF hardware, software, and science expertise
- Multi-year, multi-million CPU hour allocation

End Station

- Computationally intense research projects, also dedicated to development of community applications
- Multi-year, multi-million CPU hour allocation

Pilot Project

- Small allocations for projects, in preparation for future Grand Challenge or End Station submittals
- Limited in duration

Response to 2006 Call for Proposals

Life Sciences	2
Nanoscience	5
Materials	6
Computer	3
Chemical	5
Environmental	2
Engineering Physics	2
Computational Mechanics	1
Combustion	3
Climate and Carbon Research	6
Fusion	6
Astrophysics	6
Accelerator	2
Nuclear Physics	1
Turbulence	1
High Energy Physics	1

The majority of the proposals are led by DOE and University PIs.

Outline

- Who are we: NLCF at ORNL
- Science Motivators (first results)
- Allocations and Support for Break-Through Science
- 2006-2010 Roadmap
- Wrapup

Delivering Science and Leadership in High-end Computing

Nanoscience

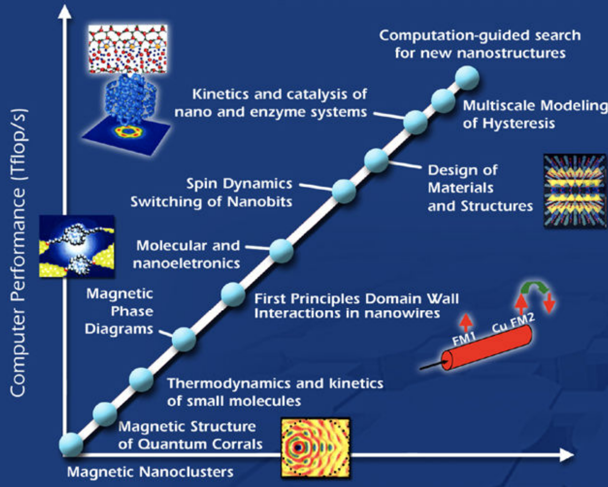

Expected Outcomes

5 years

- Realistic simulation of self-assembly and single-molecule electron transport
- Finite temperature properties of nanoparticles/quantum corrals

10 years

- Multi-scale modeling of molecular electronic devices
- Computation-guided search for new materials/nanostructures

Climate

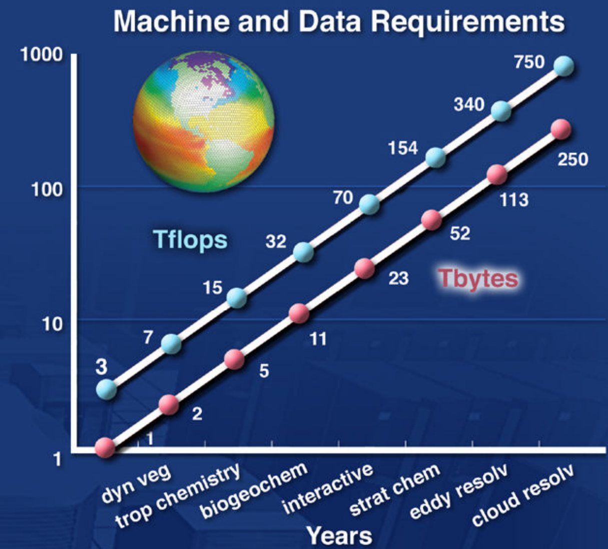

Expected Outcomes

5 years

- Fully coupled carbon-climate simulation
- Fully coupled sulfur-atmospheric chemistry simulation

10 years

- Cloud-resolving 30-km spatial resolution atmosphere climate simulation
- Fully coupled, physics, chemistry, biology Earth system model

Biology

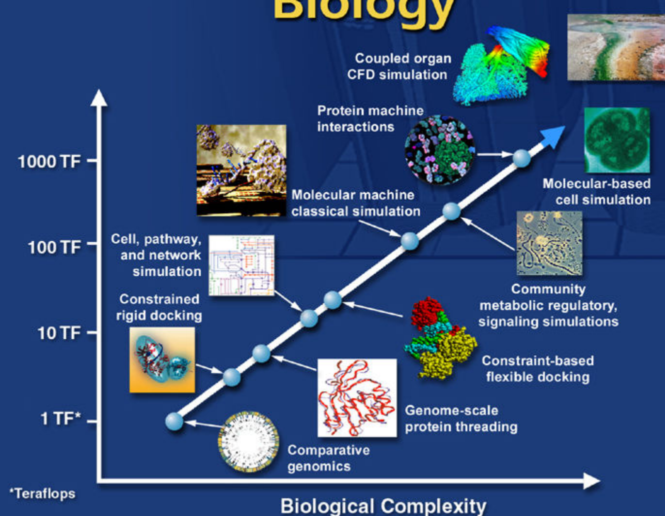

High-performance computing is essential to the high-throughput experimental approach.

