HPC @ ORNL
Where do we go from here?

Presented to:
SC'10
November 16, 2010

Presented by:
Buddy Bland
Project Director
Oak Ridge Leadership Computing Facility

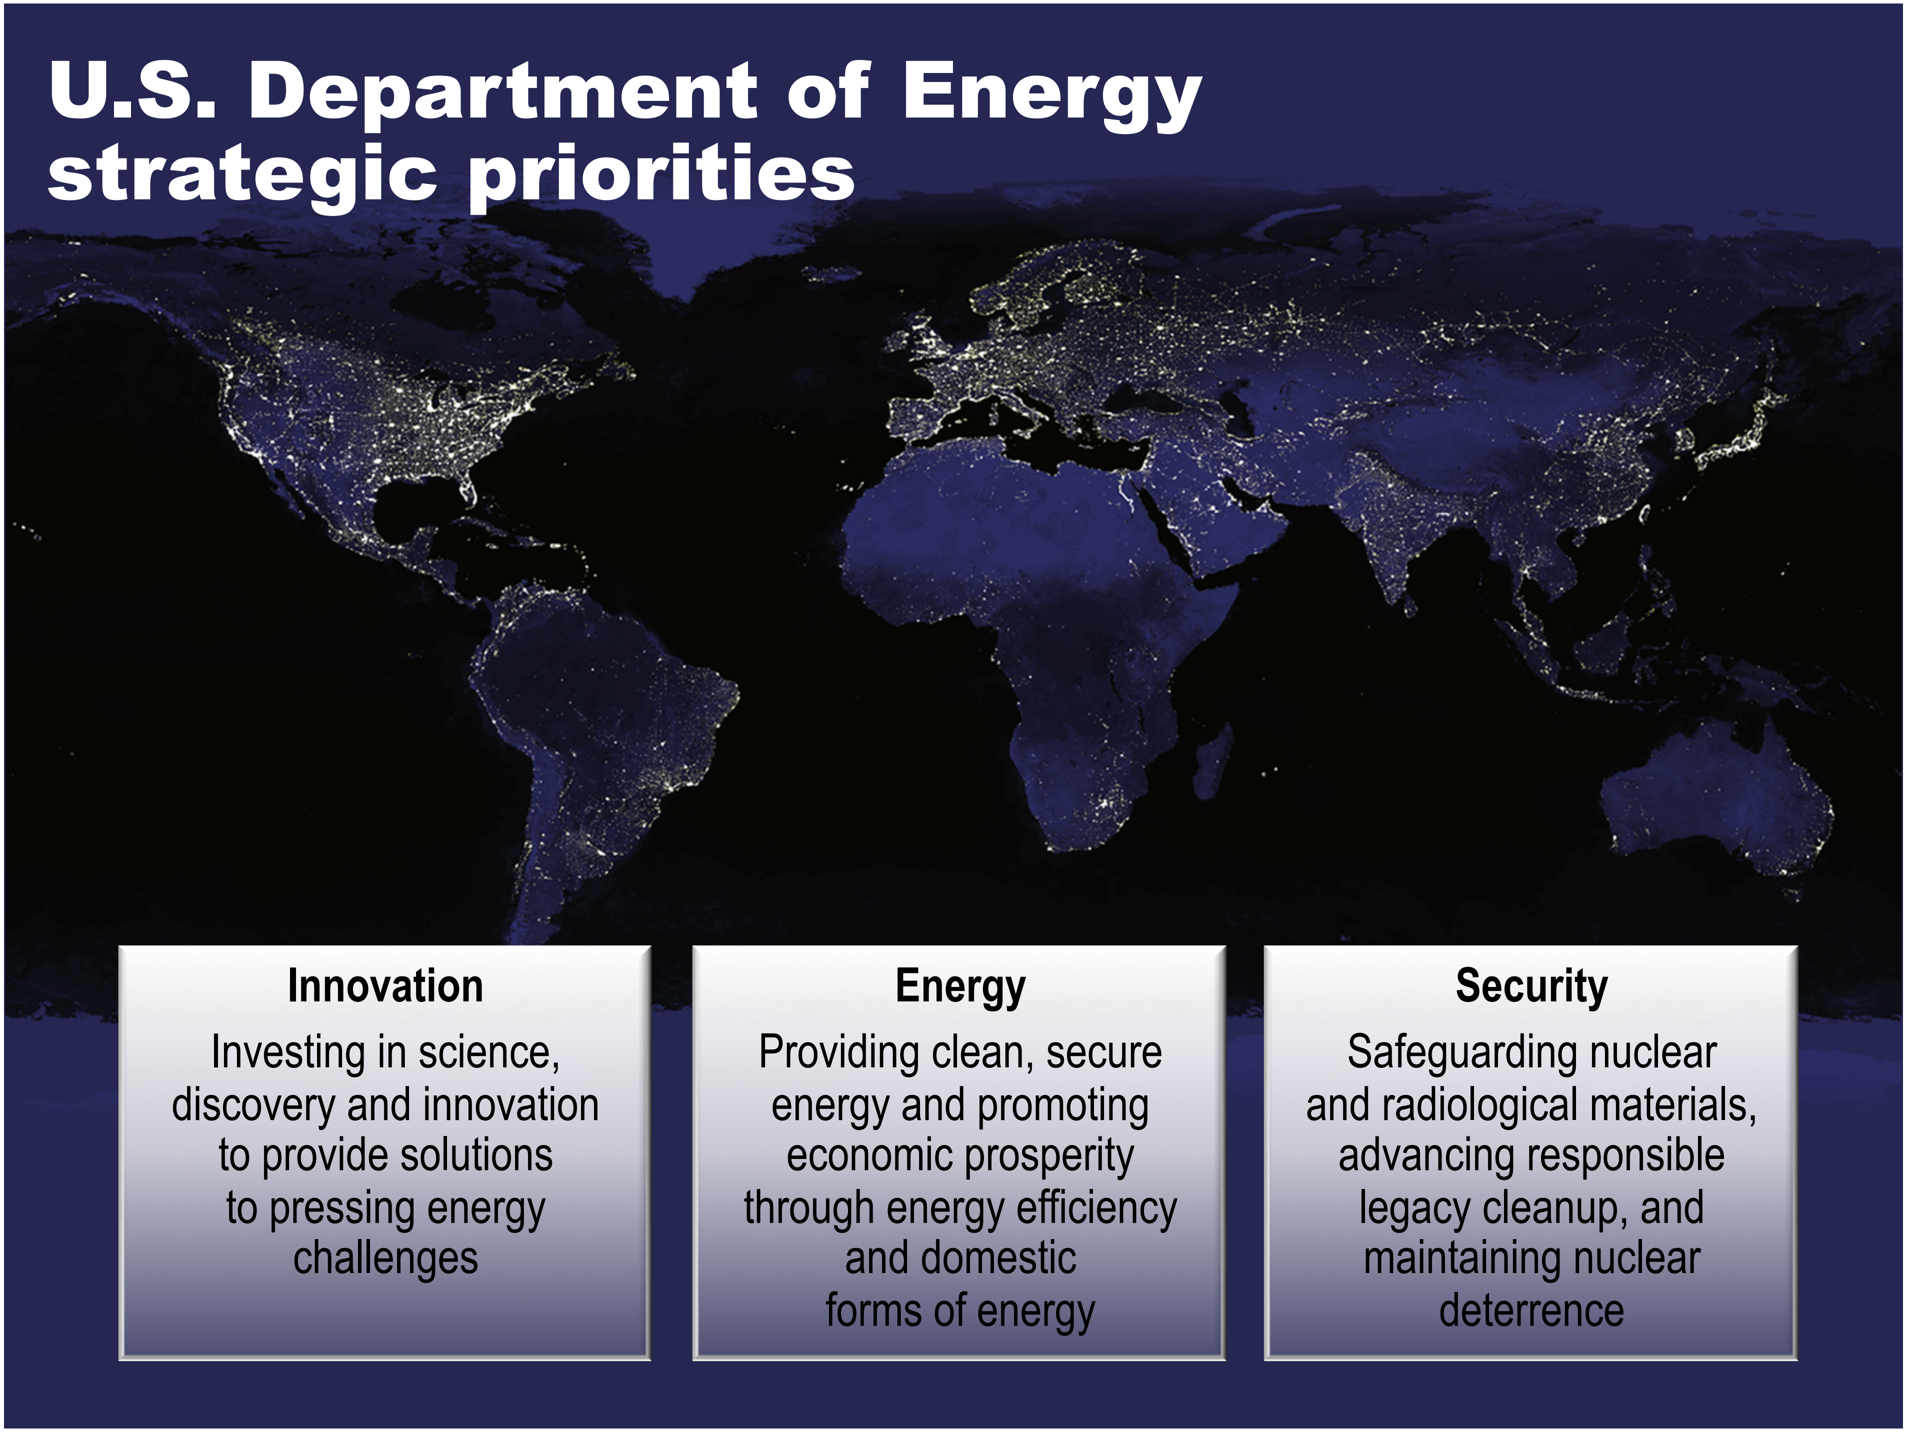
U.S. Department of Energy strategic priorities
Innovation
Investing in science, discovery and innovation to provide solutions to pressing energy challenges
Energy
Providing clean, secure energy and promoting economic prosperity through energy efficiency and domestic forms of energy
Security
Safeguarding nuclear and radiological materials, advancing responsible legacy cleanup, and maintaining nuclear deterrence

