Leadership Computing at the National Center for Computational Science: Transitioning to Heterogeneous Architectures

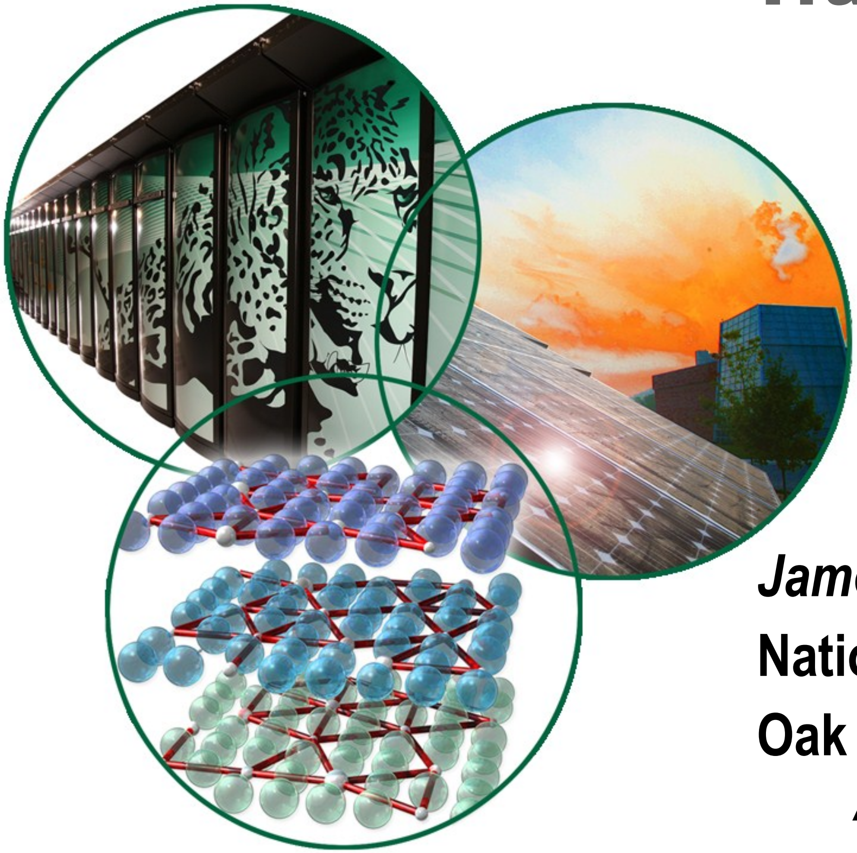

15th ECMWF Workshop on High Performance Computing in Meteorology
4 October 2012

James J. Hack, Director

National Center for Computational Sciences

Oak Ridge National Laboratory

Arthur S. "Buddy" Bland, OLCF Project Director

Bronson Messer, OLCF Scientific Computing

James Rogers, NCCS Director of Operations

Jack Wells, NCCS Director of Science

U.S. Department of Energy strategic priorities

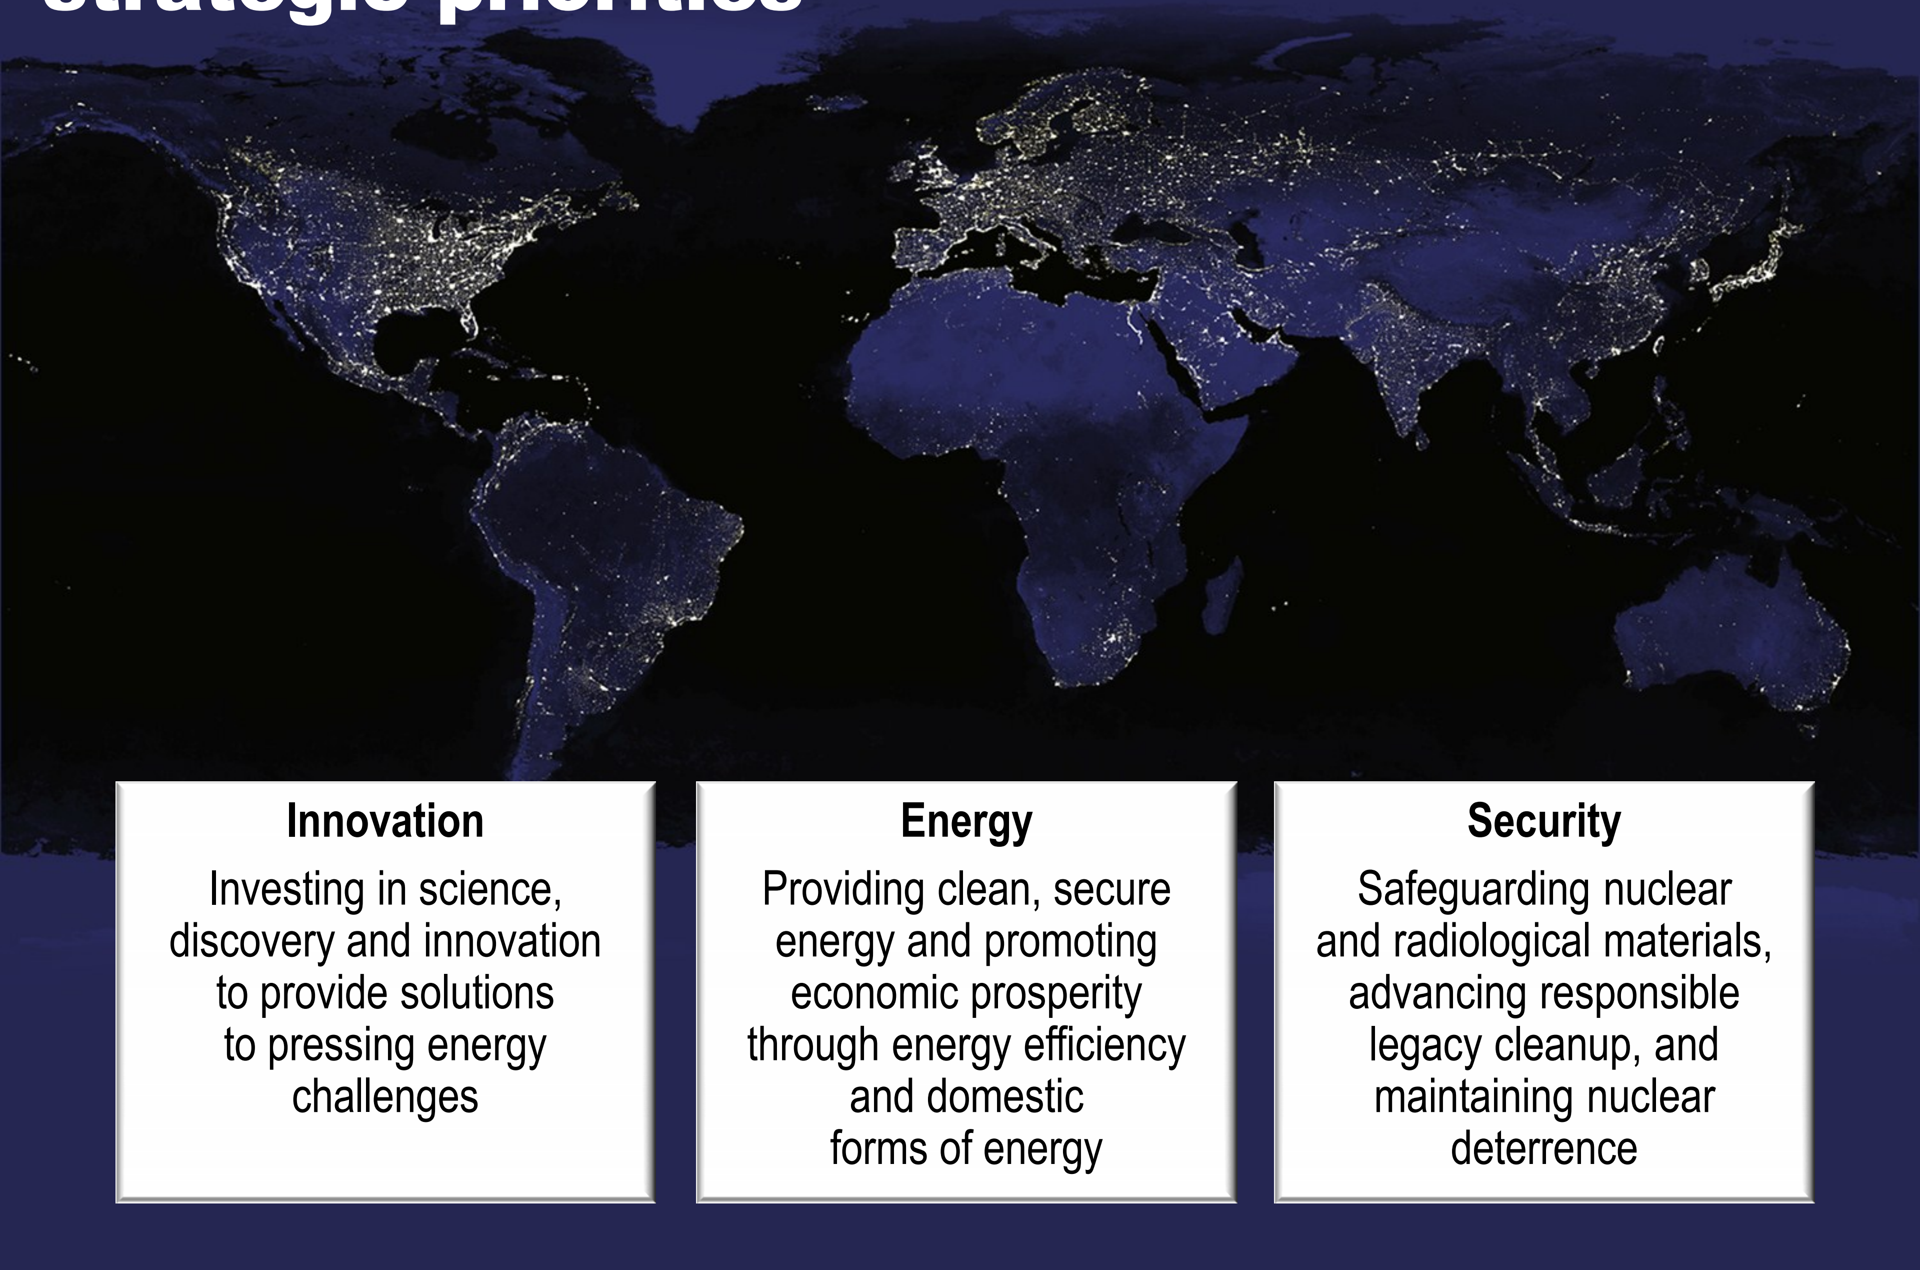

Innovation

Investing in science, discovery and innovation to provide solutions to pressing energy challenges

Energy

Providing clean, secure energy and promoting economic prosperity through energy efficiency and domestic forms of energy

Security

Safeguarding nuclear and radiological materials, advancing responsible legacy cleanup, and maintaining nuclear deterrence

Energy is the defining challenge of our time

Global energy consumption will increase 50% by 2030

- The major driver for
  - *Climate change*
  - National security
  - Economic competitiveness
  - Quality of life

- Incremental changes to existing technologies cannot meet this challenge
  - Transformational advances in energy technologies are needed
  - Transformational adaptation strategies will need to be implemented
  - Transformational changes to tools enabling virtualization of strategies

ORNL has a long history in High Performance Computing

ORNL has had 20 systems

on the TOP500® SUPERCOMPUTER SITES lists

1954
ORACLE

1969
IBM 360/9

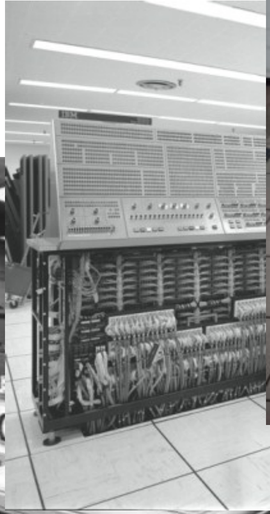

1985
Cray X-MP

1992-1995
Intel Paragons

1996-2002
IBM Power 2/3/4

2003-2005
Cray X1/X1E

2007
IBM Blue Gene/P

ORNL has increased system performance by 1,000 times since 2004

Hardware scaled from single-core through dual-core to quad-core and dual-socket , 12-core SMP nodes	Scaling applications and system software is the biggest challenge

- NNSA and DoD have funded much of the basic system architecture research
  - Cray XT based on Sandia Red Storm
  - IBM BG designed with Livermore
  - Cray X1 designed in collaboration with DoD

- DOE SciDAC and NSF PetaApps programs are funding scalable application work, advancing many apps
- DOE-SC and NSF have funded much of the library and applied math as well as tools
- Computational Liaisons key to using deployed systems

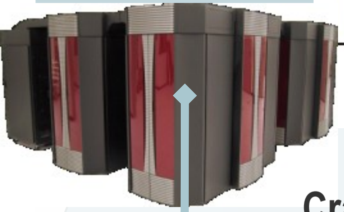

Cray X1
3 TF

2005

Cray XT3
Single-core
26 TF

2006

Cray XT3
Dual-Core
54 TF

2007

Cray XT4
119 TF

Cray XT4
Quad-Core
263 TF

2008

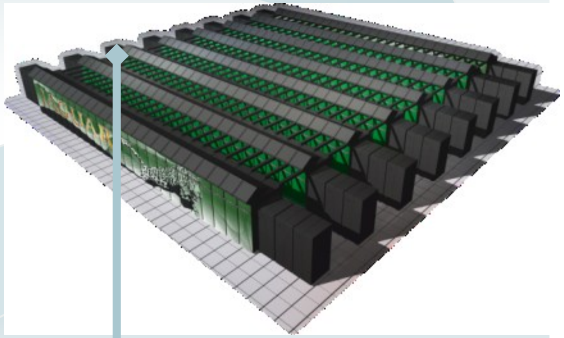

Cray XT5 Systems
12-core, dual-socket SMP
2000+ TF and 1000 TF

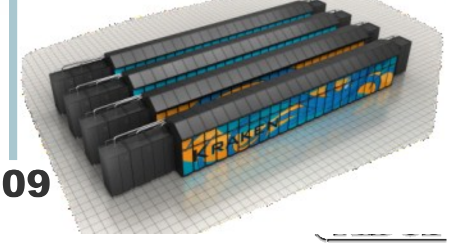

2009

Today, ORNL has one of the world's most powerful computing facilities

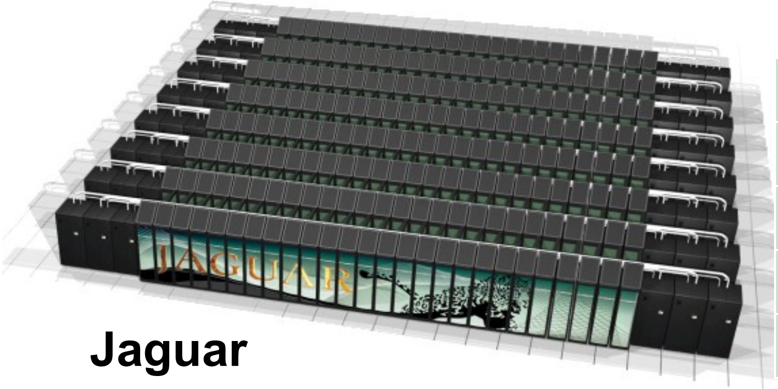

Jaguar

Peak performance	3.27 PF/s
Memory	600 TB
Disk bandwidth	> 240 GB/s
Square feet	5,000
Power	5.1 MW

TOP500 SUPERCOMPUTER SITES

#6

Dept. of Energy's most powerful computer

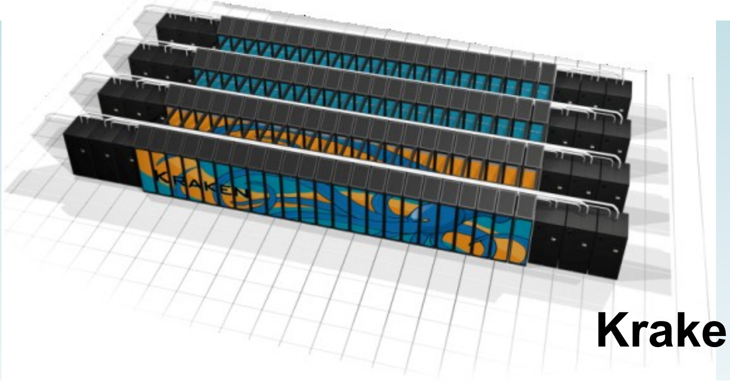

Kraken

Peak performance	1.17 PF/s
Memory	151 TB
Disk bandwidth	> 50 GB/s
Square feet	2,300
Power	3.5 MW

TOP500 SUPERCOMPUTER SITES

#21

National Science Foundation's most powerful computer

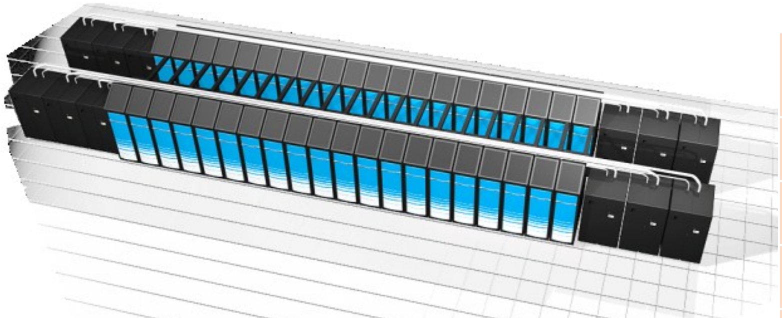

NOAA Gaea

Peak Performance	1.1 PF/s
Memory	248 TB
Disk Bandwidth	104 GB/s
Square feet	1,600
Power	2.2 MW

TOP500 SUPERCOMPUTER SITES

#33

National Oceanic and Atmospheric Administration's most powerful computer

NOAA Gaea System

C2 – 721TF Cray *XE6* (Jan 2012)

Cray XE6 LC (A Separate Compute Partition)

- ***Released to users in January 2012***
- 4,896 AMD 2.3 GHz 16-core Interlagos processors
- **78,336 compute cores**, 2,448 32-core nodes
- 156.7 TB DDR3 memory, 64 GB/node
- Peak performance: 721 TF
- Sustained performance: 565 TF
- Gemini High Speed Network
- Footprint: 26 cabinets
- Peak Electrical Consumption: 1,455 kVA

C1 – 386TF Cray *XE6* (September 2012)

- Architecturally identical to C2 system
- 2,624 AMD 2.3 GHz 16-core Interlagos processors
- **41,984 compute cores**, 1,312 32-core nodes
- 84 TB DDR3 memory, 64 GB/node, 2.0 GB/core
- Peak performance: 386 TF