Expected Outcomes

5 years

- Metabolic flux modeling for Hydrogen and Carbon fixation pathways
- Constrained flexible docking simulations of interacting proteins

10 years

- Multi-scale stochastic simulations of combined microbial metabolic, regulatory and protein interaction networks
- Dynamics simulations of complex molecular machines

Fusion

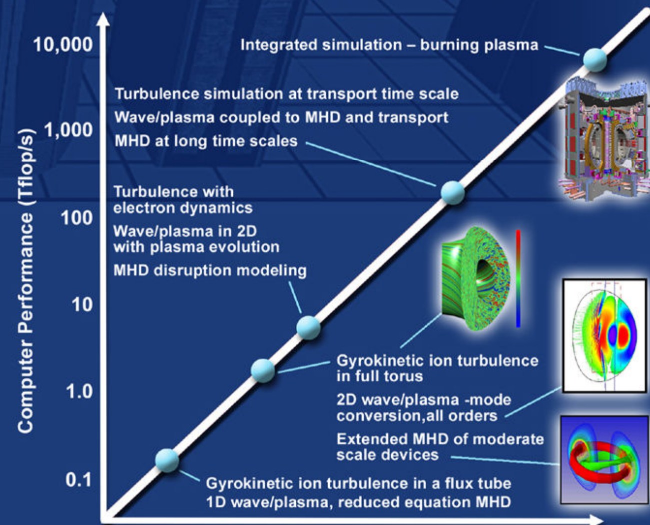

Expected Outcomes

5 years

- Full-torus, electromagnetic simulation of turbulent transport with kinetic electrons for simulation times approaching transport time-scale
- Develop understanding of internal reconnection events in extended MHD, with assessment of RF heating and current drive techniques for mitigation

10 years

- Develop quantitative, predictive understanding of disruption events in large tokamaks
- Begin integrated simulation of burning plasma devices - multi-physics predictions for ITER

Center for Computational Sciences
Oak Ridge National Laboratory

Office of Science
U.S. Department of Energy

NLCF plan for the next 5 years:

Vector Arch
Global memory
Powerful CPU

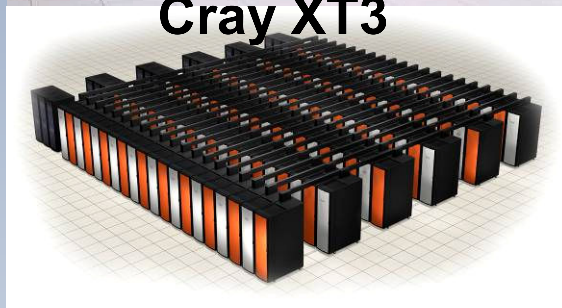

Cluster Arch
Low latency
High bandwidth

Scalability
100K CPU
MB/CPU

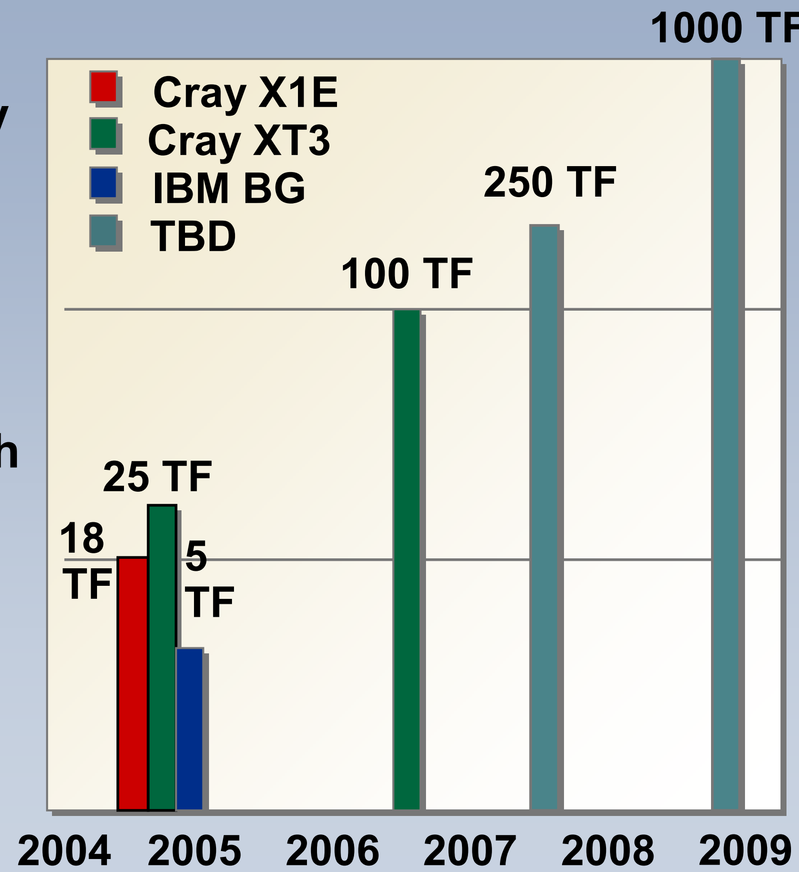

ORNL Path to Leadership Computing

Scalable, High-Bandwidth Computing

Performance
Programmability
Robustness

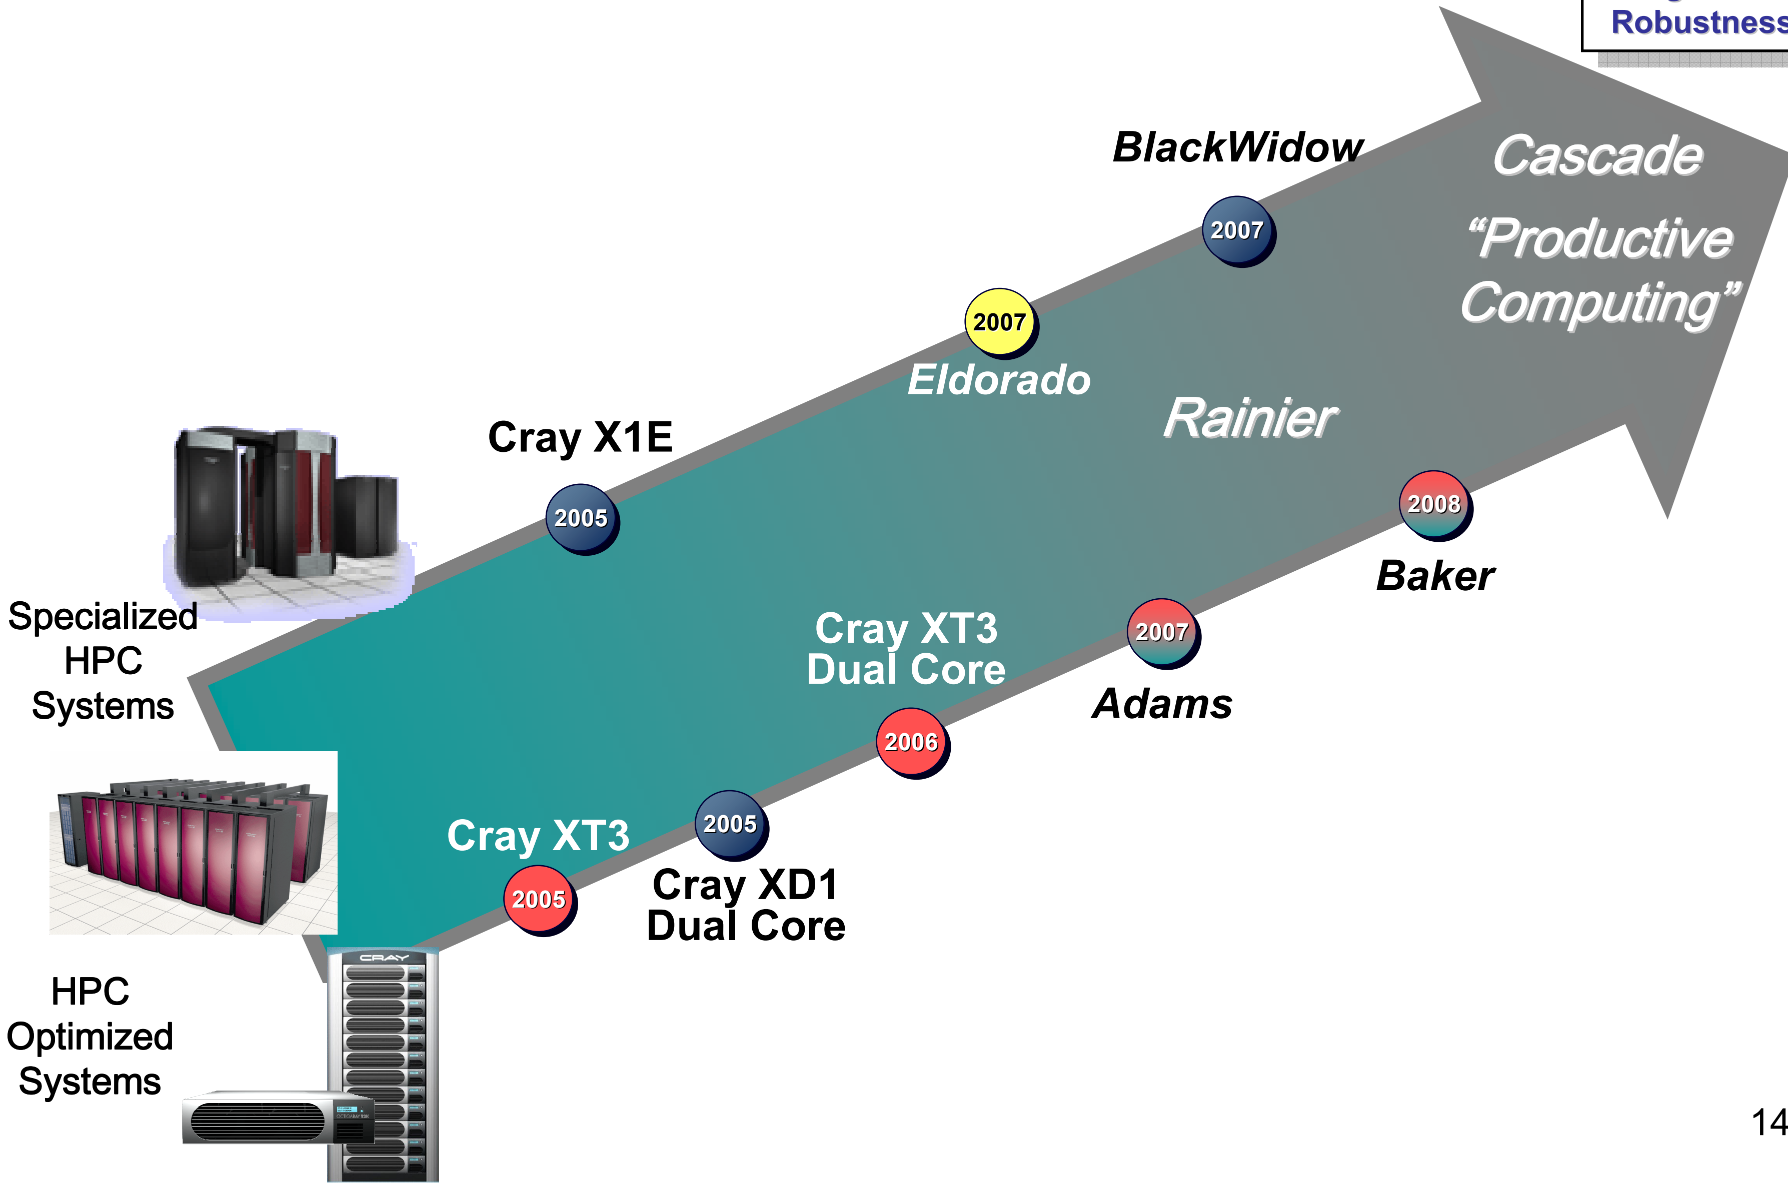

Future development of Jaguar