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

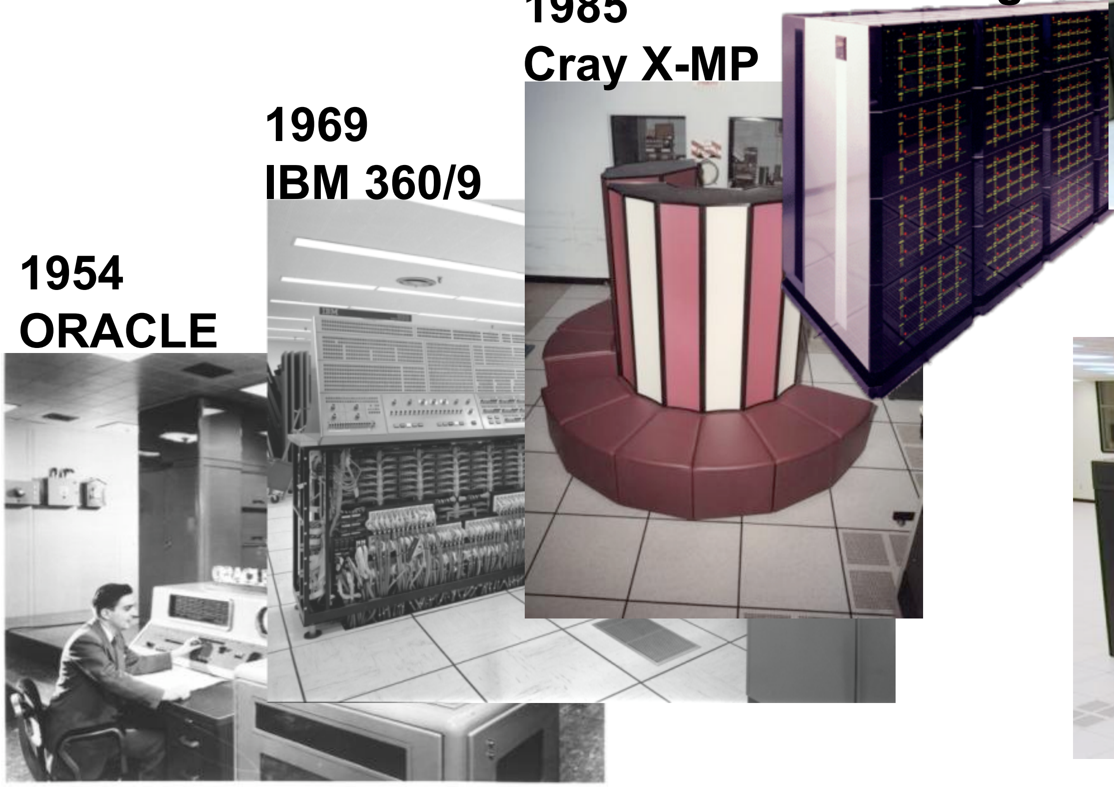

1996-2002 IBM Power 2/3/4

2007 IBM Blue Gene/P

2003-2005 Cray X1/X1E

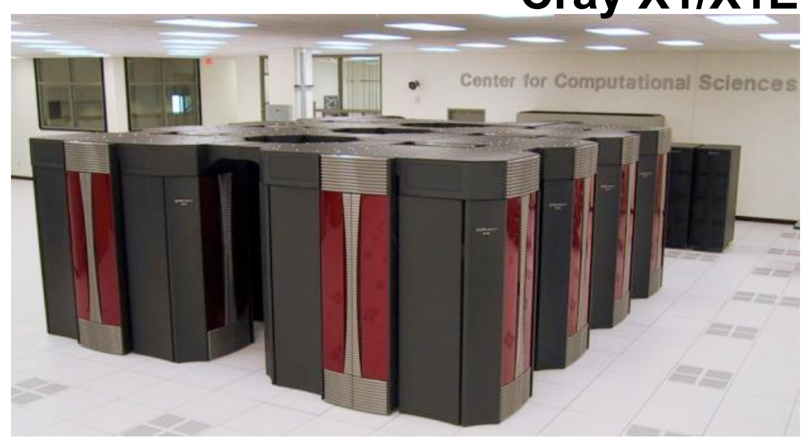

Today, we have the world's most powerful computing facility

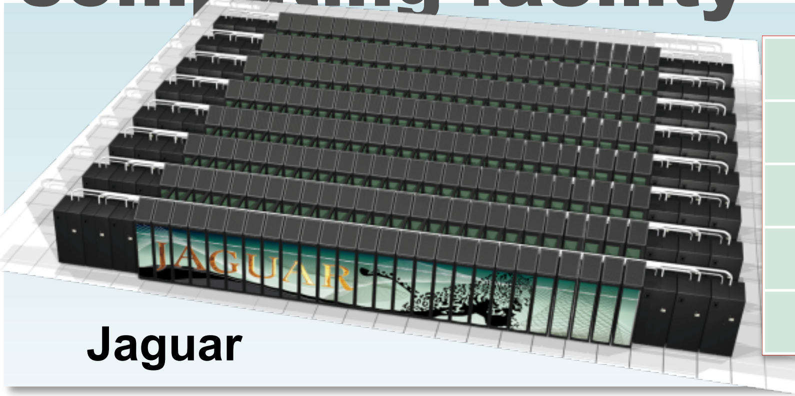

Jaguar

Peak performance	2.33 PF/s
Memory	300 TB
Disk bandwidth	> 240 GB/s
Square feet	5,000
Power	7 MW