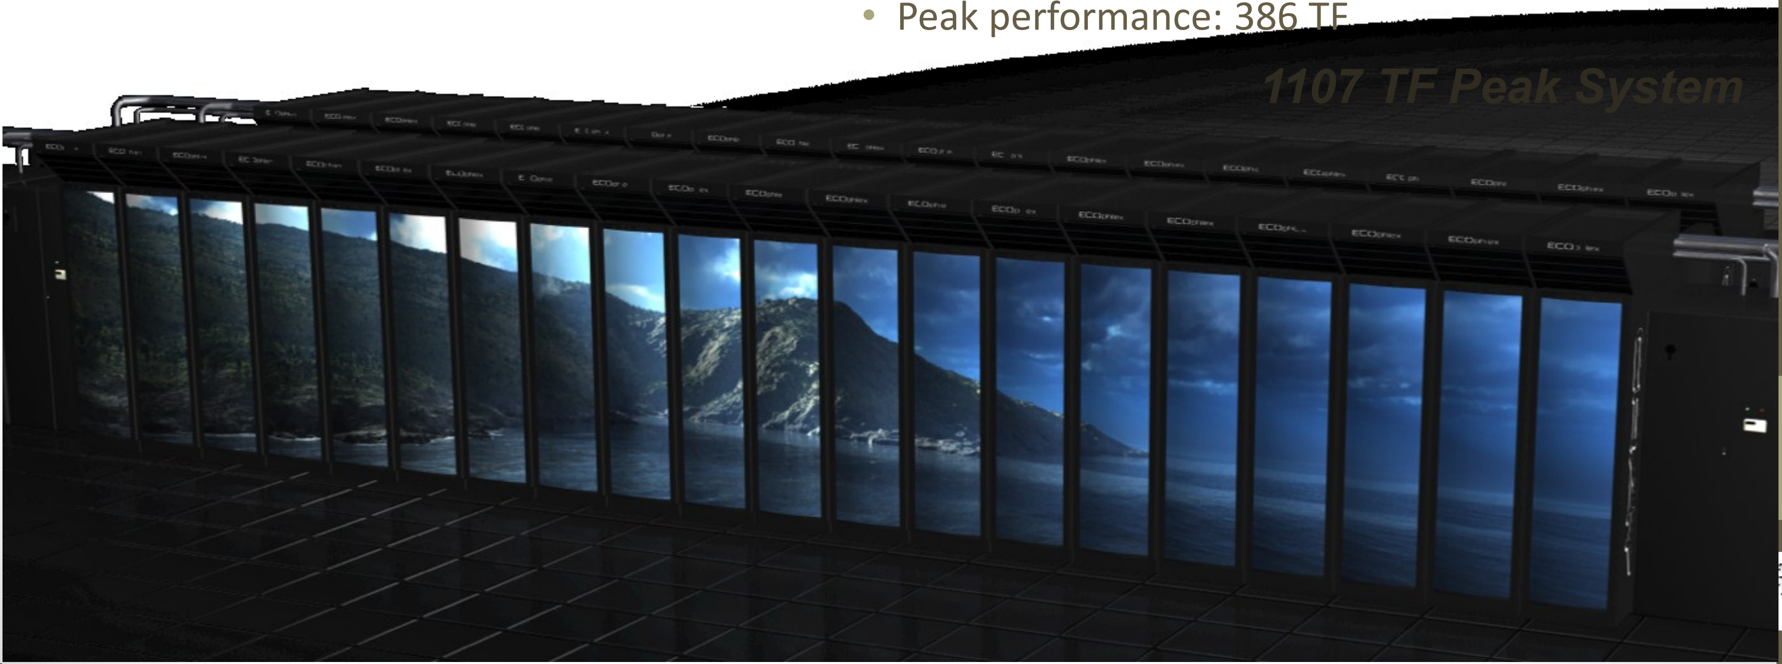

1107 TF Peak System

Oak Ridge Leadership Com Mission

The OLCF is a DOE Office of Science National User Facility whose mission is to enable breakthrough science by:

- Fielding the most powerful capability computers
  for scientific research,
- Building the required infrastructure to facilitate user access to these computers,
- Selecting a few time-sensitive problems of national importance that can take advantage of these systems,
- And partnering with these teams to deliver breakthrough science.

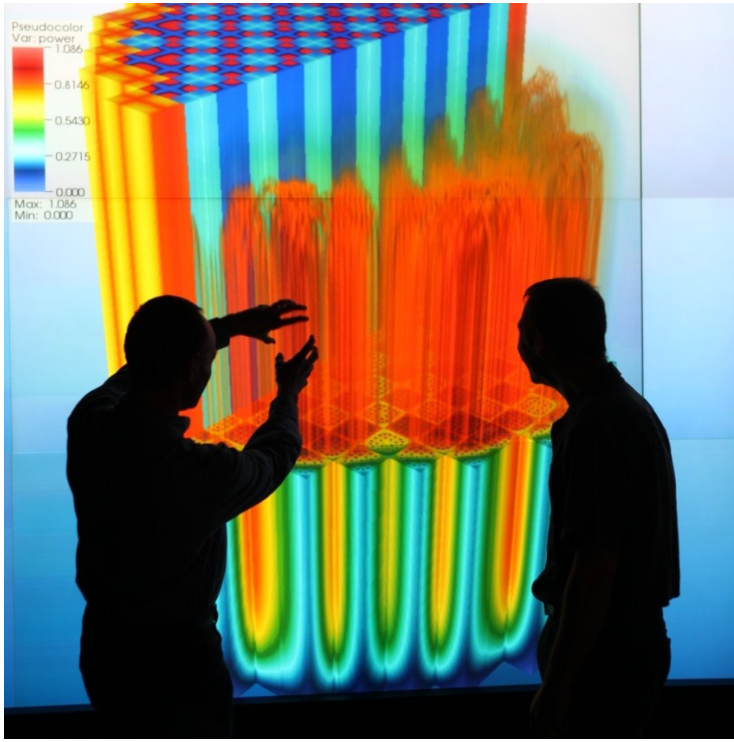

Innovative and Novel Computational Impact on Theory and Experiment

INCITE provides awards of time on the Oak Ridge and Argonne Leadership Computing Facility systems for researchers to pursue transformational advances in science and technology: **1.7 billion hours** were awarded in 2012.

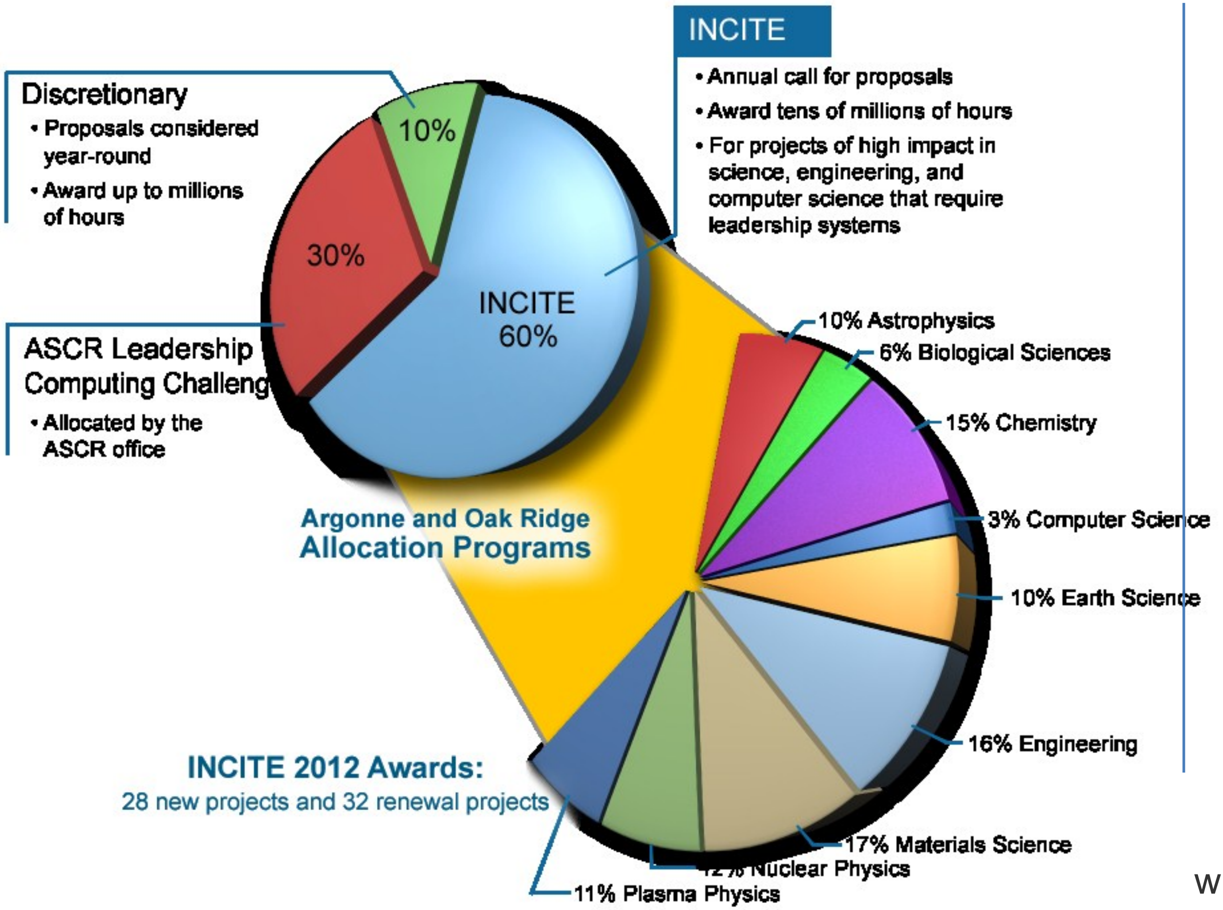

Discretionary
- Proposals considered year-round
- Award up to millions of hours

ASCR Leadership Computing Challeng
- Allocated by the ASCR office

INCITE
- Annual call for proposals
- Award tens of millions of hours
- For projects of high impact in science, engineering, and computer science that require leadership systems

INCITE 2012 Awards:
28 new projects and 32 renewal projects

Call for Proposals

The INCITE program seeks proposals for high-impact science and technology research challenges that require the power of the leadership-class systems. Allocations will be for calendar year 2013.

April 11 – June 27, 2012

Contact information
Julia C. White, INCITE Manager
whitejc@DOEleadershipcomputing.org

Computational Science has been our deliverable

OLCF has delivered important breakthroughs

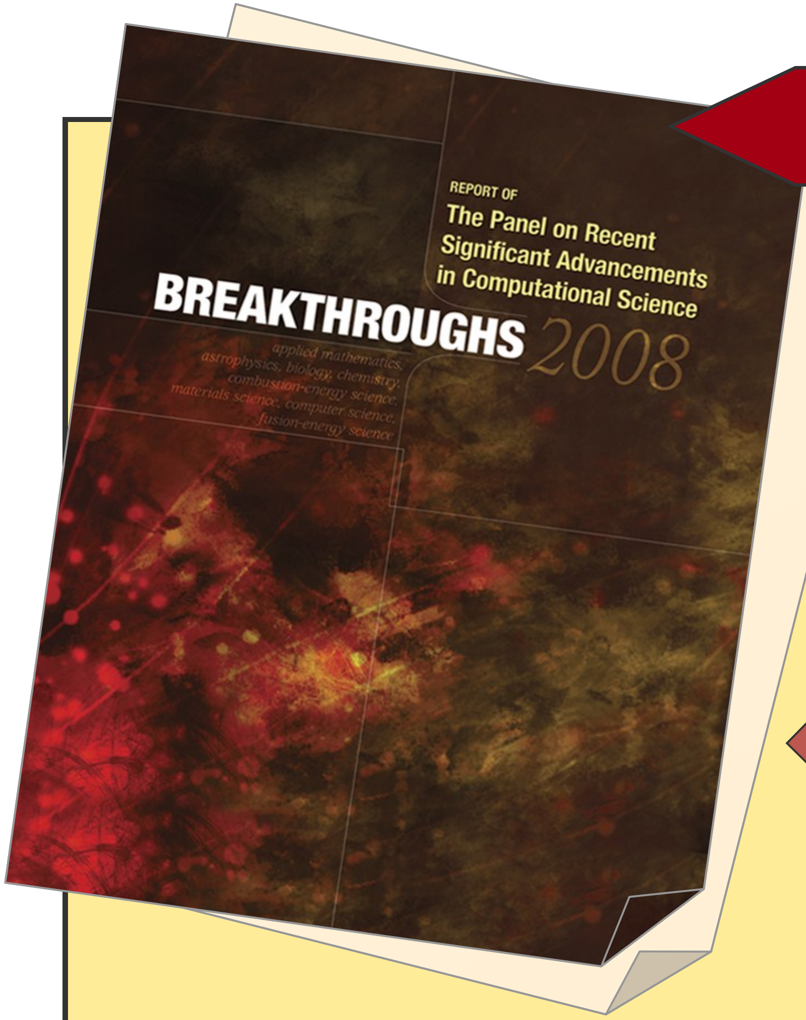

REPORT OF
The Panel on Recent
Significant Advancements
in Computational Science

BREAKTHROUGHS 2008

applied mathematics,
astrophysics, biology, chemistry,
combustion-energy science,
materials science, computer science,
fusion-energy science

OLCF makes mark on DOE Science

- **A mid-life LCF document showcasing ten scientific computing milestones includes five projects conducted at the NCCS**
- The document, titled “Breakthroughs 2008,” chronicles major recent advances in simulation under the auspices of the DOE’s ASCR program
- The OLCF accomplishments, which represent half of the total list, took place on the Cray Jaguar system

Breakthrough Science at Every Scale

Nuclear Physics
Vary et al., discover that nuclear structure and lifetimes using first-principles nuclear theory requires accounting for the complex nuclear interactions known as the three-body force

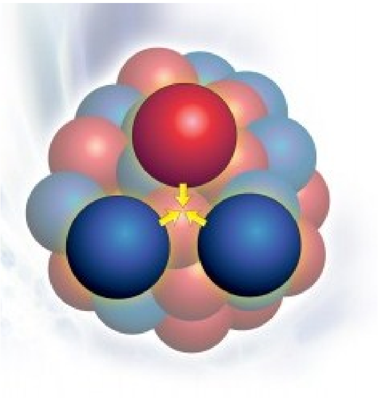

New Materials
Lopez-Bezanilla et al., discover that boron-nitride monolayers are an ideal dielectric substrate for future nanoelectronic devices constructed with graphene as the active layer

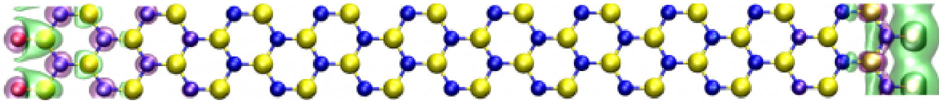

Biofuels
Smith et al., reveal the surface structure of lignin clumps down to 1 angstrom

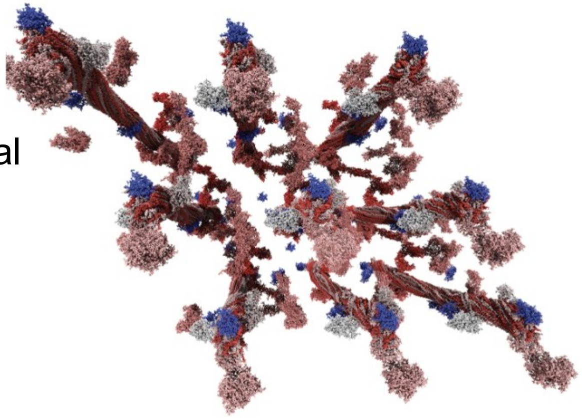

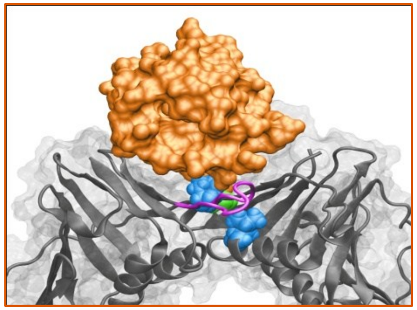

Biochemistry
Ivanov et al., illuminate how how DNA replication continues past a damaged site so a lesion can be repaired later

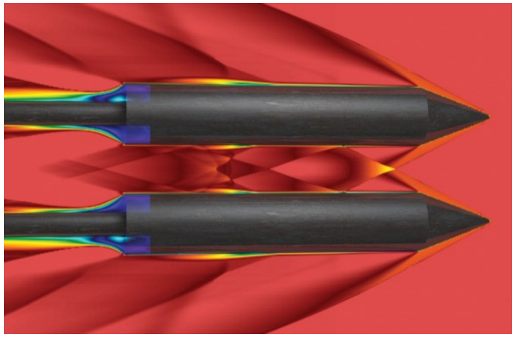

Design Innovation
Ramgen Power Systems accelerates their design of shock wave turbo compressors for carbon capture and sequestration

Turbo Machinery Efficiency
General Electric, for the first time, simulated unsteady flow in turbo machinery, opening new opportunities for design innovation and efficiency improvements.