ORNL intends to expand Jaguar to a 100 Teraflop system in 2006 by doubling the number of cabinets and going to dual-core processors.

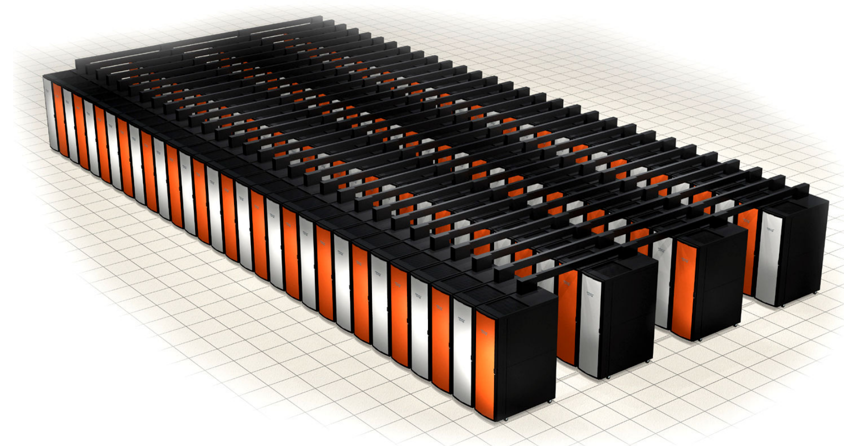

Pending approval by the U.S. Department of Energy and appropriation of the money by the Congress.

Cabinets	120
Compute Processors	Approximately 22,456
Memory	Approx. 45 TB (2 GB per processor)
Disk	480 TB
Peak Performance	100+ TeraFLOP/s