TOP500 SUPERCOMPUTER SITES #2

Dept. of Energy's most powerful computer

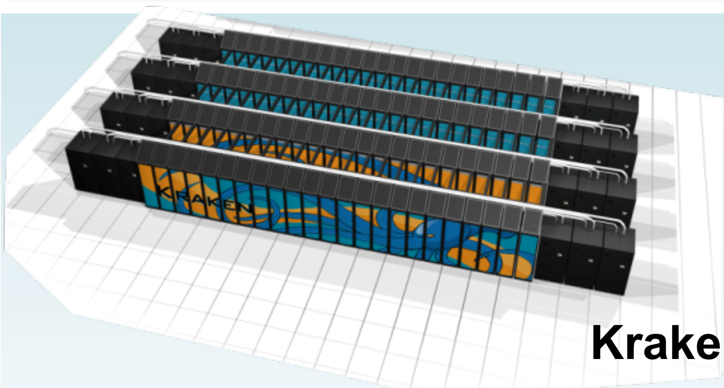

Kraken

Peak performance	1.03 PF/s
Memory	132 TB
Disk bandwidth	> 50 GB/s
Square feet	2,300
Power	3 MW

TOP500 SUPERCOMPUTER SITES #8

National Science Foundation's most powerful computer

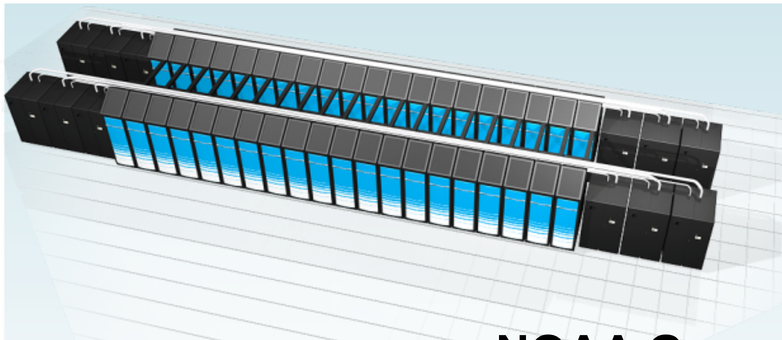

NOAA Gaea

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 #32

National Oceanic and Atmospheric Administration's most powerful computer

These systems are part of a comprehensive Simulation Environment to meet our science missions