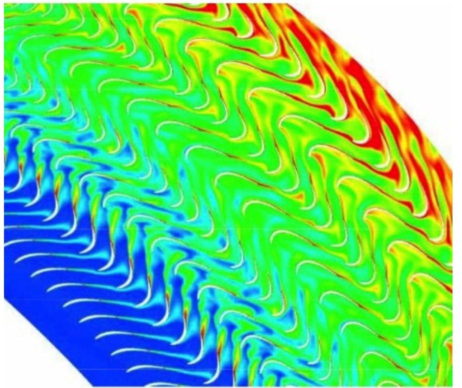

High Resolution Earth System Modeling

Objectives and Impact

- **Strategy**: Develop predictive global simulation capabilities for addressing climate change consequences
- **Goal**: Higher fidelity simulations with improved predictive skill on decadal time scales on regional space scales
- **Tactical**: Configurable high-resolution *scalable* atmospheric, ocean, terrestrial, cryospheric, and carbon component models to answer policy and planning relevant questions about climate change
- **Impact**: Exploration of renewable energy resource deployment, carbon mitigation strategies, climate adaptation scenarios (agriculture, energy and water resource management, protection of vulnerable infrastructure, national security)

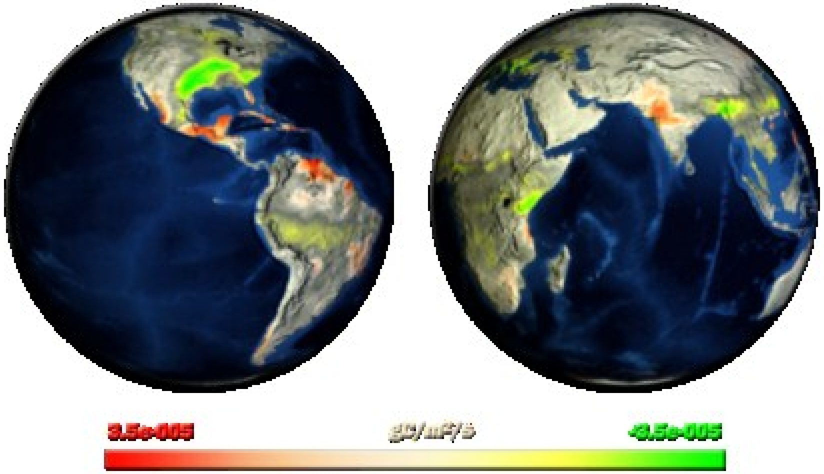

Net ecosystem exchange of CO_2

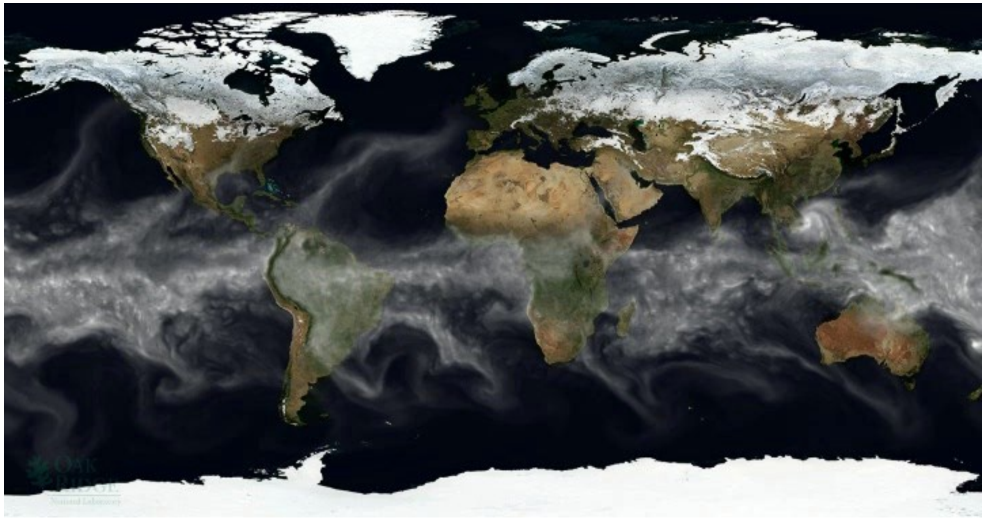

Mesoscale-resolved column integrated water vapor
Jaguar XT5 simulation

Eddy-resolved sea surface temperature
Jaguar XT5 simulation

Examples of climate consequences questions

- **Water Resources**
  - management and maintenance of existing water supply systems, development of flood control systems and drought plans
- **Agriculture and food security**
  - Erosion control, dam construction (irrigation), optimizing planting/harvesting times, introduction of tolerant/resistant crops (to drought, insect/pests, etc.)
- **Human health**
  - Public health management reform, improved urban and housing design, improved disease/vector surveillance and monitoring
- **Terrestrial ecosystems**
  - Improvement of management systems (deforestation, reforestation,…), development/improvement of forest fire management plans
- **Coastal zones and marine ecosystems**
  - Better integrated coastal zone planning and management
- **Human-engineered systems**
  - Better planning for long-lived infrastructure investments

New Scientific Opportunities

Continued scientific and computational science research coupled with advances in computational technology opens new opportunities

Exploring population dynamics in the context of adaptation to water availability, coastal vulnerability using Landscan population dataset

CSSEF Testbed

Computationally demanding, including workflow

Formal incorporation of uncertainty quantification

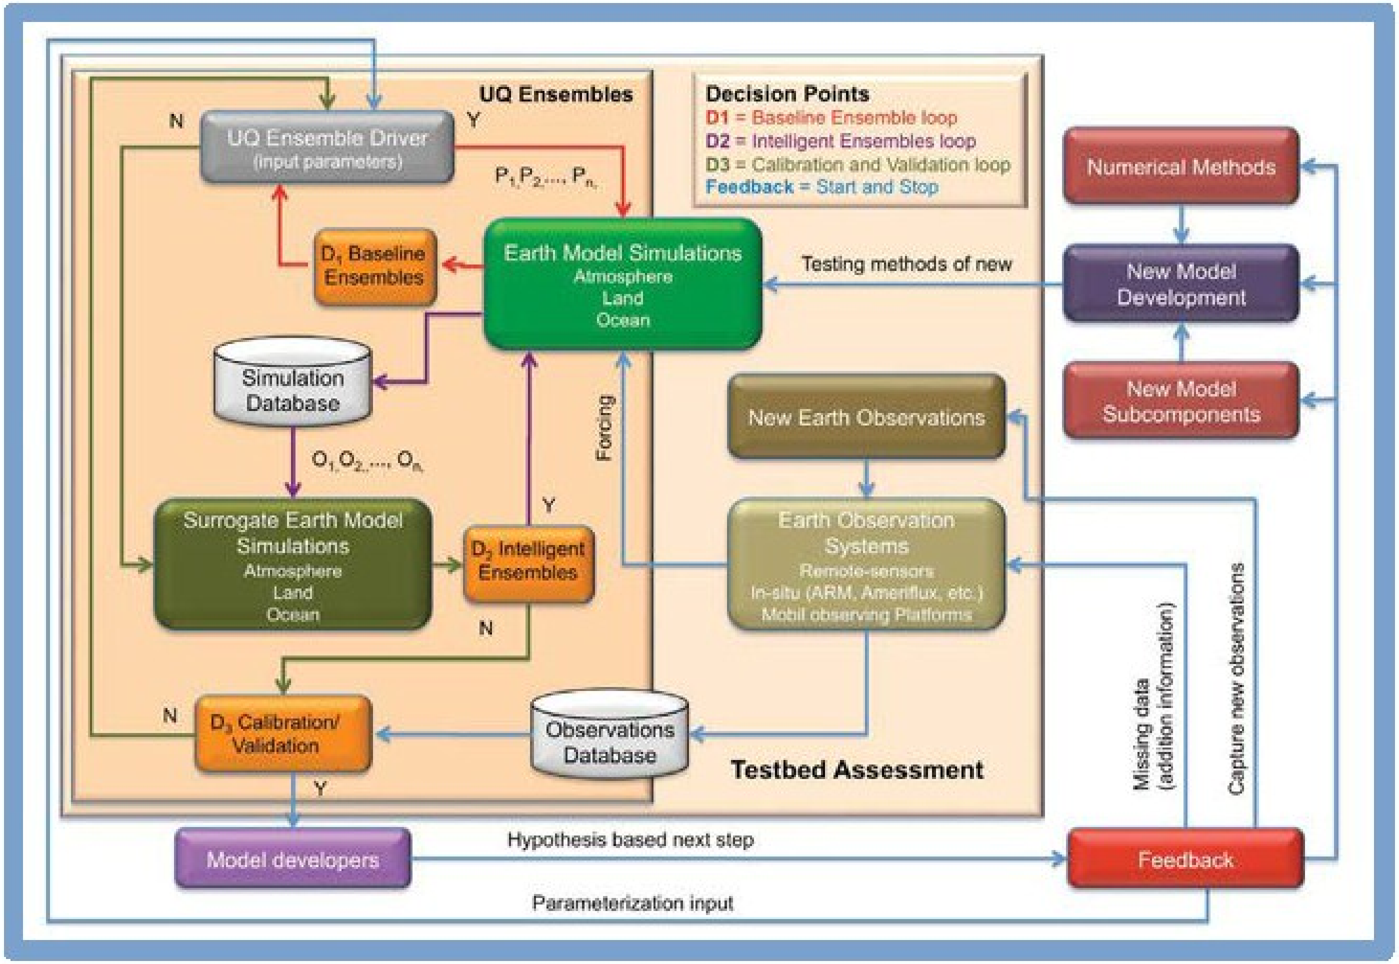

Advances in Predictive Science will need continued advances in computational facilities

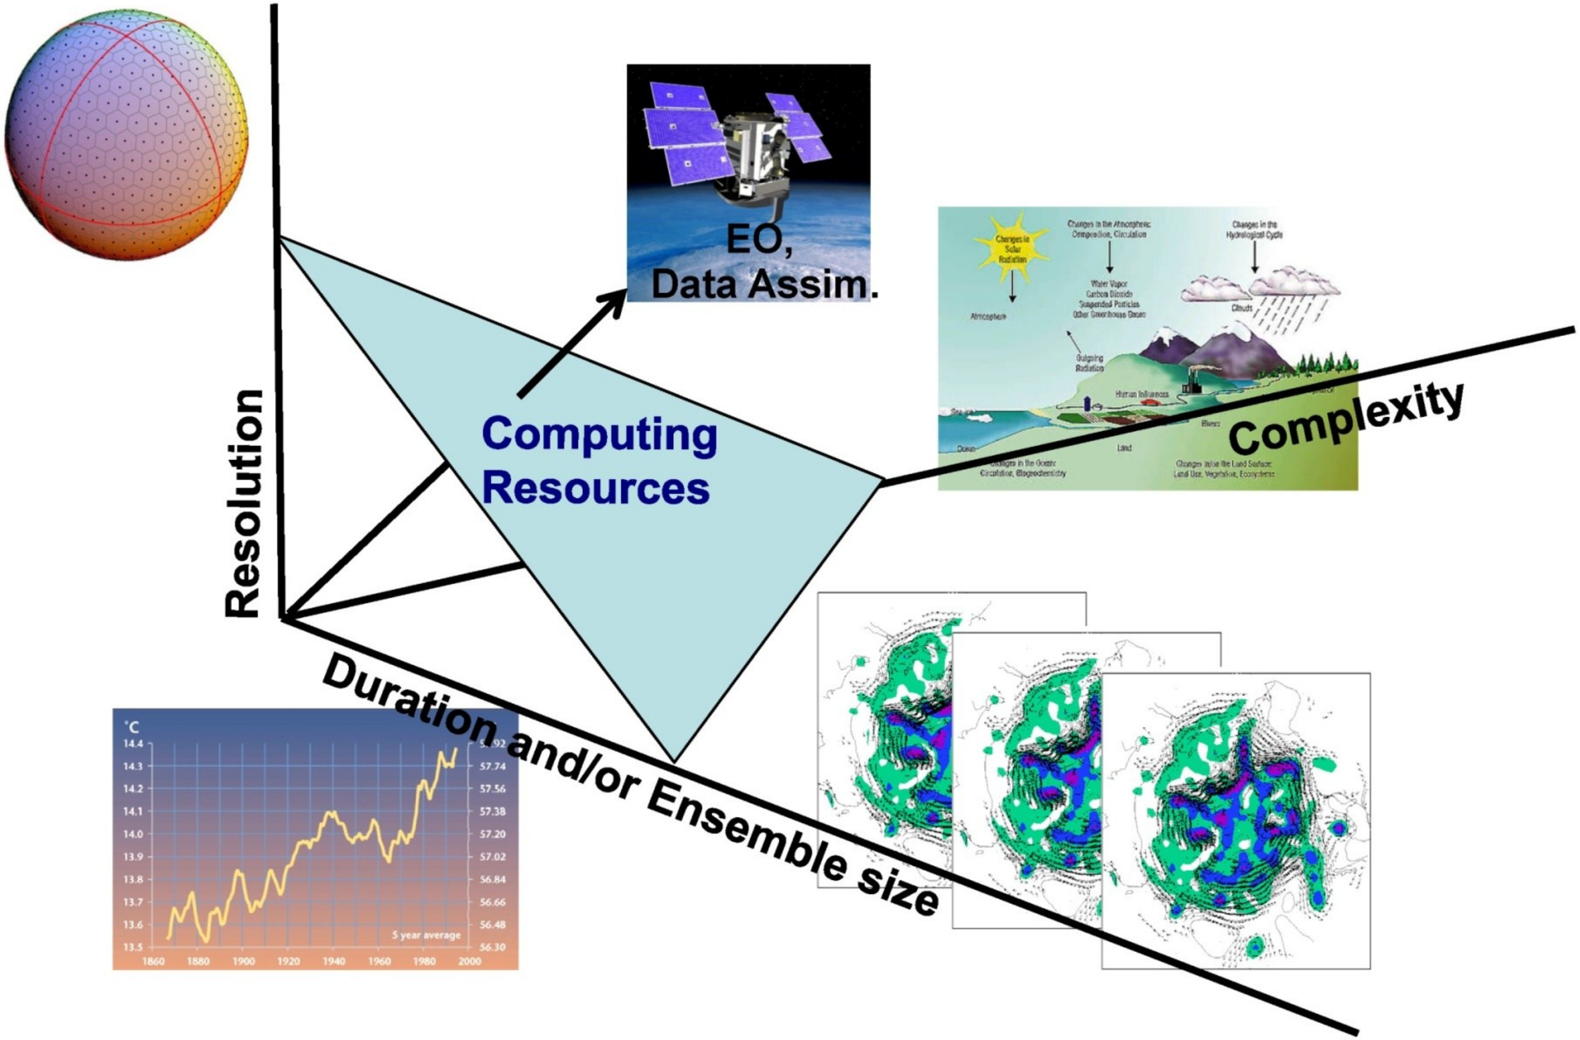

Resolution

Computing Resources

EO, Data Assim.

Complexity

Duration and/or Ensemble size

Broad requirement for advancing computational capability 1000x over the next decade

Mission: Deploy and operate the computational resources required to tackle global challenges	Vision: Maximize scientific productivity and progress on the largest scale computational problems
• Deliver transforming discoveries in climate, materials, biology, energy technologies, etc. • Ability to investigate otherwise inaccessible systems, from regional climate impacts to energy grid dynamics	• Providing world-class computational resources and specialized services for the most computationally intensive problems • Providing stable hardware/software path of increasing scale to maximize productive applications development

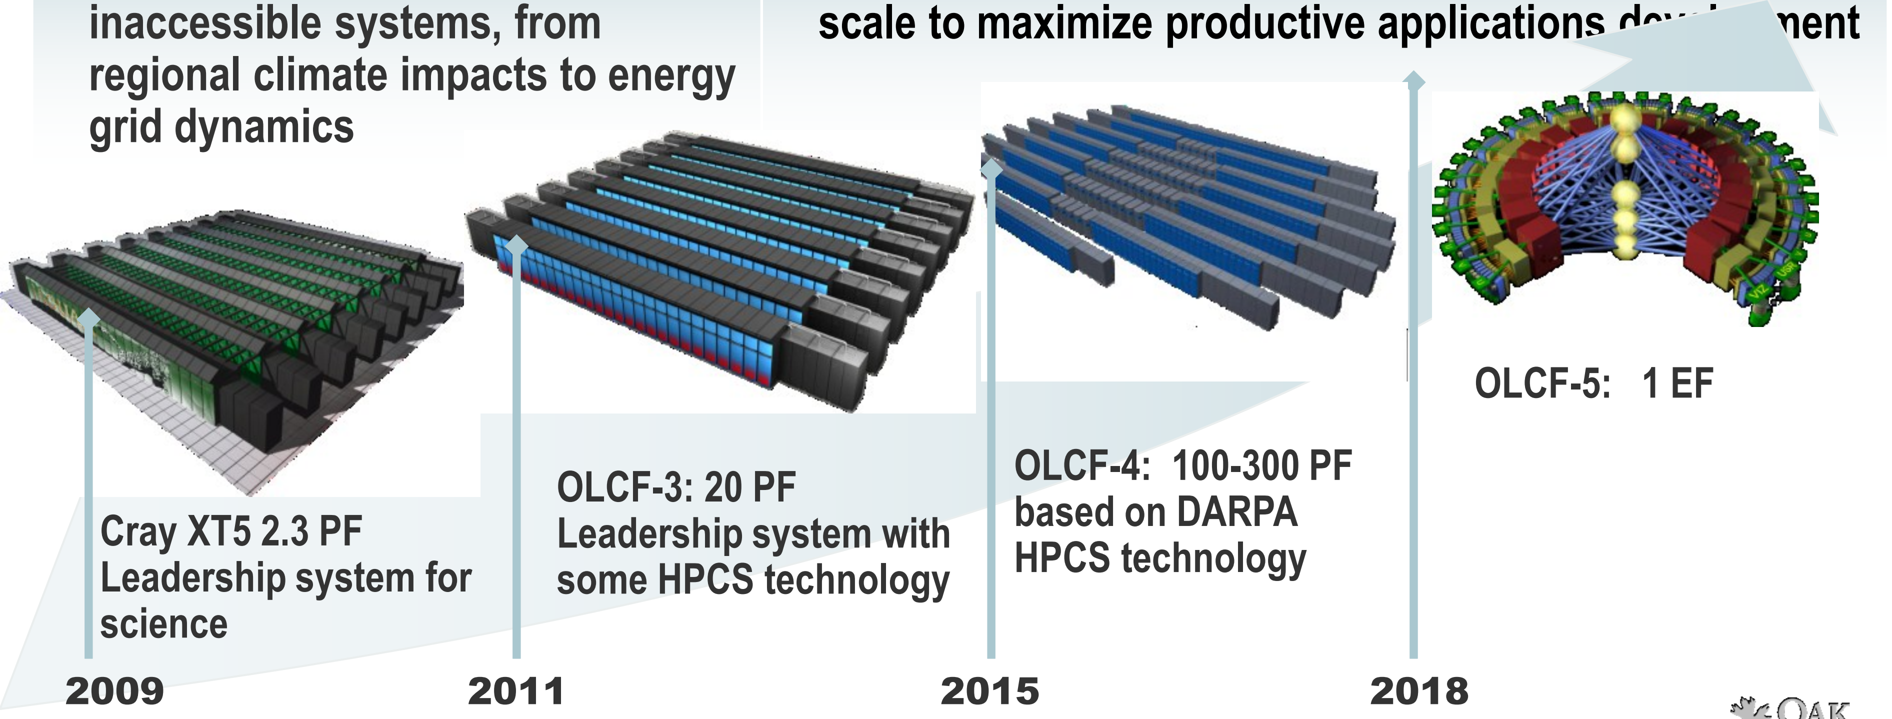

Cray XT5 2.3 PF
Leadership system for science

OLCF-3: 20 PF
Leadership system with some HPCS technology

OLCF-4: 100-300 PF
based on DARPA HPCS technology

OLCF-5: 1 EF

But, the High-Performance Computing Community has a Power Problem!