Where we plan to be in 2006

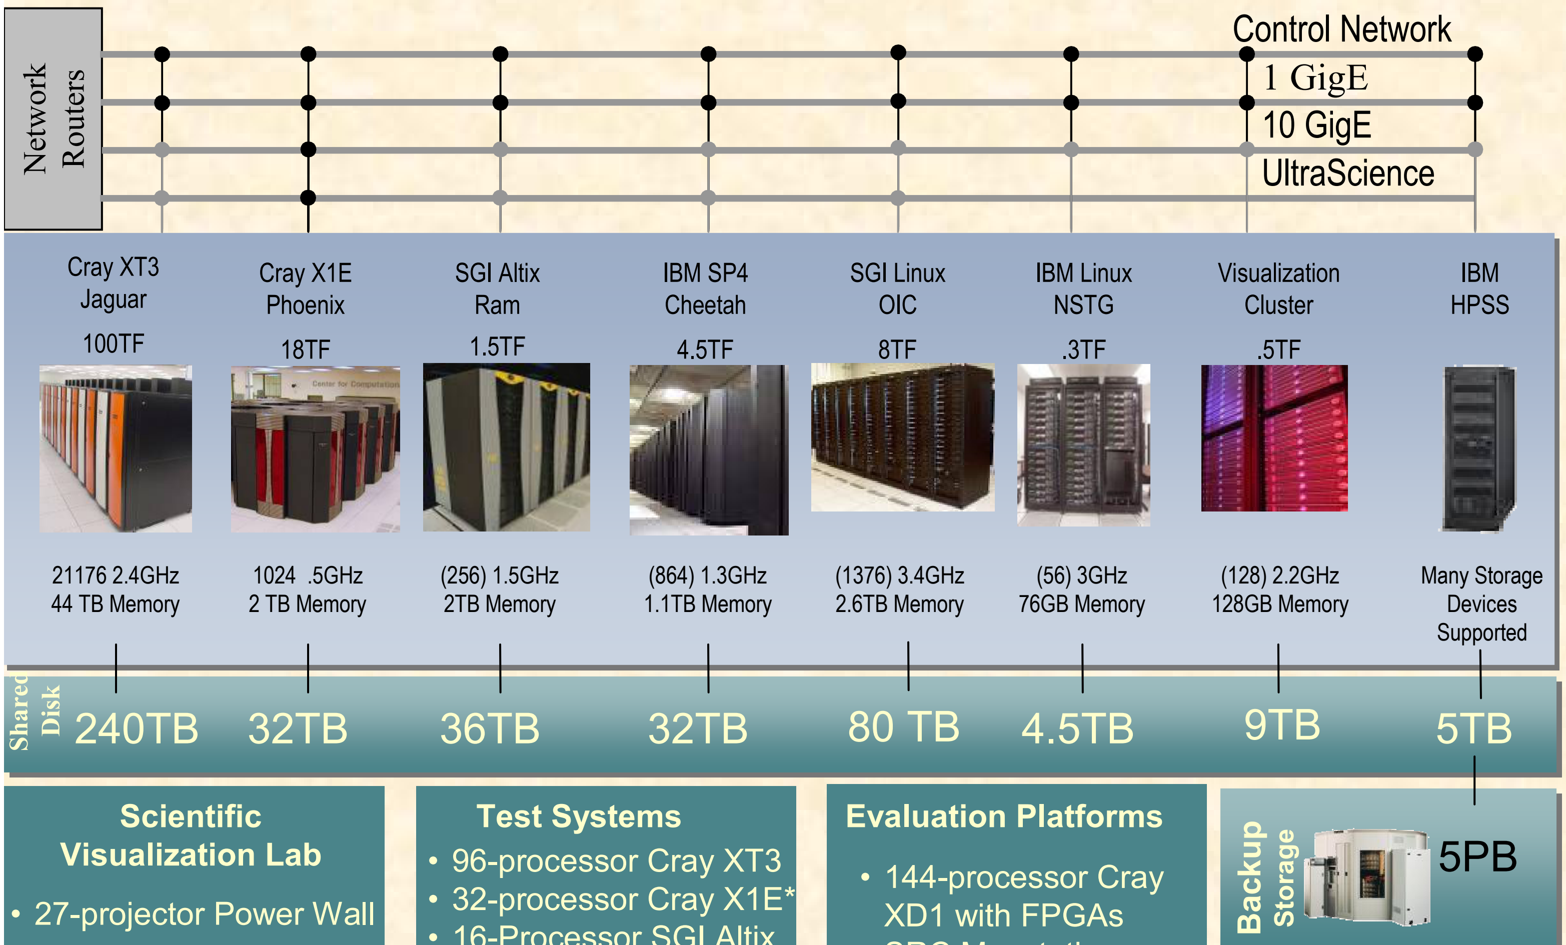

February 2006 Summary

8 Systems

Supercomputers
24,880 CPUs
52TB Memory
133 TFlops

438.5 TB

5 PB

Scientific Visualization Lab
- 27-projector Power Wall

Test Systems
- 96-processor Cray XT3
- 32-processor Cray X1E*
- 16-Processor SGI Altix

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

Impact of sustained exponential growth

- We are only beginning to realize the transforming power of computing as an enabler of innovation and discovery.
- A characteristic of exponential growth is that we will make as much progress in the next doubling cycle as we've made since the birth of the field:
  - 64, 8192, 1048576, 134217728; 1073741824
- We will have to run faster to stay in place.

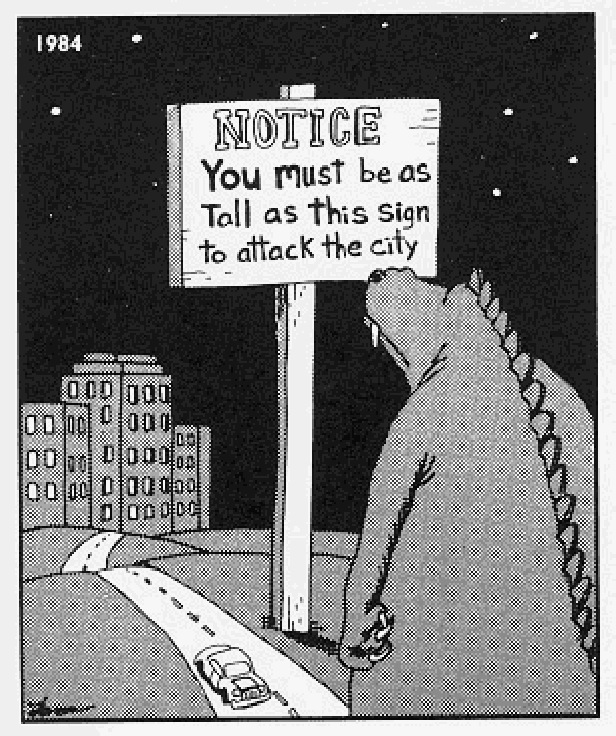

Information technology growth rate

Exponential growth creates the potential for revolutionary changes in what we do and how we do it