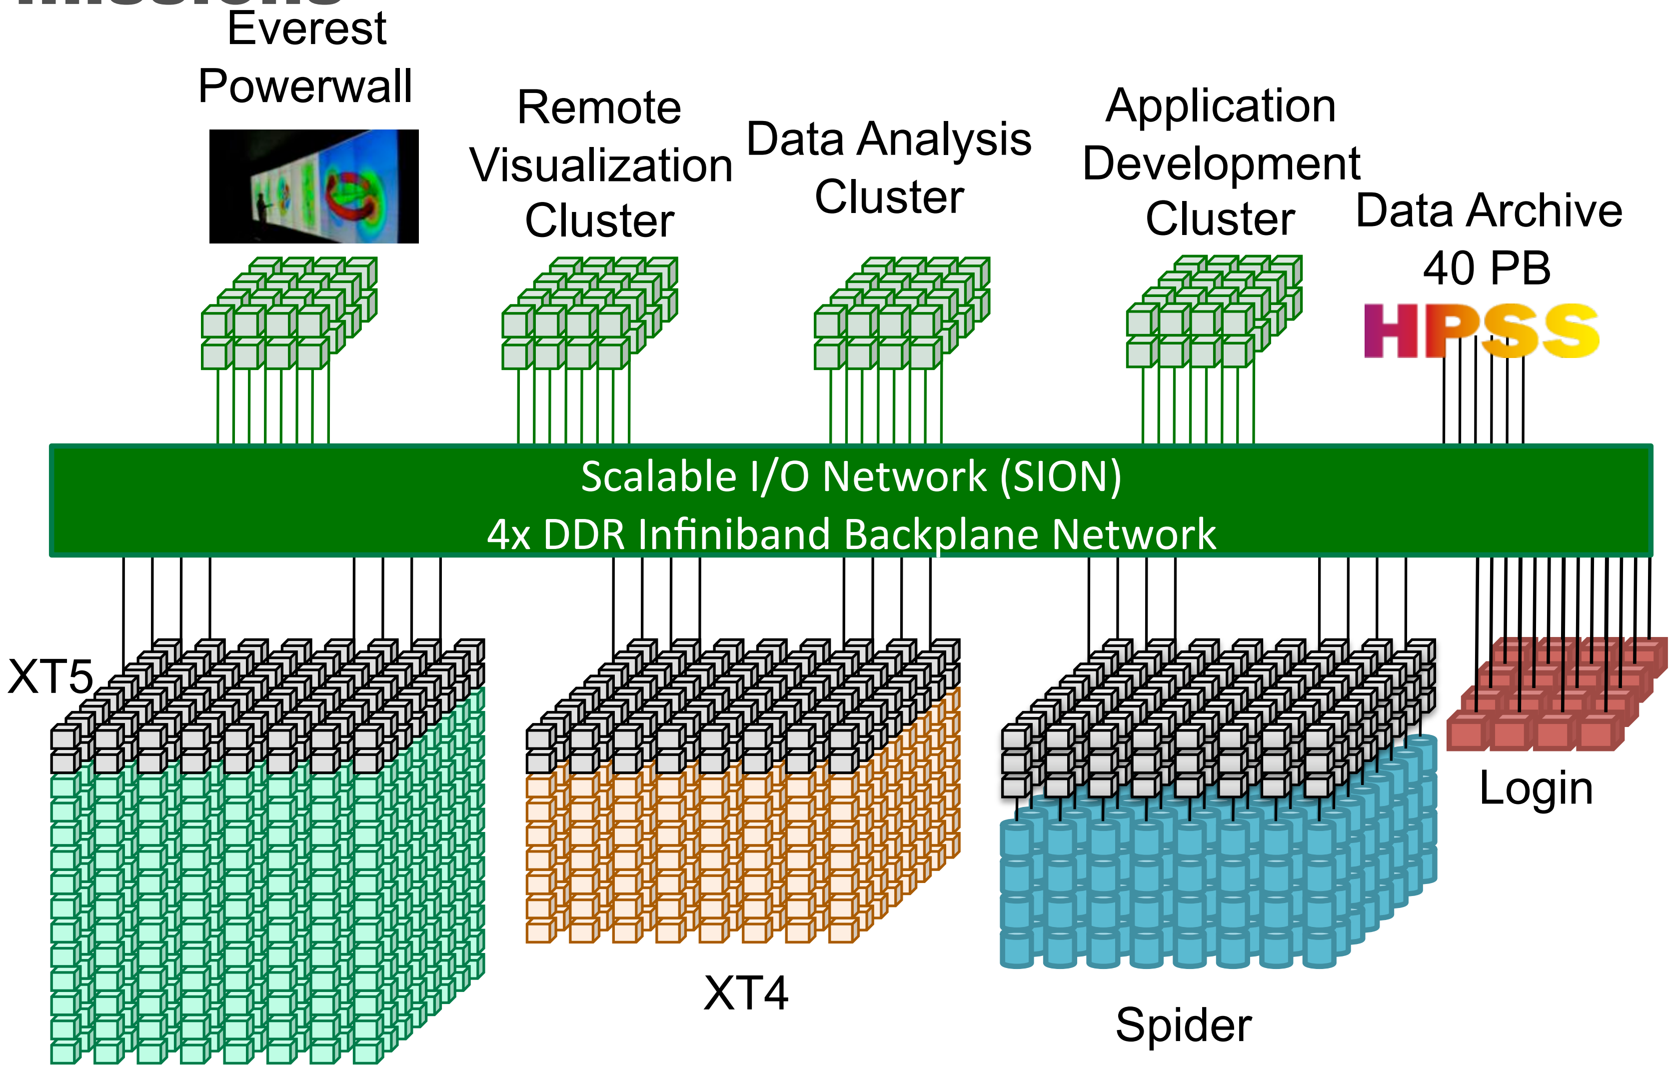

We have increased system performance by 1,000 times since 2004