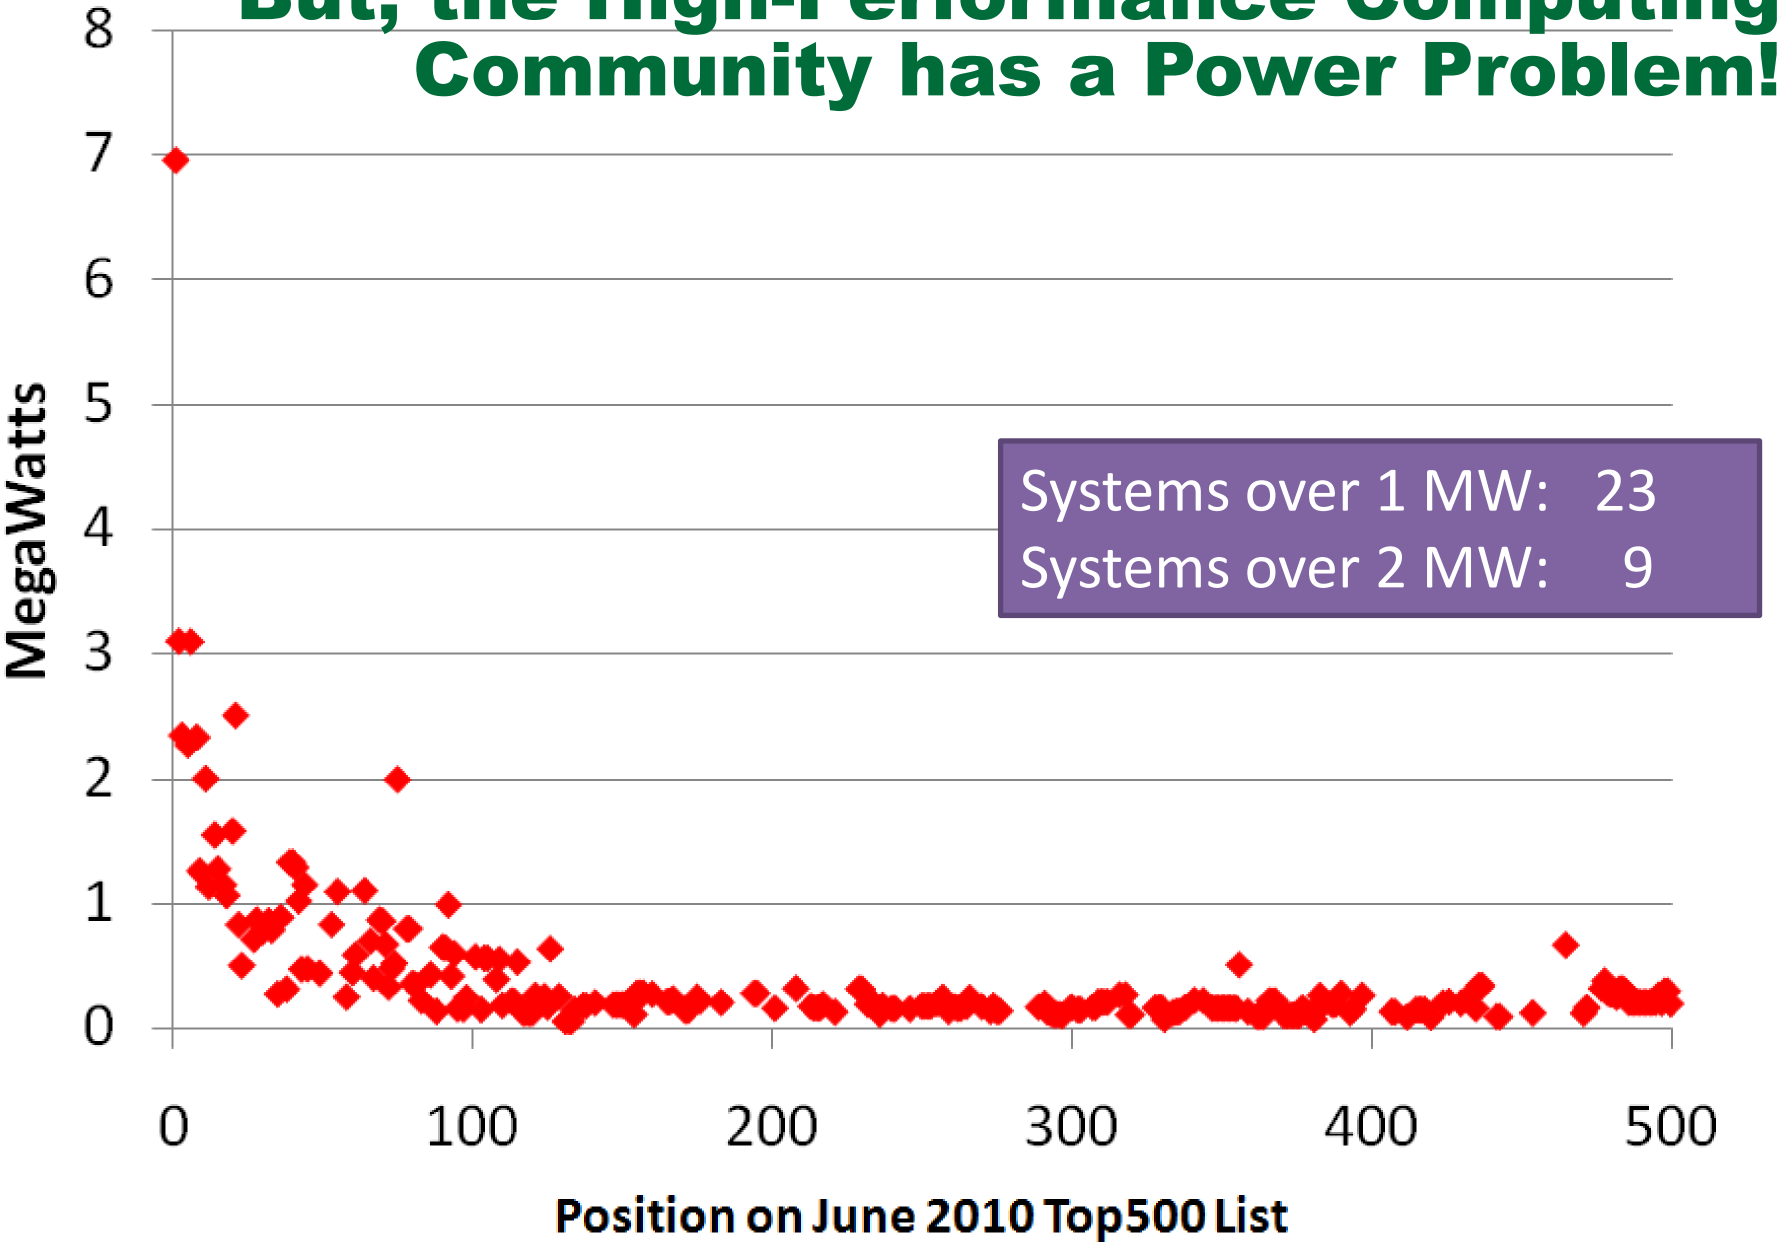

Systems over 1 MW: 23
Systems over 2 MW: 9

Data from http://www.Top500.org

Trends in power efficiency

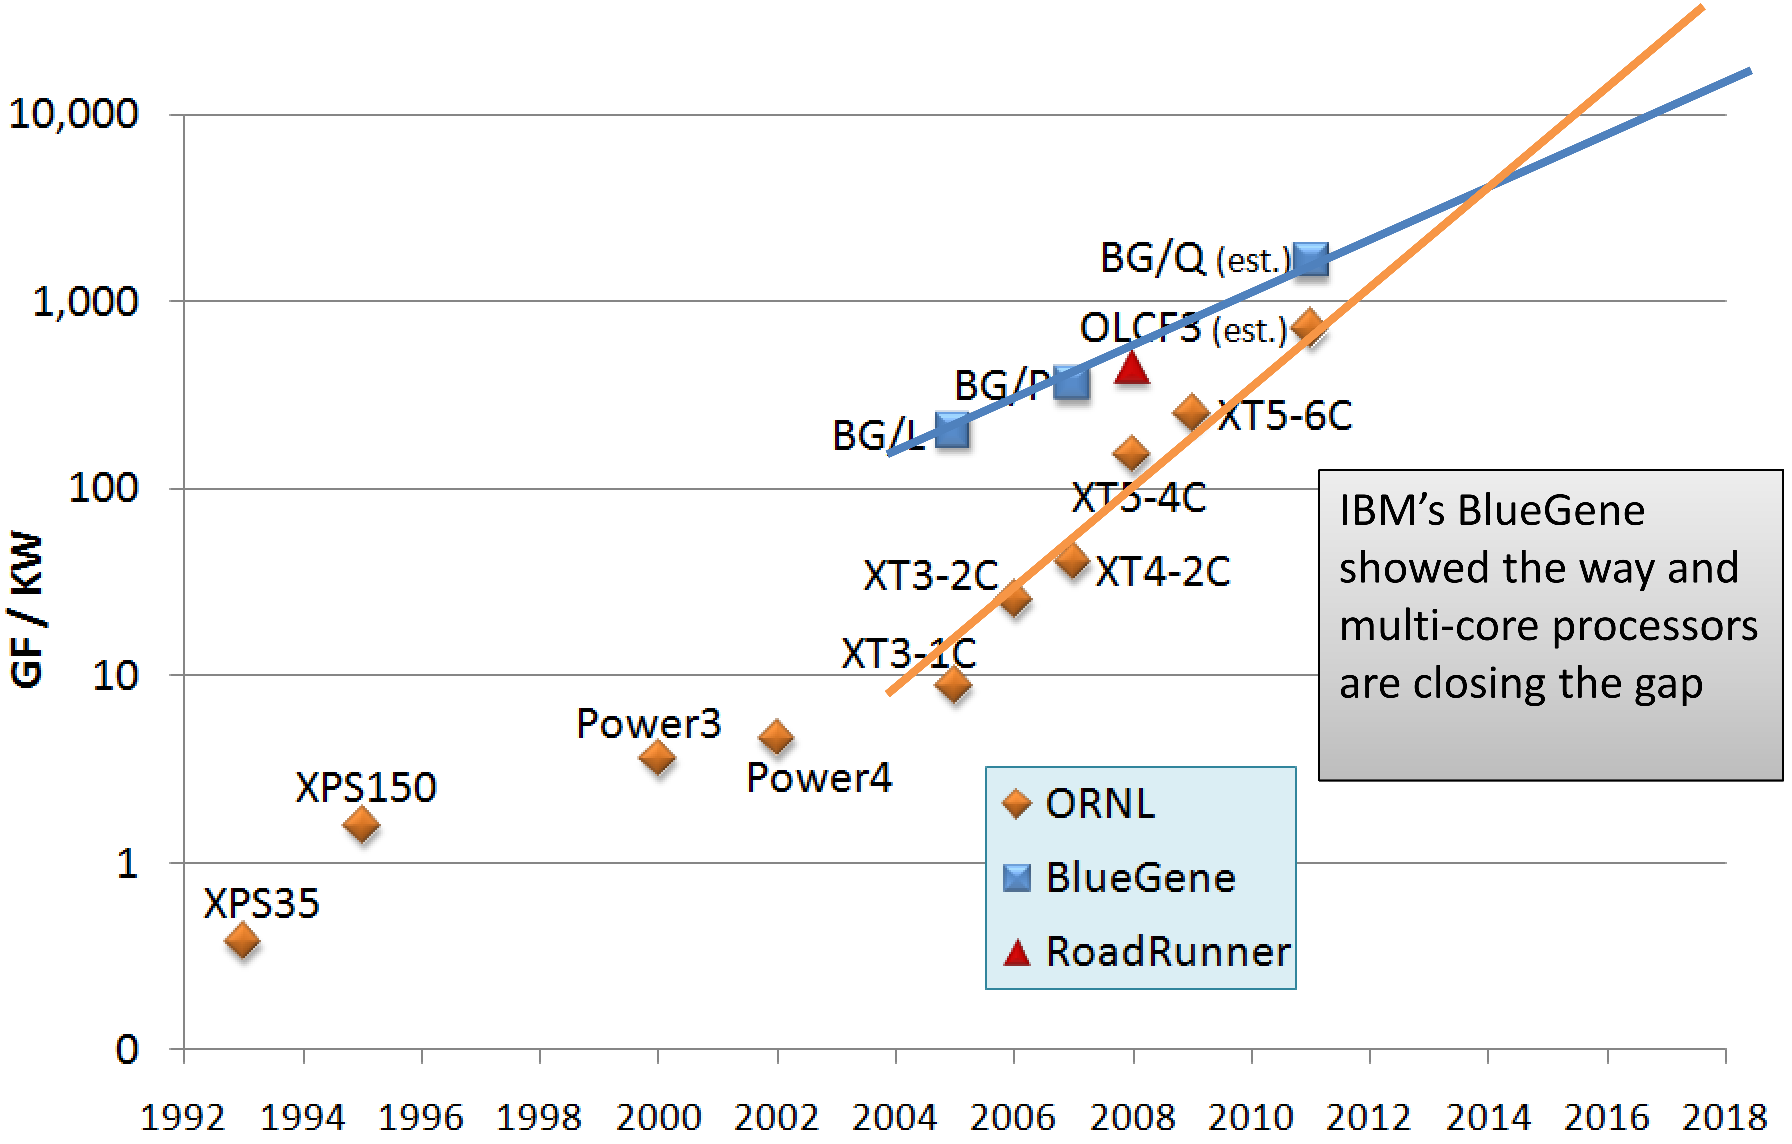

IBM's BlueGene showed the way and multi-core processors are closing the gap

Current Technology will require huge amounts of power for Exascale systems

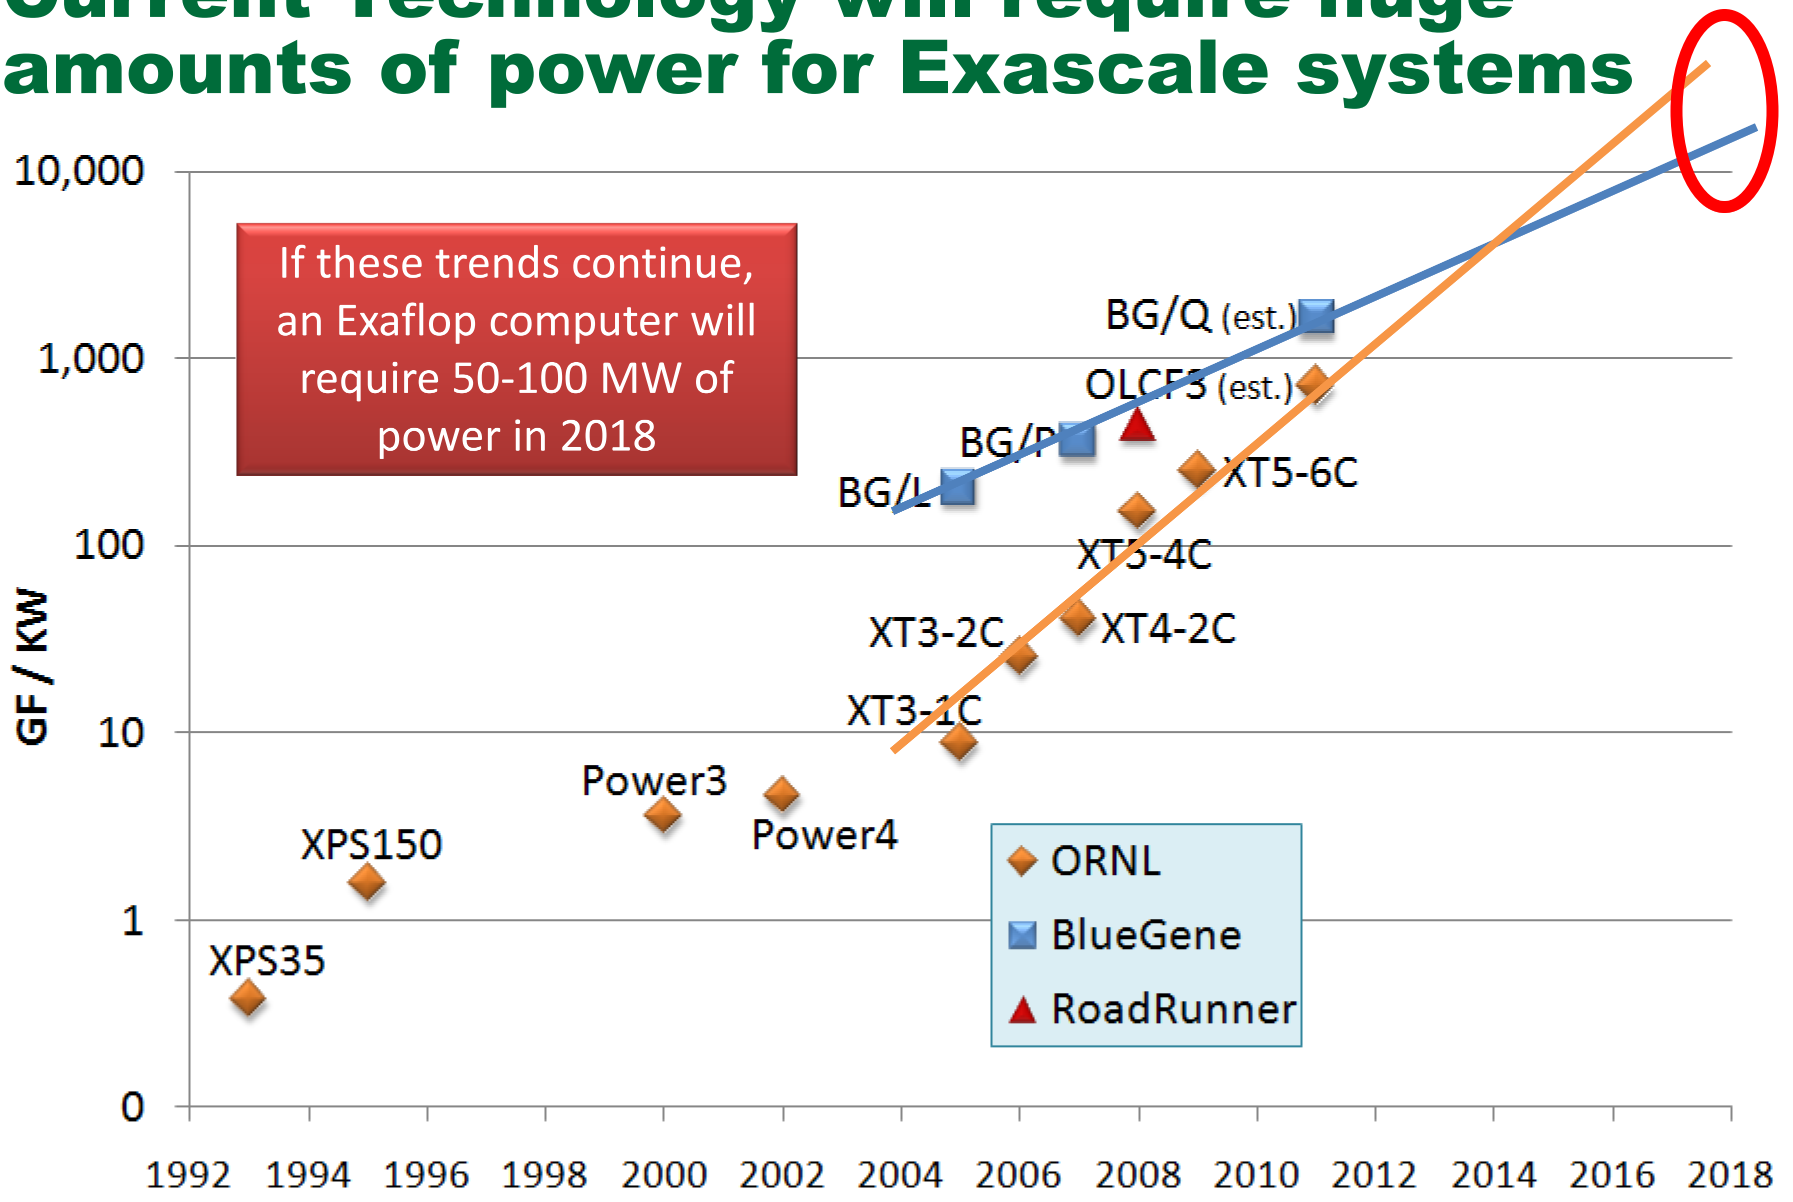

If these trends continue, an Exaflop computer will require 50-100 MW of power in 2018

Why has clock rate scaling ended?