- **Processing power**
  - Doubling every 18 months
  - 60% improvement each year
  - Factor of 100 every decade
- **Disk Capacity**
  - Doubling every 12 months
  - 100% improvement each year
  - Factor of 1000 every decade
    - 10X as fast as processor performance!
- **Optical bandwidth today**
  - Doubling every 9 months
  - 150% improvement each year
  - Factor of 10,000 every decade
    - 10X as fast as disk capacity!
    - 100X as fast as processor performance!!

Computational science (processing power)

Data science (disk capacity, optical bandwidth)

Exponential growth in performance
Computing, storage, and networks

- **Rapid innovation in data storage and networking is driving innovation in science, engineering, media, etc.…**
- **Prioritize and integrate data and network science into our computational science initiative**

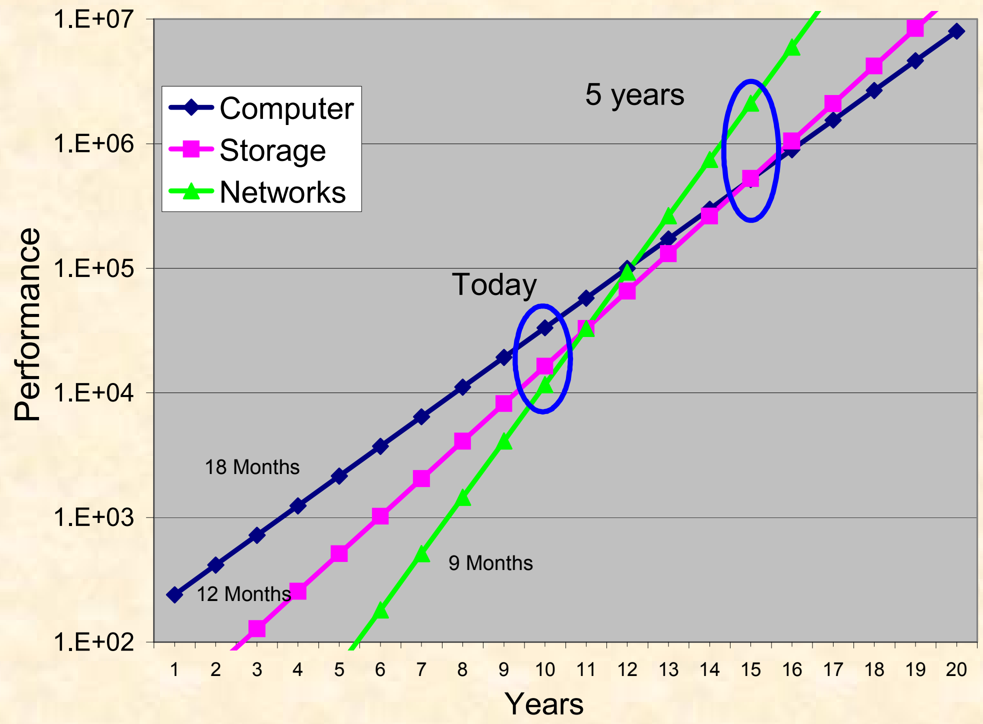

How Big Is Big?

- **Every 10X brings new challenges**
  - 64 processors was once considered large
    It hasn't been "large" for quite a while.
  - 1024 processors is today's "medium" size
  - 2048-16192 processors is today's "large"
    We are struggling even here.
- **100K processor systems**
  - are being designed/deployed
  - have fundamental challenges …

  … and no integrated research programs
- **Petascale data archives**
  - the "personal petabyte" very near
- **See recent PITAC report**
  - www.nitrd.gov