Power $\propto$ Capacitance * Frequency * Voltage^2 + Leakage

- Traditionally, as Frequency increased, Voltage decreased, keeping the total power in a reasonable range
- But we have run into a wall on voltage
    - As the voltage gets smaller, the difference between a “one” and “zero” gets smaller. Lower voltages mean more errors.
    - While we like to think of electronics as digital devices, inside we use analog voltages to represent digital states. But this is imperfect!

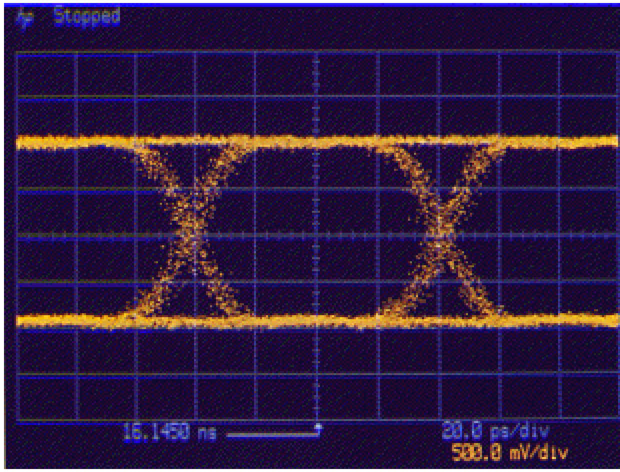

- Capacitance increases with the complexity of the chip
- Total power dissipation is limited by cooling

Power to move data

Energy_to_move_data $\propto$ ***bitrate * length***^{2} ***/ cross_section_area_of_wire***

- Energy consumed increases proportionally to the bit-rate, so as we move to ultra-high bandwidth links, the power requirements will become an increasing concern.
- Energy consumption is highly distance-dependent (the square of the length term), so bandwidth is likely to become increasingly localized as power becomes a more difficult problem.
- Improvements in chip lithography (making smaller wires) will not improve the energy efficiency or data carrying capacity of electrical wires.

D. A. B. Miller and H. M. Ozaktas, “Limit to the Bit-Rate Capacity of Electrical Interconnects from the Aspect Ratio of the System Architecture,” Journal of Parallel and Distributed Computing, vol. 41, pp. 42-52 (1997) article number PC961285.

And Then There's the Memory Wall

"FLOPS are 'free'. In most cases we can now compute on the data as fast as we can move it." - Doug Miles, The Portland Group

What we observe today:

- Logic transistors are free
- The von Neumann architecture is a bottleneck
- Exponential increases in performance will come from increased concurrency not increased clock rates if the cores are not starved for data or instructions

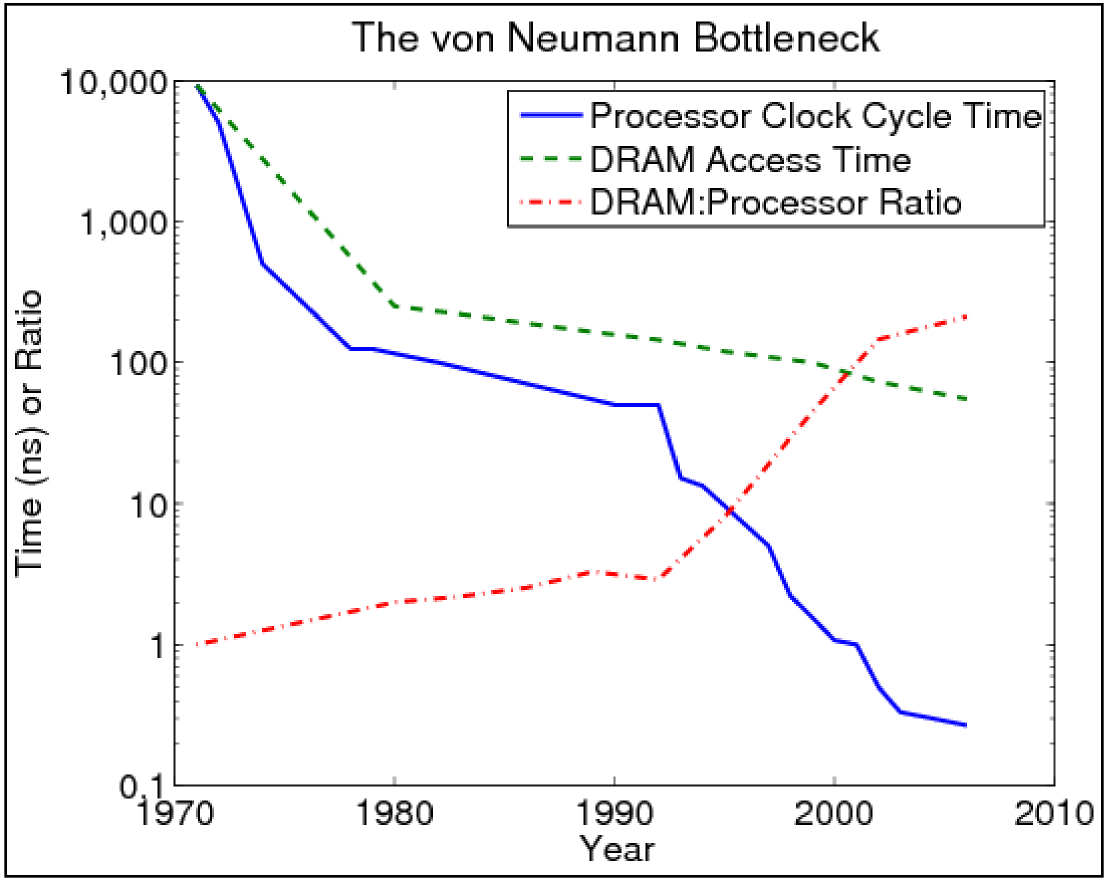

We also observe a scaling problem!

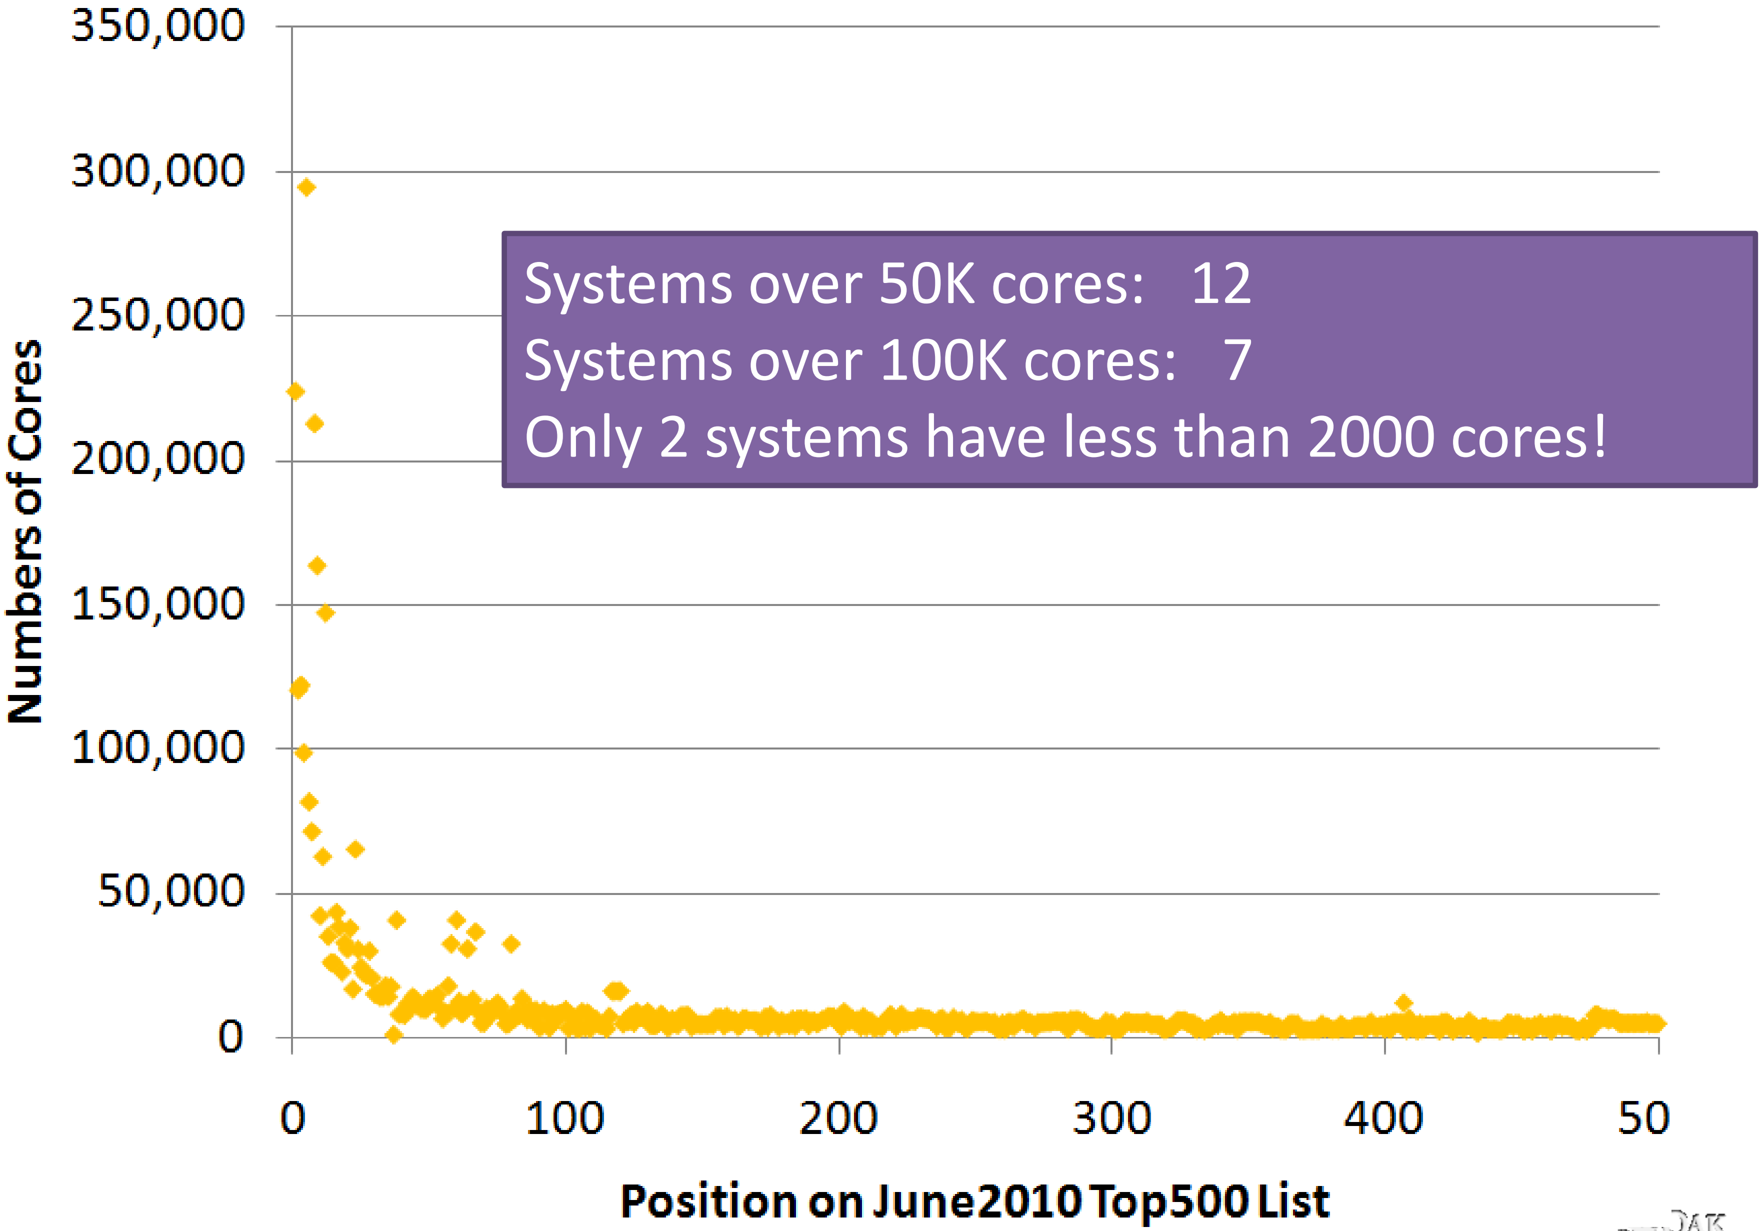

Systems over 50K cores: 12
Systems over 100K cores: 7
Only 2 systems have less than 2000 cores!

Exascale Systems

- 1-10 billion way parallelism
  - Requires hierarchical parallelism to manage
    - MPI between nodes
    - OpenMP or other threads model on nodes
    - SIMD / Vector parallelism within thread
- Power will dominate architectures
    - Takes more power to go to memory for data than to recompute it
- Traditional “balance ratios” can’t continue to be met
    - Memory size is not growing as fast as processor performance
    - Memory bandwidth is growing even more slowly
    - So compute becomes relatively cheaper while memory and data movement becomes the expensive resource in a system

Technology transitions have been observed over time

Logistic change is characterized by an initial period of slow growth, followed by a period of exponential growth, then a point of inflection, and finally a period of asymptotic growth as the technology approaches a limit. This pattern of change was first observed in population studies [28], and it has since been found to be descriptive of change in a remarkably diverse set of circumstances, including technological evolution in general and the evolution of electronic and computer technologies in particular.

Worlton (1988)

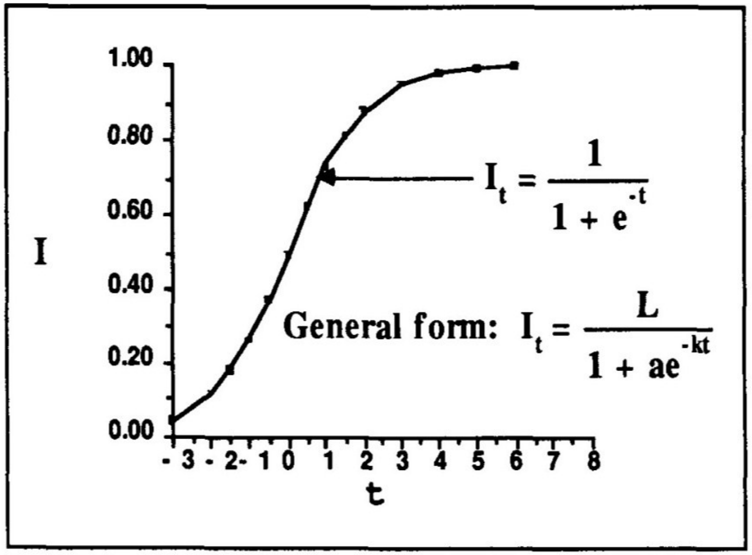

Figure 9. Logistic change.

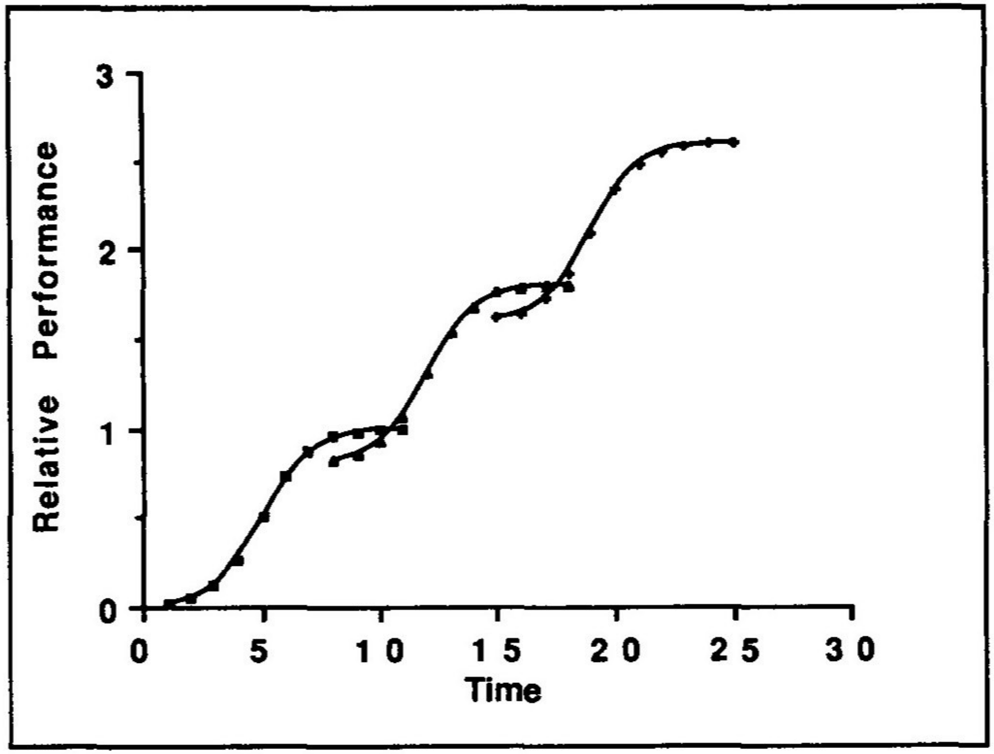

Figure 10. Piecewise-logistic patterns of change.