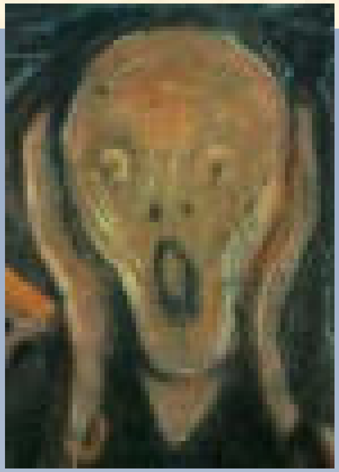

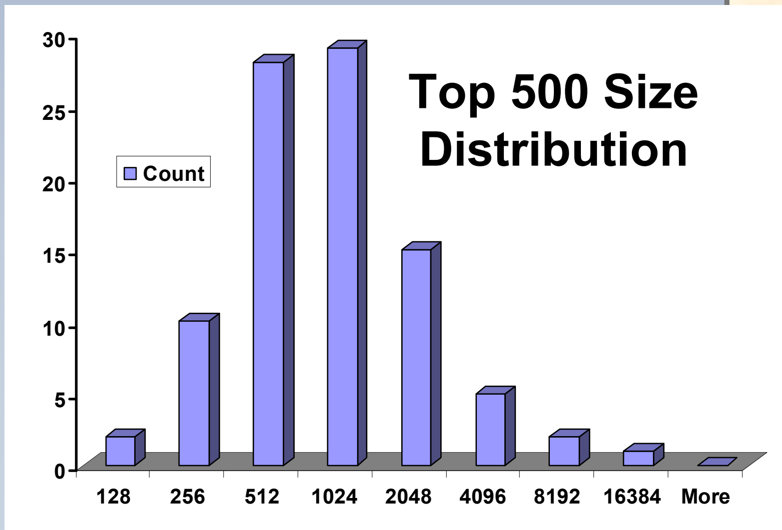

Preparing for Big:
Math and CS challenges

- **Theoretical Models (existing)**
  - May not perform well on petascale computers
  - May not have needed fidelity
  - May be inadequate to describe new phenomena revealed by experiment or simulation
- **Scientific Modeling and Simulation Codes (existing)**
  - Do not take advantage of new architectures (5%-10% of peak)
  - New computing capabilities lead to new simulation possibilities and, thus, new applications codes
- **Systems Software**
  - Vendor operating systems do not provide needed functionality
  - Systems software for petascale applications non-existent
    - Software to manage and visualize massive (petabyte) data sets
    - Software to accelerate development and use of petascale scientific applications
    - Techniques for porting software to the next generation inadequate
  - Few mathematical algorithms scale to thousand-million processors

ORNL computing infrastructure needs Power and cooling 2006 - 2011

- **Immediate need to add 8 MW to prepare for 2007 installs of new systems**
- **NLCF petascale system could require an additional 10 MW by 2008**
- **Need total of 40-50 MW for projected systems by 2011**
- **Numbers just for computers: add 75% for cooling**
- **Cooling will require 12,000 – 15,000 tons of chiller capacity**

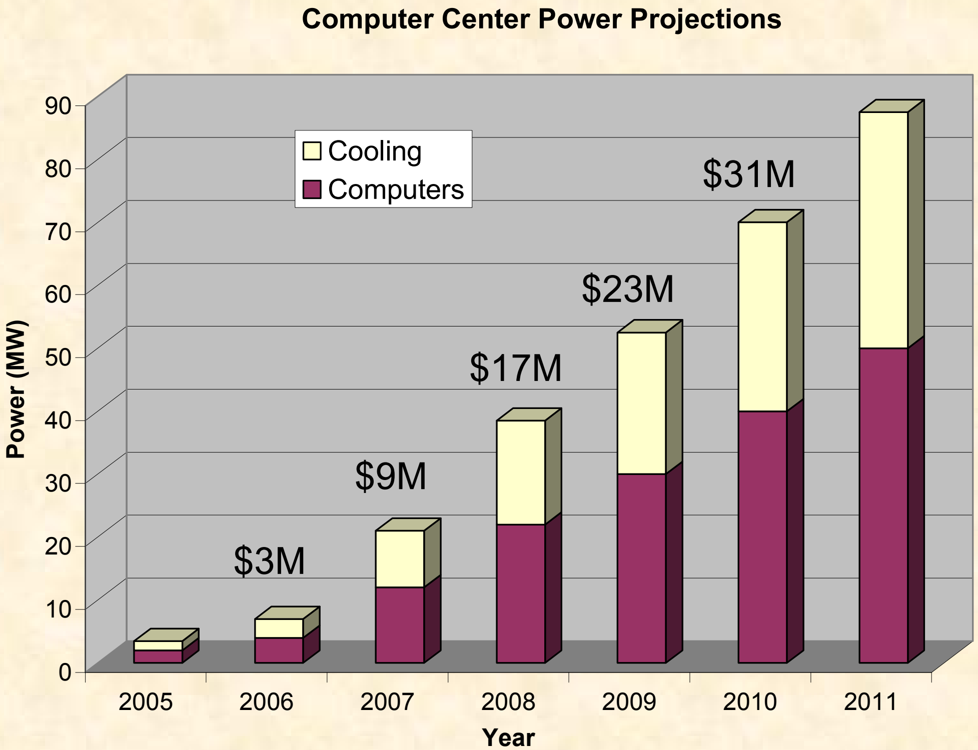

Cost estimates based on $0.05 kW/hr

But wait, there is the data explosion

Byte = 8 bits
Kilobyte = 10^3
Megabyte = 10^6
Gigabyte = 10^9
Terabyte = 10^{12}
Petabyte = 10^{15}
Exabyte = 10^{18}
Zettabyte = 10^{21}
Yottayte = 10^{24}

U.S. broadcast media	14,893 TB
Worldwide filmed content	420,254 TB
Worldwide printed content	1,633 TB
Internet	432,897 TB
World telephone calls	17,300,000 TB
Worldwide magnetic content	4,999,230 TB
Worldwide optical content	103 TB
Electronic flows of new info	17,905,340 TB