Examples of inflection points where technology changed

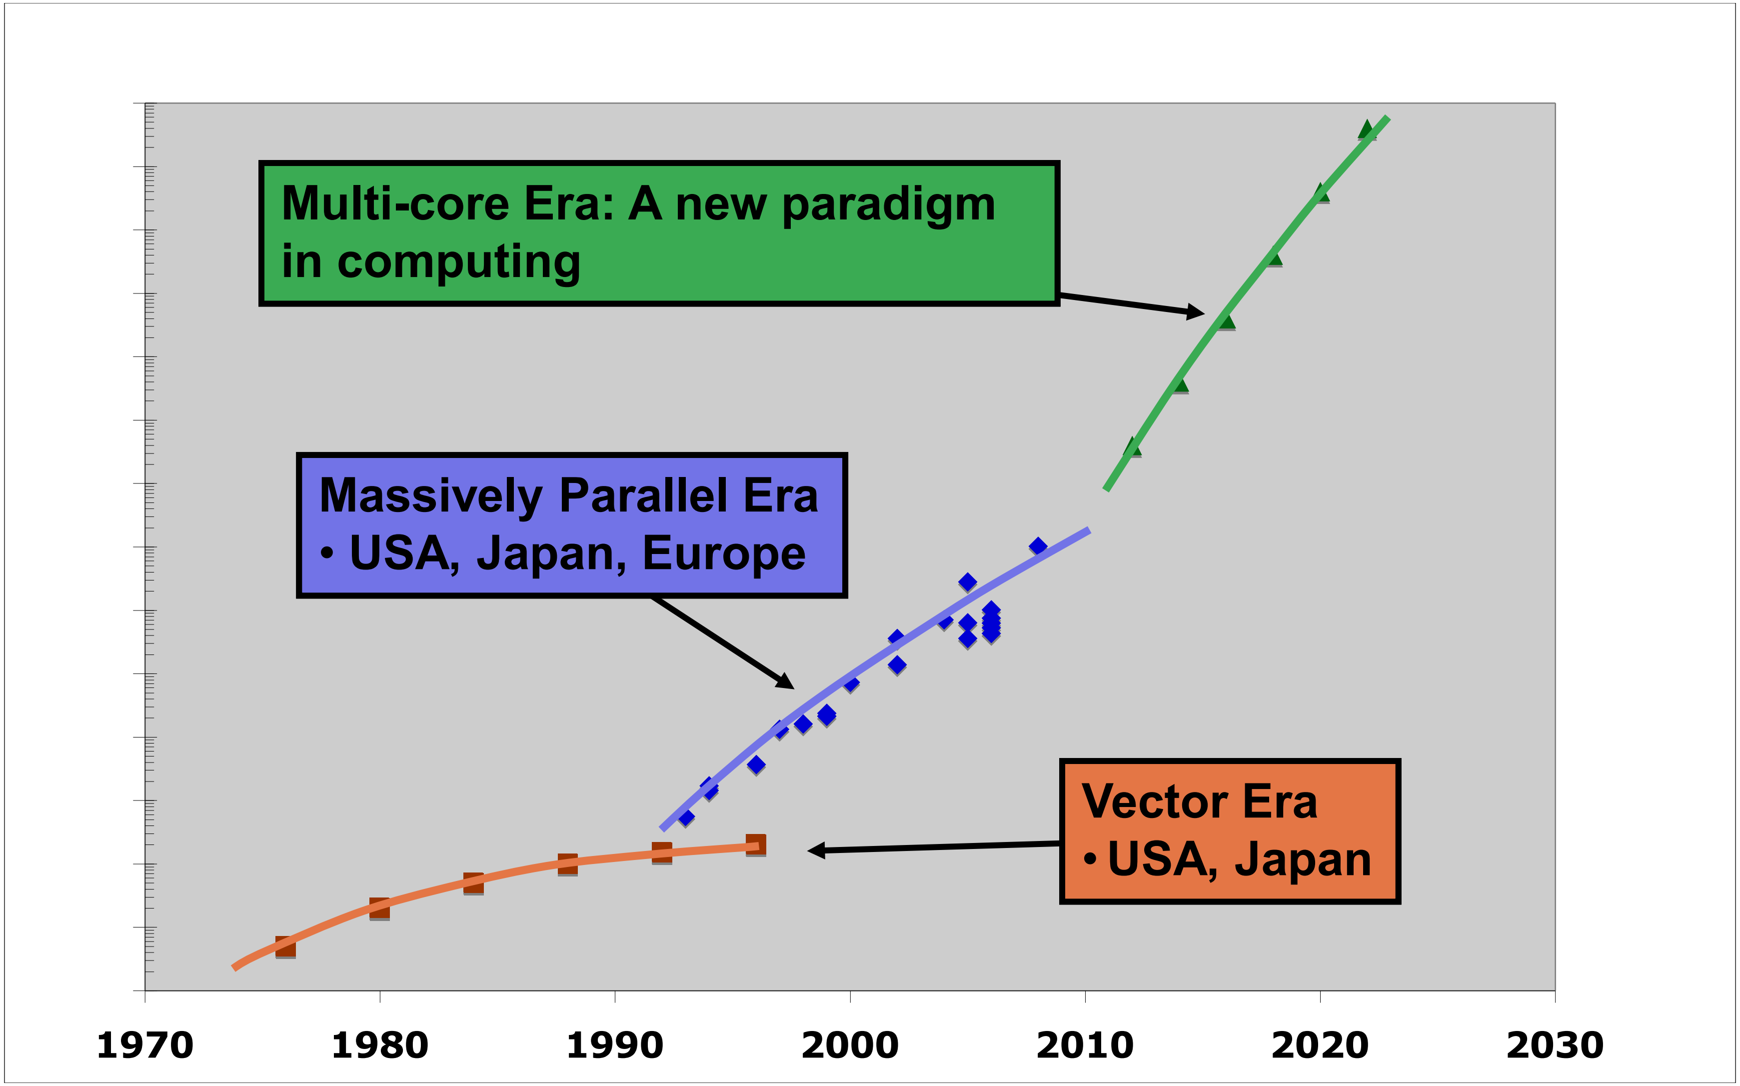

Multi-core Era: A new paradigm in computing

Massively Parallel Era
- USA, Japan, Europe

Vector Era
- USA, Japan

What is needed to move beyond Jaguar

- Weak scaling of apps has run its course
- Need more powerful nodes for strong scaling
  - Faster processors but using much less power per GFLOPS
  - More memory
  - Better interconnect
- Hierarchical programming model to expose more parallelism
  - Distributed memory among nodes 100K – 1M way parallelism
  - Threads within nodes 10s – 100s of threads per node
  - Vectors within the threads 10s – 100s of vector elements/ thread

ORNL's "Titan" System

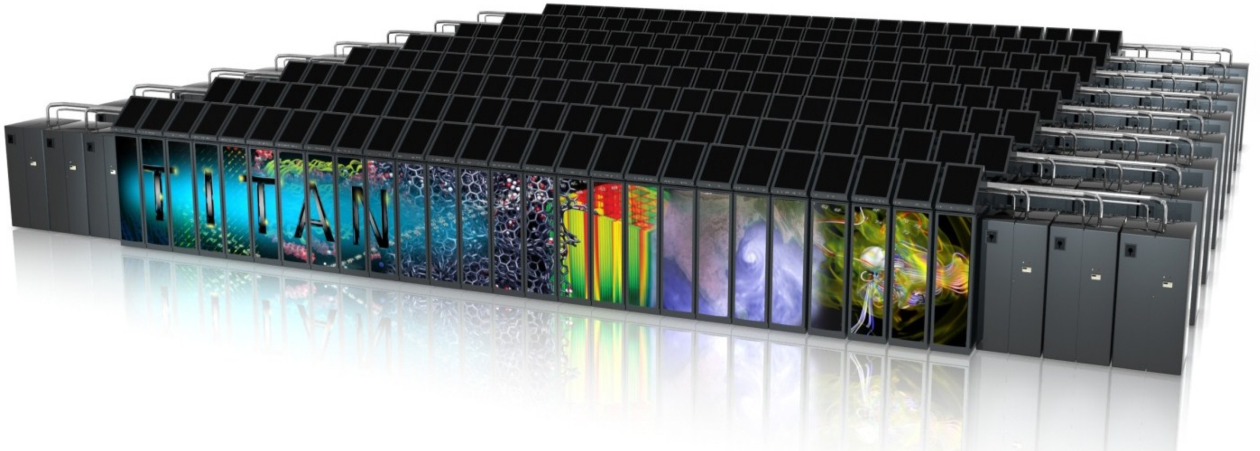

- Upgrade of Jaguar from Cray XT5 to XK6
- Cray Linux Environment operating system
- Gemini interconnect
  - 3-D Torus
  - Globally addressable memory
  - Advanced synchronization features
- AMD Opteron 6274 processors (Interlagos)
- New accelerated node design using NVIDIA multi-core accelerators
  - 2011: 960 NVIDIA x2090 "Fermi" GPUs
  - 2012: 14,592 NVIDIA K20 "Kepler" GPUs
- 20+ PFlops peak system performance
- 600 TB DDR3 mem. + 88 TB GDDR5 mem

Titan Specs	
Compute Nodes	18,688
Login & I/O Nodes	512
Memory per node	32 GB + 6 GB
# of Fermi chips (2012)	960
# of NVIDIA K20 "Kepler" processor (2013)	14,592
Total System Memory	688 TB
Total System Peak Performance	20+ Petaflops

Cray XK6 Compute Node

XK6 Compute Node Characteristics

AMD Opteron 6200 Interlagos
16 core processor

Tesla X2090 @ 665 GF

Host Memory
16 or 32GB
1600 MHz DDR3

Tesla X090 Memory
6GB GDDR5 capacity

Gemini High Speed Interconnect

Upgradeable to NVIDIA's KEPLER "K20" many-core processor

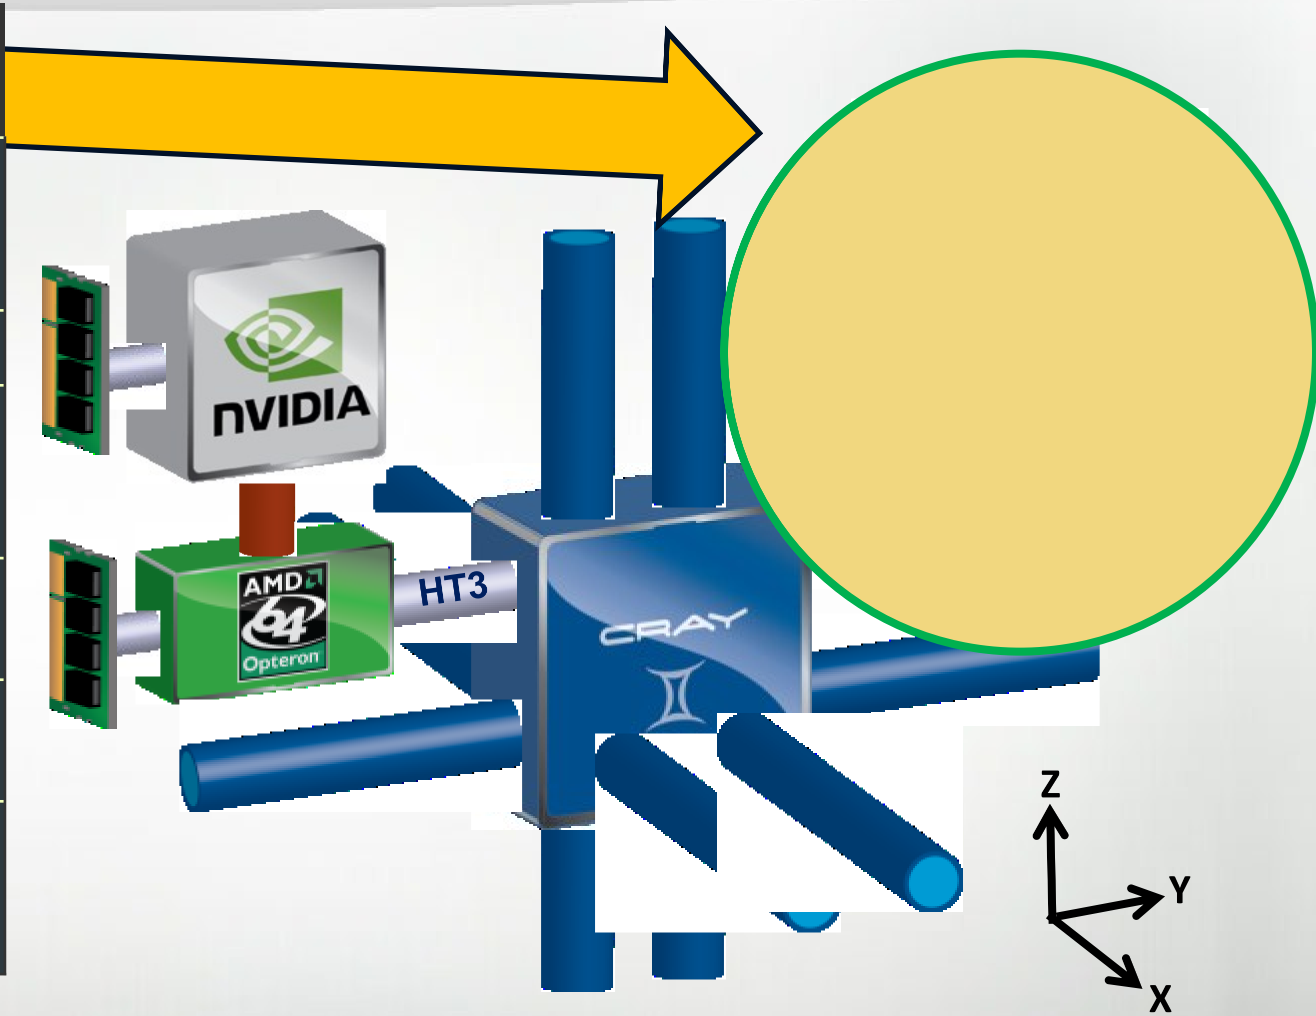

Slide courtesy of Cray, Inc.

Why use an accelerator?

- Best way to get to a very powerful node
  - Our users tell us that they want fewer, much more powerful nodes
  - Titan nodes will be greater than 1.5 TeraFLOPS per node
- Power consumption per GF is much better than a conventional processor

Processor type	GigaFLOPS / Watt
Cray XE6 (Magny-Cours)	1
Titan (Projected)	6.3

- Explicitly managed memory hierarchy
  - Programmer places the data in the appropriate memory and manages to save energy

Isn’t this a risky strategy?

- **Hardware risk is low**

Component	Risk
System Scale	**Same as Jaguar**
Processor	**Opteron follow-on**
Accelerator	**Fermi follow-on**
Interconnect	**Deployed in Cielo & Franklin at scale**
Operating System	**Incremental changes from Jaguar**

- **Applications are where there is risk**

Component	Risk
Programmability	**Rewrite in CUDA or something else?**
Tools	**Debuggers and development tools**
Programmer reluctance to rewrite for GPUs	**Fermi follow-on will provide a lot of capability**

Titan: Early Science Applications

Driven by science mission, algorithms, data structures, programming models, libraries

Facilitated by creation of the Center for Accelerated Application Readiness (CAAR)

WL-LSMS

Role of material disorder, statistics, and fluctuations in nanoscale materials and systems.

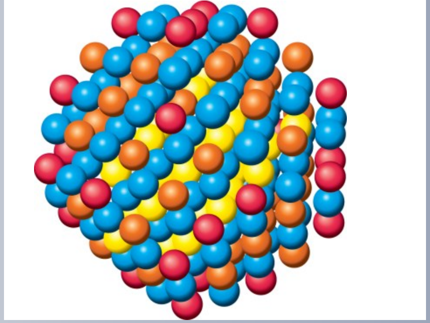

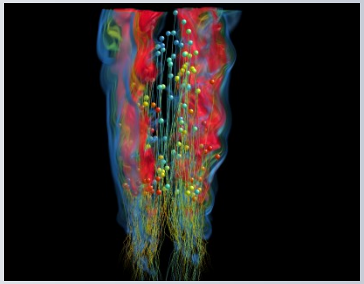

S3D

How are going to efficiently burn next generation diesel/bio fuels?

.

NRDF

Radiation transport – important in astrophysics, laser fusion, combustion, atmospheric dynamics, and medical imaging – computed on AMR grids.

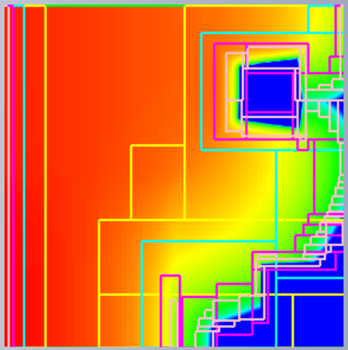

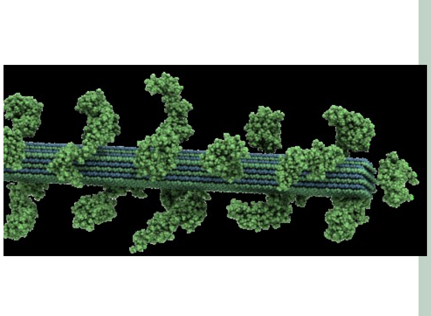

LAMMPS

A parallel particle simulator that can simulate soft materials (biomolecules, polymers), solid-state materials (metals, semiconductors) and coarse-grained or mesoscopic systems

CAM / HOMME

Answer questions about specific climate change adaptation and mitigation scenarios; realistically represent features like precipitation patterns/statistics and tropical storms

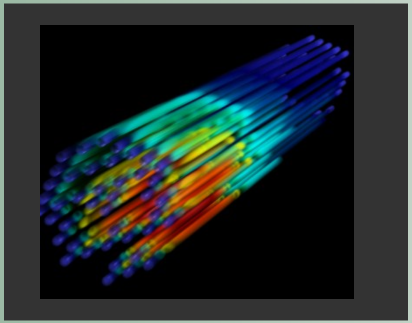

Denovo

Unprecedented high-fidelity radiation transport calculations that can be used in a variety of nuclear energy and technology applications.

Need to Exploit Hybrid Programming Model

- **On Jaguar today with 299,008 cores, we are seeing the limits of a single level of MPI scaling for most applications**

- **To take advantage of the vastly larger parallelism in Titan, users need to use hierarchical parallelism in their codes**
  - **Distributed memory: MPI, SHMEM, PGAS**
  - **Node Local: OpenMP, Pthreads, local MPI communicators**
  - **Within threads: Vector constructs on GPU, libraries, CPU SIMD**

- ***These are the same types of constructs needed on all multi-PFLOPS computers to scale to the full size of the systems!***

Hierarchical Parallelism

- MPI parallelism between nodes (or PGAS)
- On-node, SMP-like parallelism via threads (or subcommunicators, or…)
- Vector parallelism
  - SSE/AVX on CPUs
  - GPU threaded parallelism

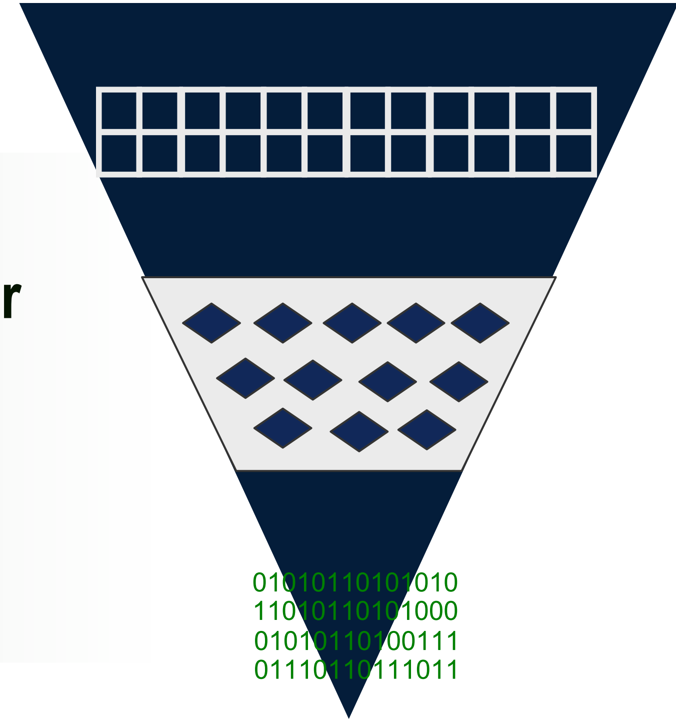

- Exposure of unrealized parallelism is essential to exploit all near-future architectures.
- Uncovering unrealized parallelism and improving data locality improves the performance of even CPU-only code.

Programming Environment

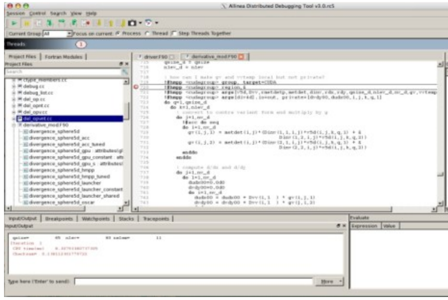

Goals:

- Full functionality hybrid programming environment
  - Compilers, Debuggers, Performance Analysis tools, Mathematical Libraries
- Hardware agnostic programming model - *portable*
  - Describe execution parallelism and data layout: expose (hierarchical) parallelism
  - Standardization through OpenMP Architecture Review Board and other industry initiatives

Exploitation of node-level parallelism

- Recognition and exploitation of hierarchical parallelism
- Development of effective programming tools to facilitate rapid refactoring of applications
- Deployment of useful performance and debugging tools to speed refactoring

Scalable Debugging for Hybrid Systems

- collaboration with Allinea to develop a scalable hybrid aware debugger based on DDT

High-productivity Hybrid-programming Through Compiler Innovation

- collaboration with HMPP to develop directive based compiler technology in CAPS compiler
  - CAPS support for OpenACC set of directives; support for all common languages used at the OLCF, …

Scalable Performance Analysis for Hybrid Systems

- collaboration with Technische Universitat Dresden to add support for Hybrid (CPU/GPU) performance analysis in Vampir

Titan Training Program

- **Goal: Enable break-through science through education**
- **Strategy: Provide conferences, workshops, tutorials, case studies, and lessons learned on tools and techniques for realizing hybrid architecture benefits. Provide content via traditional venues, online, and pre-recorded sessions.**
- **Objective: Users will be able to expose hierarchical parallelism, use compiler directive-based tools, analyze / optimize / debug codes, and use low-level programming techniques if required**

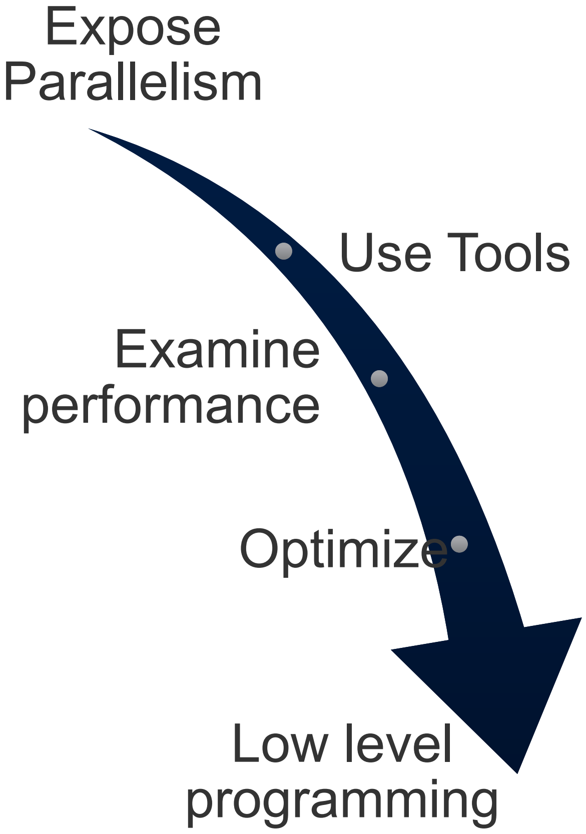

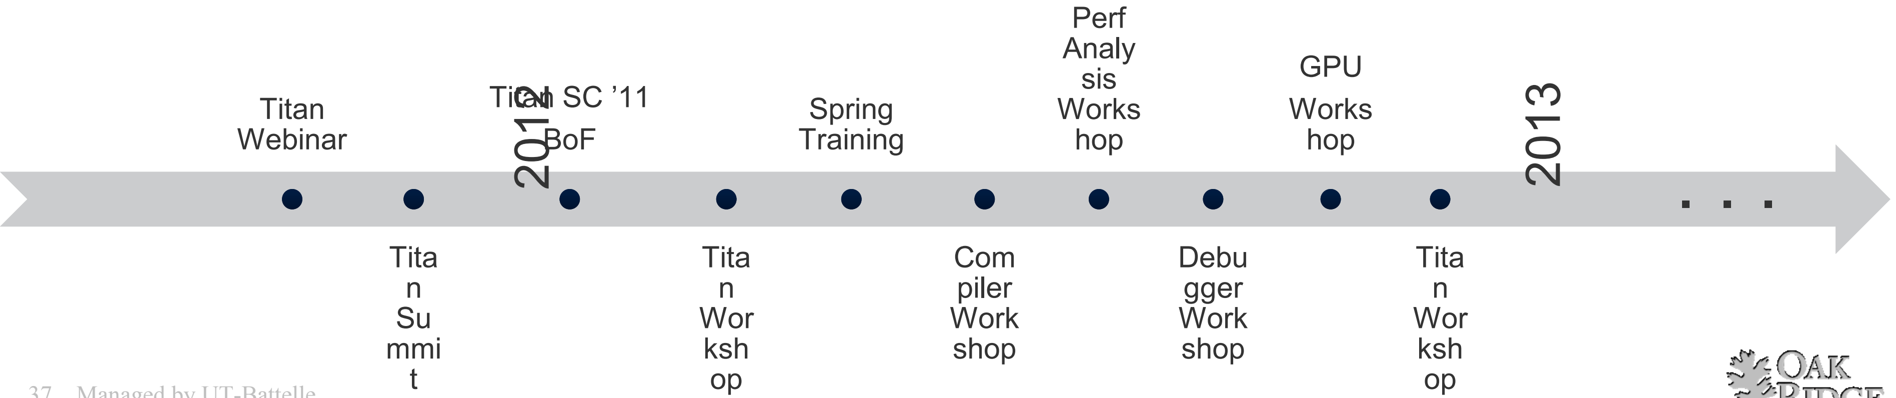

How Effective are GPUs on Scalable Applications?

March 2012 Performance Snapshot on TitanDev (Fermi accelerator)

Application	XK6 (w/ GPU) vs. XK6 (w/o GPU) Performance Ratio Titan Dev : Jaguar	XK6 (w/ GPU) vs. XE6 Performance Ratio Titan Dev : Monte Rosa
S3D Turbulent combustion	1.5	1.4
Denovo 3D neutron transport for nuclear reactors	3.5	3.3
LAMMPS High-performance molecular dynamics	6.5	3.2
WL-LSMS Statistical mechanics of magnetic materials	3.1	1.6
CAM-SE Community atmosphere model	2.6	1.5
NAMD High-performance molecular dynamics	2.6	1.4
Chroma High-energy nuclear physics	8.8	6.1
QMCPACK Electronic structure of materials	3.8	3.0
SPECFEM-3D Seismology	4.7	2.5
GTC Plasma physics for fusion-energy	2.5	1.6
CP2K Chemical physics	2.8	1.5

Cray XK6: Fermi GPU plus Interlagos CPU
Cray XE6: Dual Interlagos and no GPU

■ Gordon Bell Winner

Why did we pick CAM-SE?

CAM4 1/8° (14km) Scalability

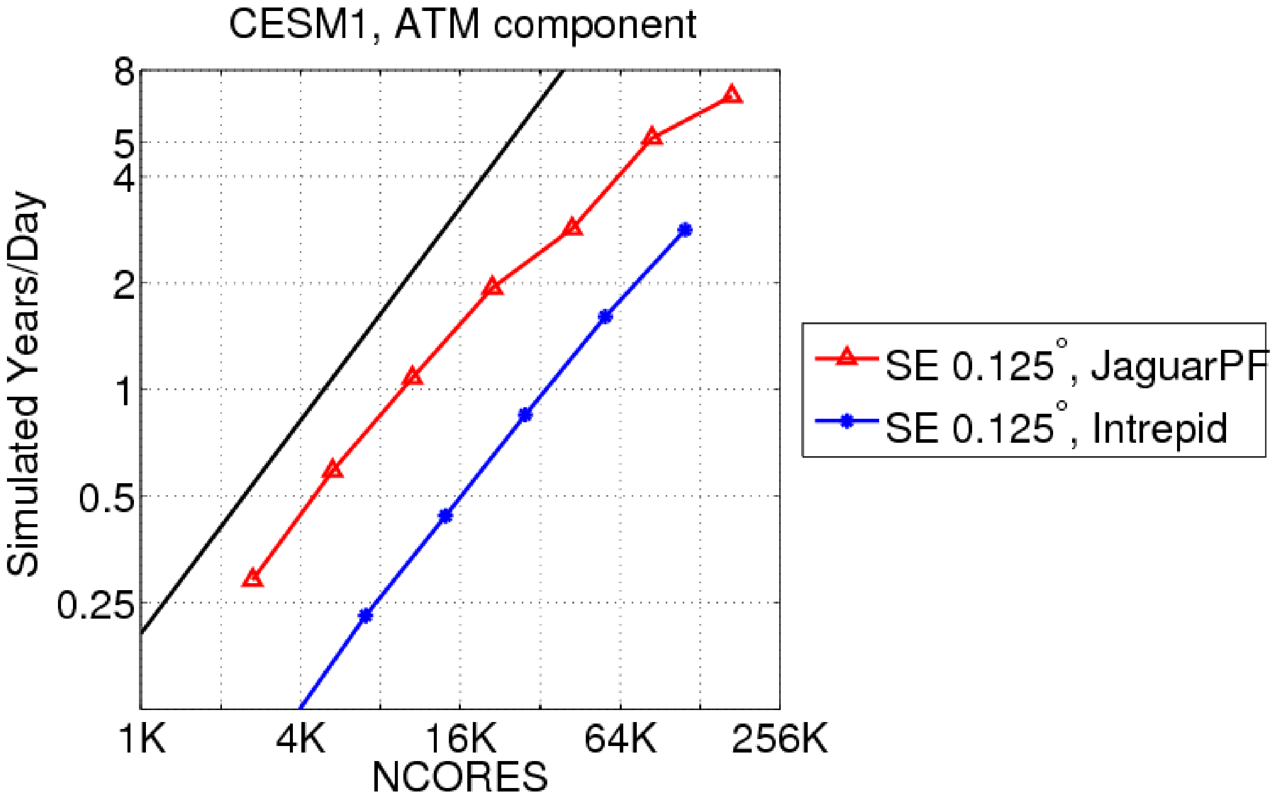

- Excellent scaling to near full machine on both LCFs:
- Intrepid (4 cores/node): Excellent scalability, peak performance at 115K cores, 3 elements per core, 2.8 SYPD.
- JaguarPF (12 cores/node): Good scalability, peak performance at 172,800 cores (2 elements per core), 6.8 SYPD.

Courtesy M. Taylor, Sandia National Laboratory

CAM-SE GPU Port

- 7 boundary exchanges per tracer advection step
- Multiple variables grouped in each exchange
- Current method
  - Buffer data for all boundary exchanges (pack)
  - Unpacked data updates memory or generates MPI calls(`bndry_exchangeV`) using communication schedule data structures
    - If generating MPI call, data is MPI buffered and sent to neighbor
    - If edge involves elements on the same MPI task, update in memory
  - Loop over MPI receives and put new data back in buffer
  - Unpack buffer with new values
- Currently optimized for 1 element / task (all MPI calls)
- Must optimize for larger volume-to-surface ratio
- Cubed-sphere decomposed by space filling curves possibly giving irregular domains

CAM-SE On Titandev (With –O2 CPU Flag)

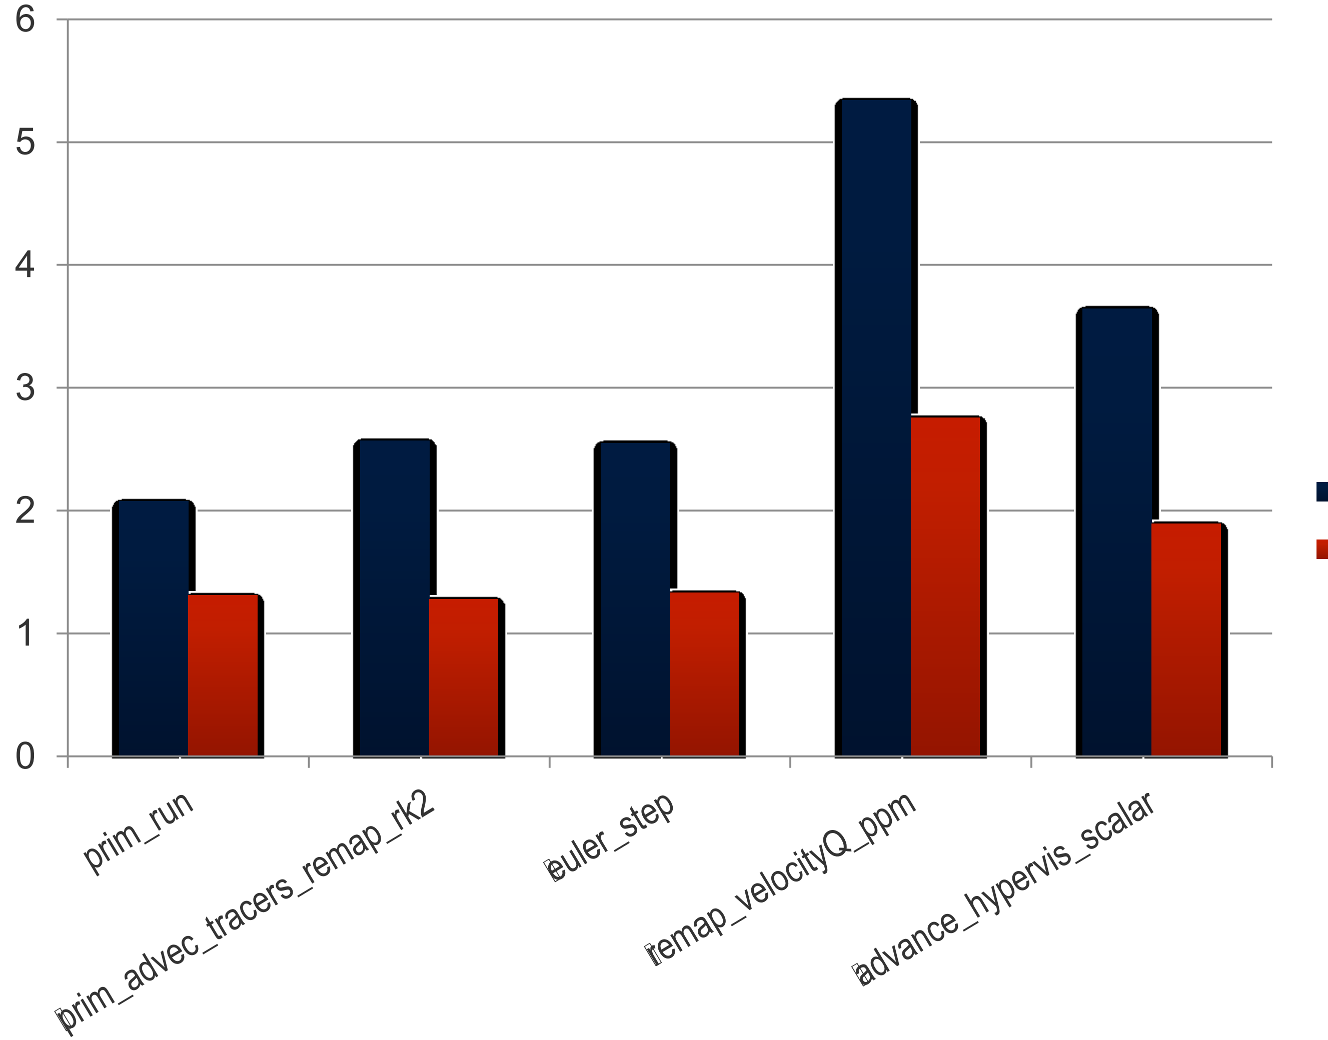

Speed-Up: Fermi GPU vs 1 Interlagos / Node

- Benchmarks performed on XK6 using end-to-end wall timers
- All PCI-e and MPI communication included

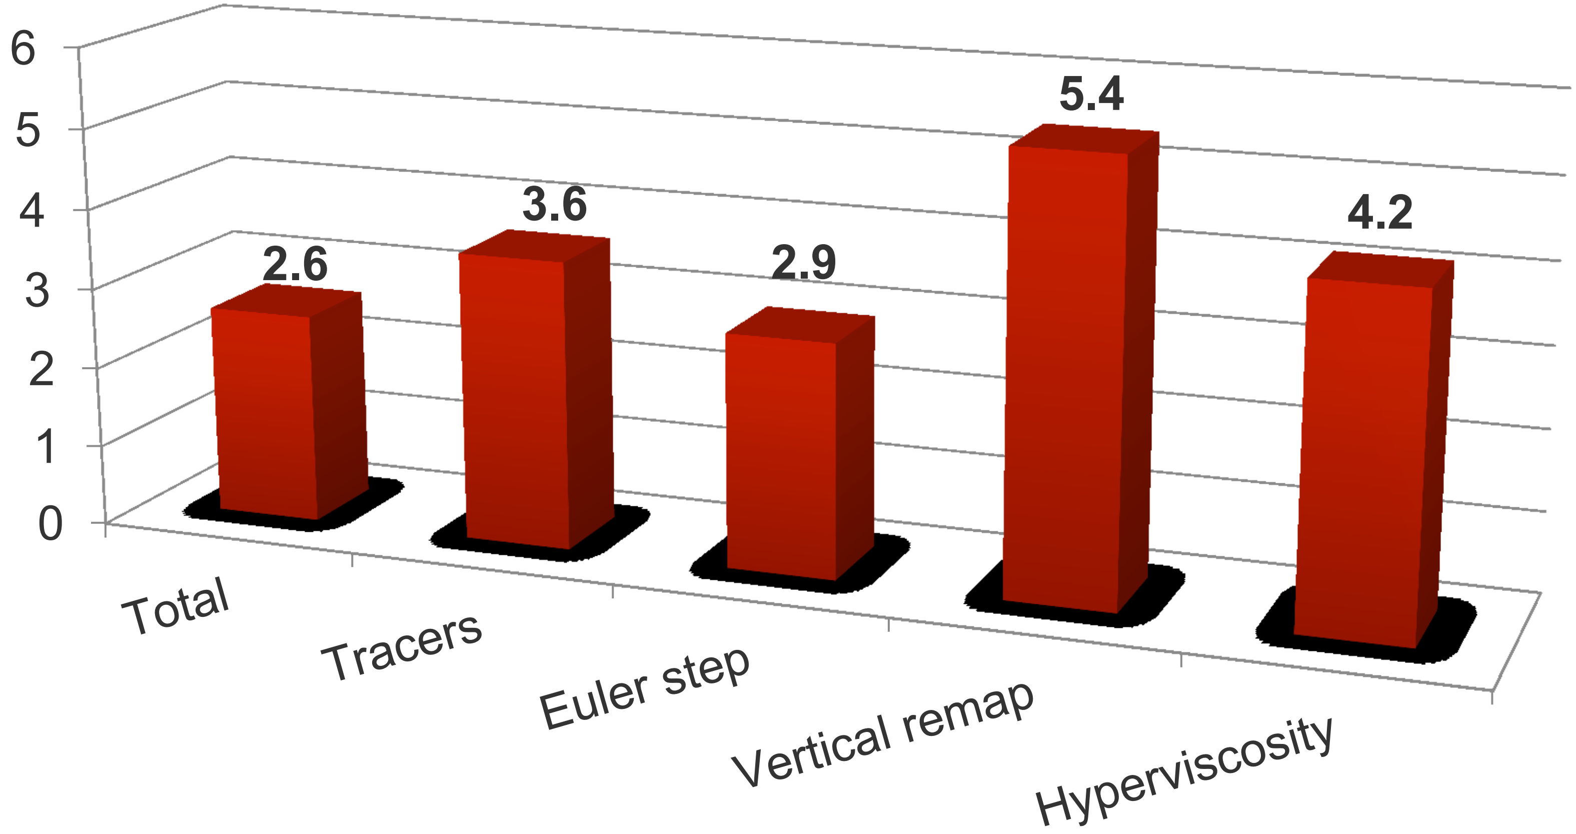

NVIDIA Kepler Processor Cards are arriving

Some Lessons Learned

- Exposure of unrealized parallelism
  - Figuring out where is often straightforward
  - Making changes to exploit it is hard work (made easier by better tools)
  - Developers can quickly learn, e.g., CUDA and put it to effective use
  - A directives-based approach offers a straightforward path to portable performance
- For those codes that already make effective use of scientific libraries, the possibility of continued use is important.
  - HW-aware choices
  - Help (or, at least, no hindrance) to overlapping computation with device communication
- Ensuring that changes are communicated back and remain in the production "trunk" is every bit as important as we initially thought.
  - Other development work taking place on all CAAR codes could quickly make acceleration changes obsolete/broken otherwise