2002 = 5 Exabytes of NEW data
5,000,000,000,000,000,000

Data sciences becoming critical to scientific and knowledge discovery

Overwhelming quantities of data generated by:

- **Experimental facilities**
- **Computer simulations**
- **Satellites**
- **Sensors**
- **Etc.**

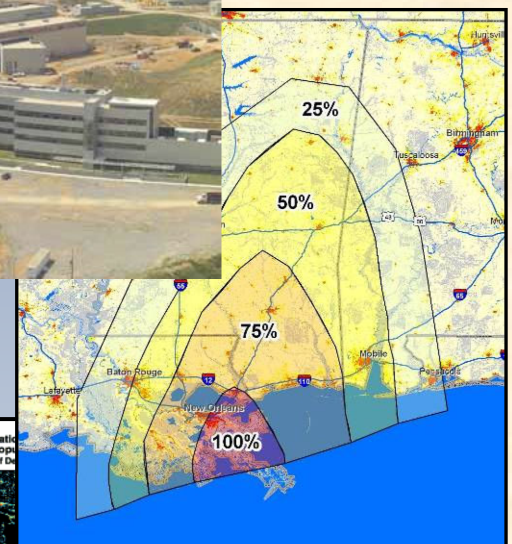

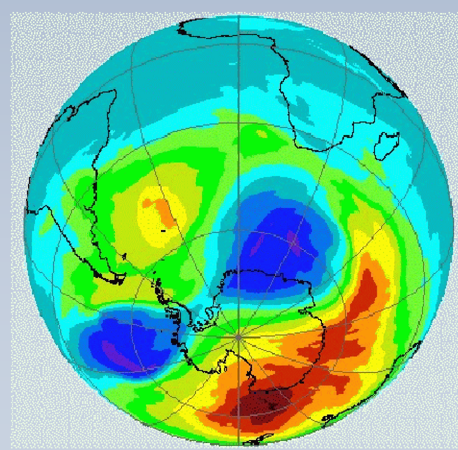

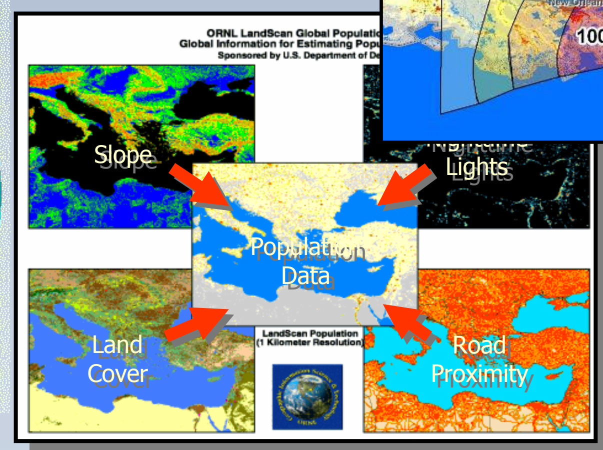

Data science challenges

Emerging scientific and sensor data sets have properties that will require new CS solutions to knowledge discovery.

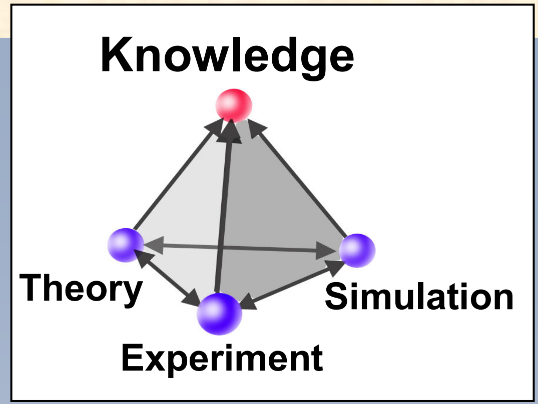

- **Massive (terabytes to petabytes)**
  - Existing methods do not scale in terms of time, storage, number of dimensions
  - Need scalable data analysis algorithms
- **Distributed (e.g., across grids, multiple files, disks, tapes)**
  - Existing methods work on a single, centralized dataset
  - Need distributed data analysis algorithm
- **Dynamically change with time**
  - Existing methods work with static datasets Any changes require re-computation
  - Need dynamic (updating and downdating) techniques
- **High-dimensional and Heterogeneous**
  - Cannot assume homogeneity or ergodicity
  - Need methods for handling heterogeneity and dimension reduction

Math and CS needs in Data Science

- Data management and data mining algorithms scalable to petabytes of distributed scientific data
- Fast retrieval of data subsets from storage systems: especially for tertiary storage
- Efficient transfer of large datasets between sites
- Easy navigation between heterogeneous, distributed data sources
- High performance I/O from leadership computers
- Visualization of massive data sets

Outline

- Who are we: NLCF at ORNL
- Science Motivators (first results)
- Allocations and Support for Break-Through Science
- 2006-2010 Roadmap
- **Wrapup**

Next 5 years:
Deliver breakthrough science and technology