Maintaining leadership: data infrastructure

- The OLCF Spider project was groundbreaking
  - Largest scale and highest performance parallel file system ever deployed
  - Success, which leveraged Lustre experience at LLNL, has served as a blueprint for HPC community
- Collaboration was key to success
  - Leveraged R&D across ORNL to architect, prototype, develop and deploy
  - Partnerships with Cray, DDN, and the Lustre development team were critical
- Titan environment demands new capabilities
  - Scalable metadata performance
  - Bandwidth scalability to over a Terabyte/sec
  - Scalable system resiliency for improved fault tolerance
- Challenges would not be met without our efforts
  - Leadership role in Lustre – Open Scalable File Systems Consortium
  - Continued partnerships both internally and externally to ensure solutions that meet our requirements are available and affordable

Data and experimental workflow

- Simulation and sampling plan
    - Need to archive and export >300 TB to partners
- Development of automated workflow for high resolution production simulations
    - Manual system not scalable
    - Initial configuration for Jaguar, but can be exported hardened and exported

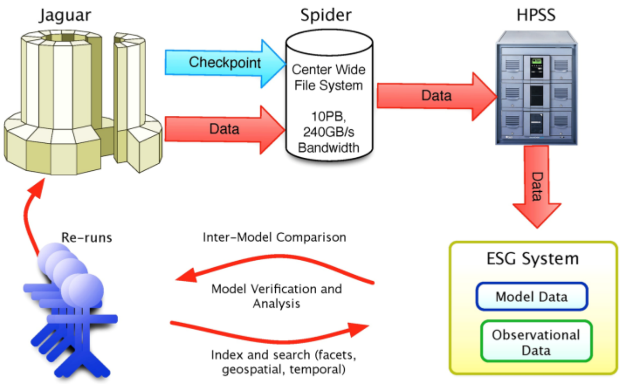

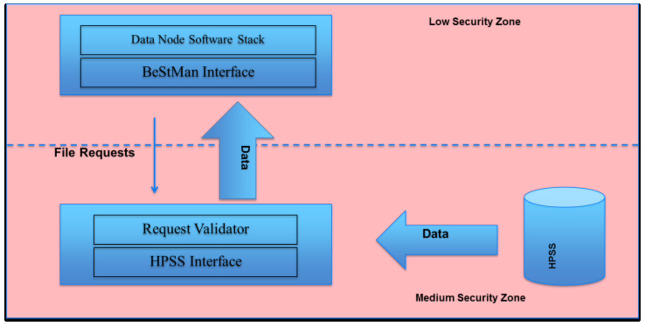

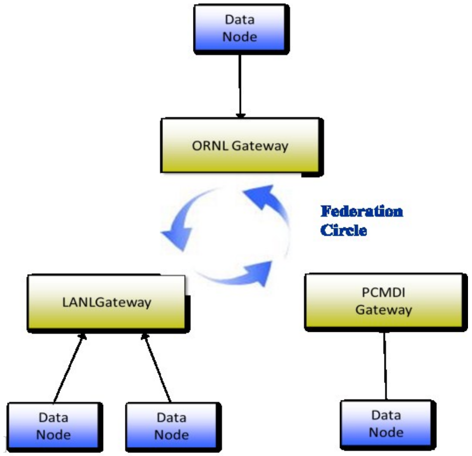

- High Volume Climate Data Server
    - http://cds.ccs.ornl.gov
    - Integrated with Earth System Grid (ESG)

Shipman and collaborators leveraging ESG effort See Flanery at al. poster

HPSS – Managing Exponential Growth in Storage

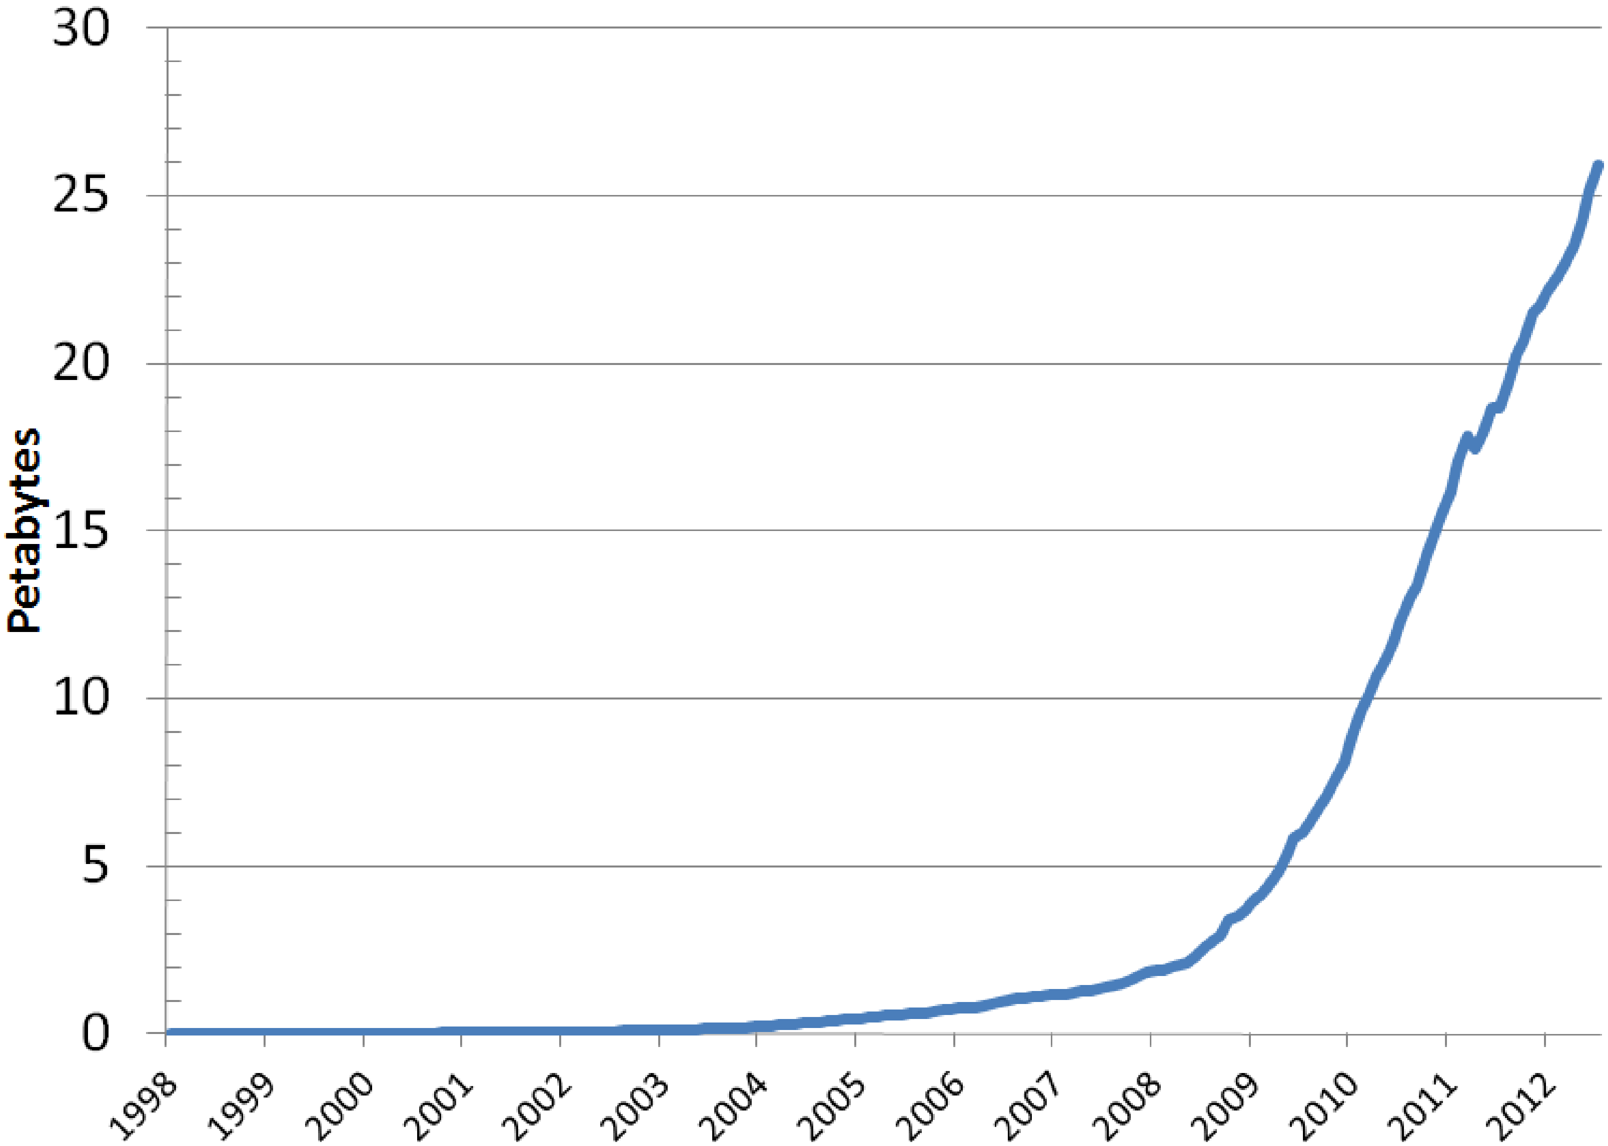

HPSS Growth:

PB	Months
1	101.5
2	20.3
3	6.2
4	4.3
5	3.1
6	2.6
7	2.8
8	2.1
9	1.4
10	1.4
11	2.1
12	1.8
13	1.5
14	1.7
15	1.5
16	1.5
17	1.4
18	3.2
19	2.4
20	1.3
21	1.7
22	2.2
23	2.5
24	1.8
25	1.1
26	1.3

HPSS is managing 26+ PB and growing at more than 30 TB per day.

The rate of increase in stored data is less due to aggressive steps to limit growth

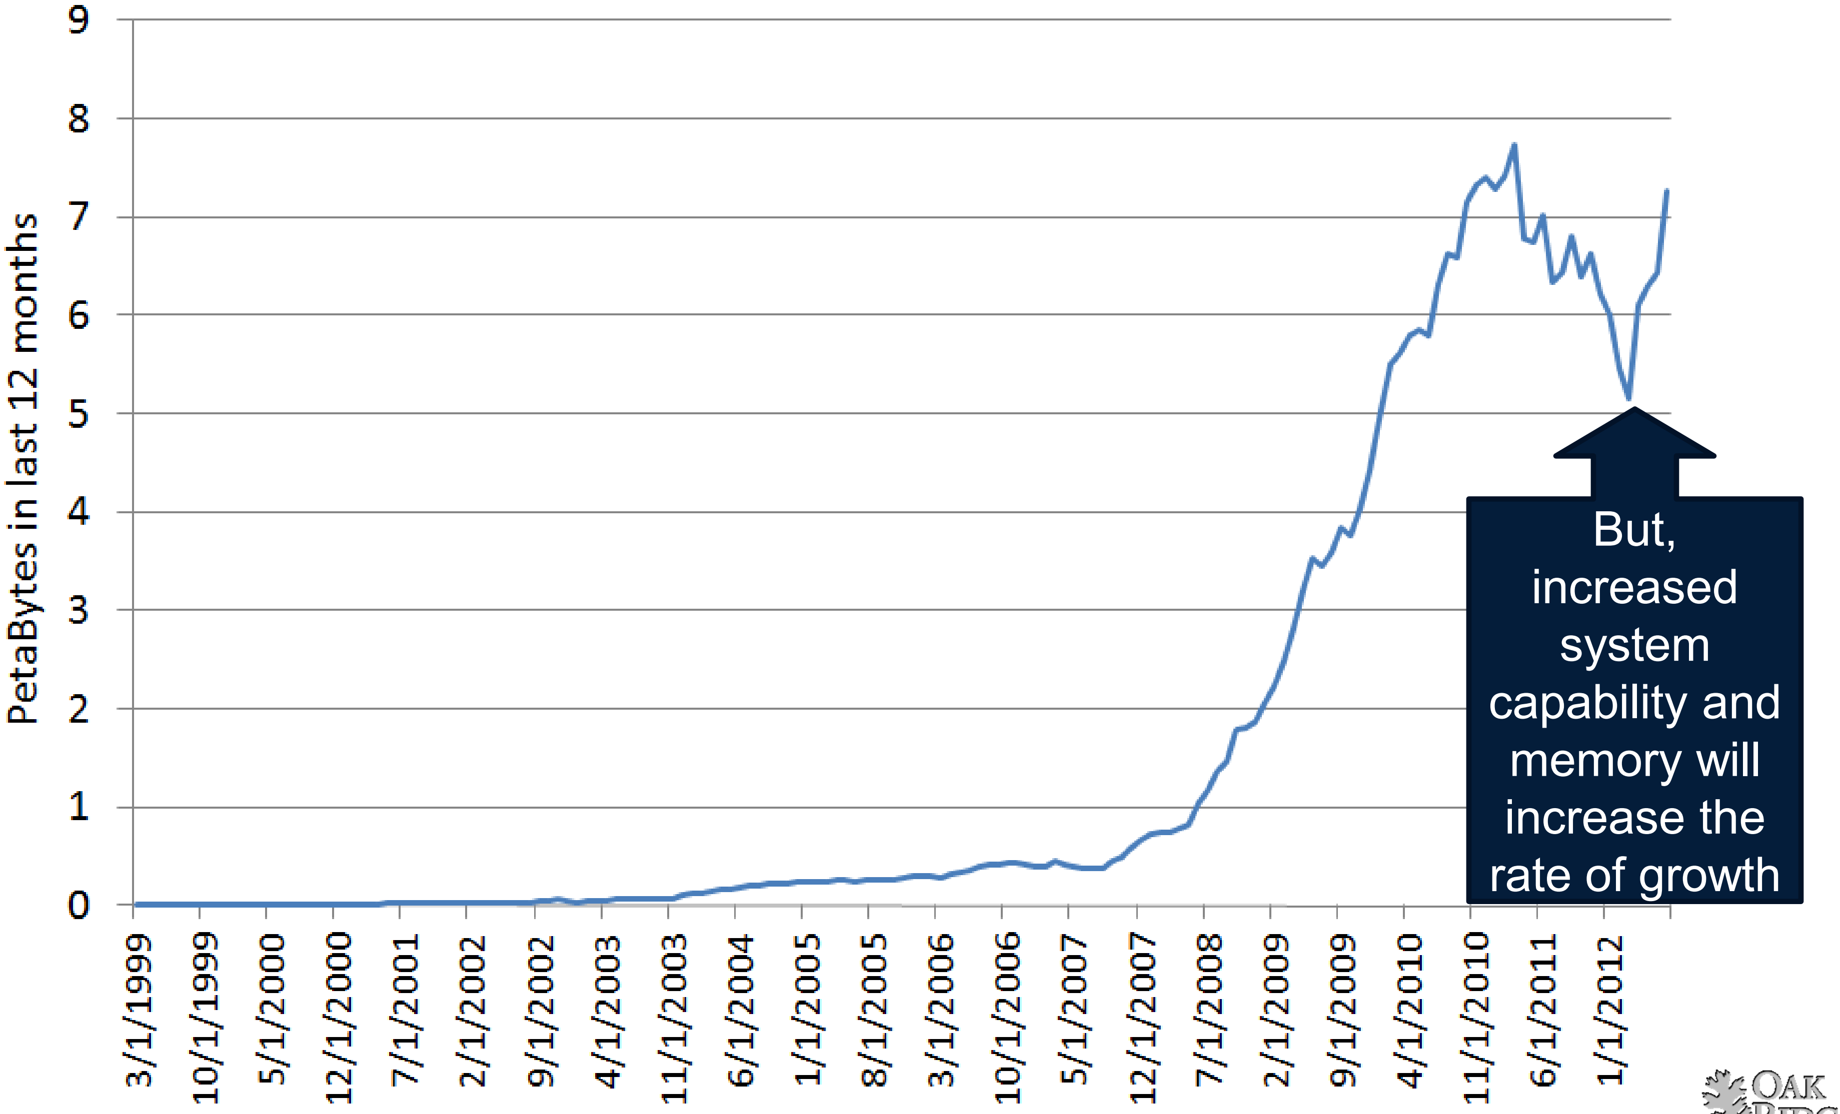

But, increased system capability and memory will increase the rate of growth

Summary

- Partnering has demonstrated value in navigating architectural transition
  - highly integrated engagement with user community has led to early success
  - CAAR application effort already demonstrating advantages of hybrid architectures
  - user assistance and outreach will help codify best practices and inform the broader community via educational opportunities

- Important investments in and collaborations with technology providers
  - Scalable Debugging for Hybrid Systems
    - collaboration with Allinea to develop a scalable hybrid aware debugger based on DDT
  - High-productivity Hybrid-programming Through Compiler Innovation
    - collaboration with HMPP to develop directive based compiler technology in CAPS compiler
      - CAPS support for OpenACC set of directives; support for all common languages used at the OLCF
  - Scalable Performance Analysis for Hybrid Systems
    - collaboration with Technische Universitat Dresden to add support for Hybrid (CPU/GPU) performance analysis in Vampir

Questions?

http://www.nccs.gov
http://www.nics.tennessee.edu/