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

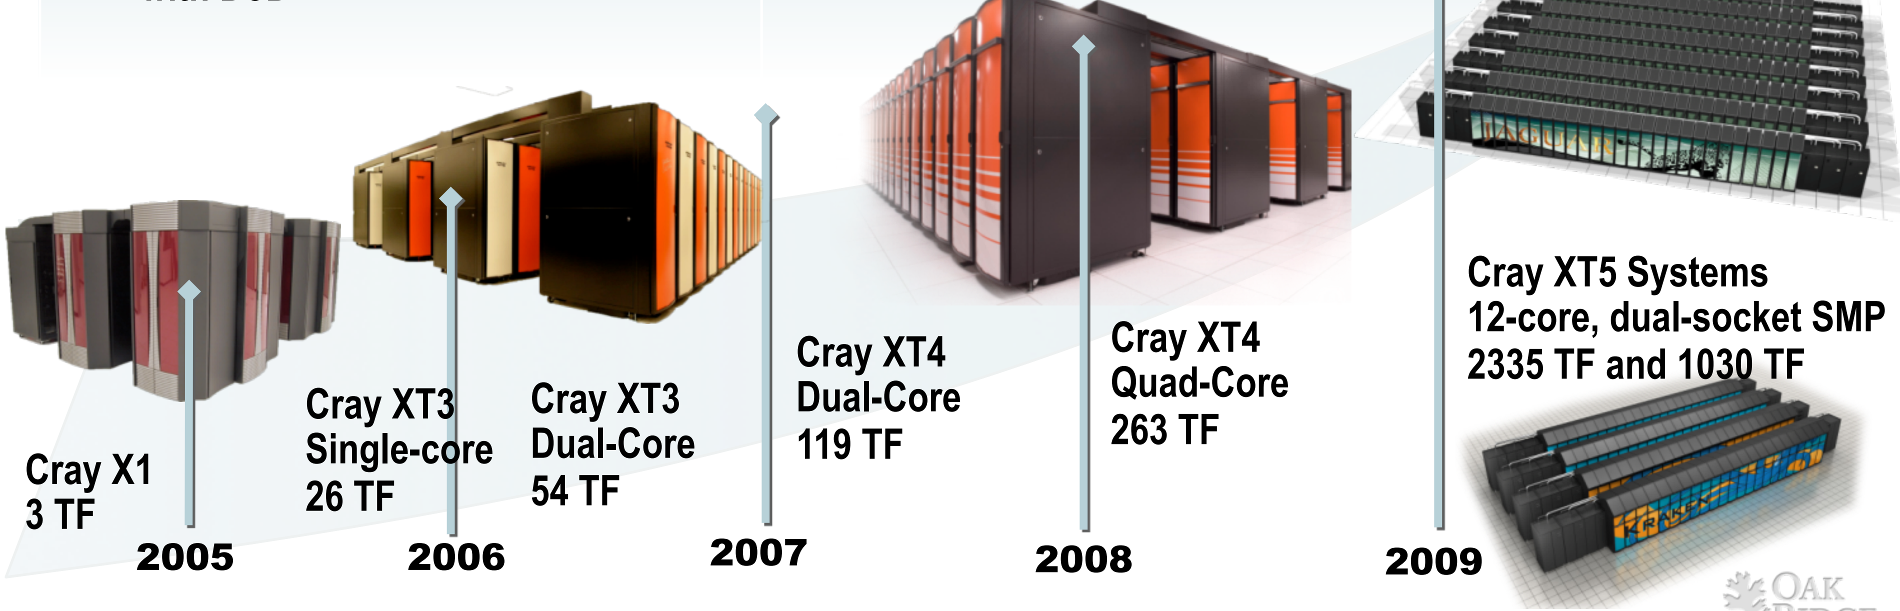

Our science requires that we advance computational capability 1000x over the next decade

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

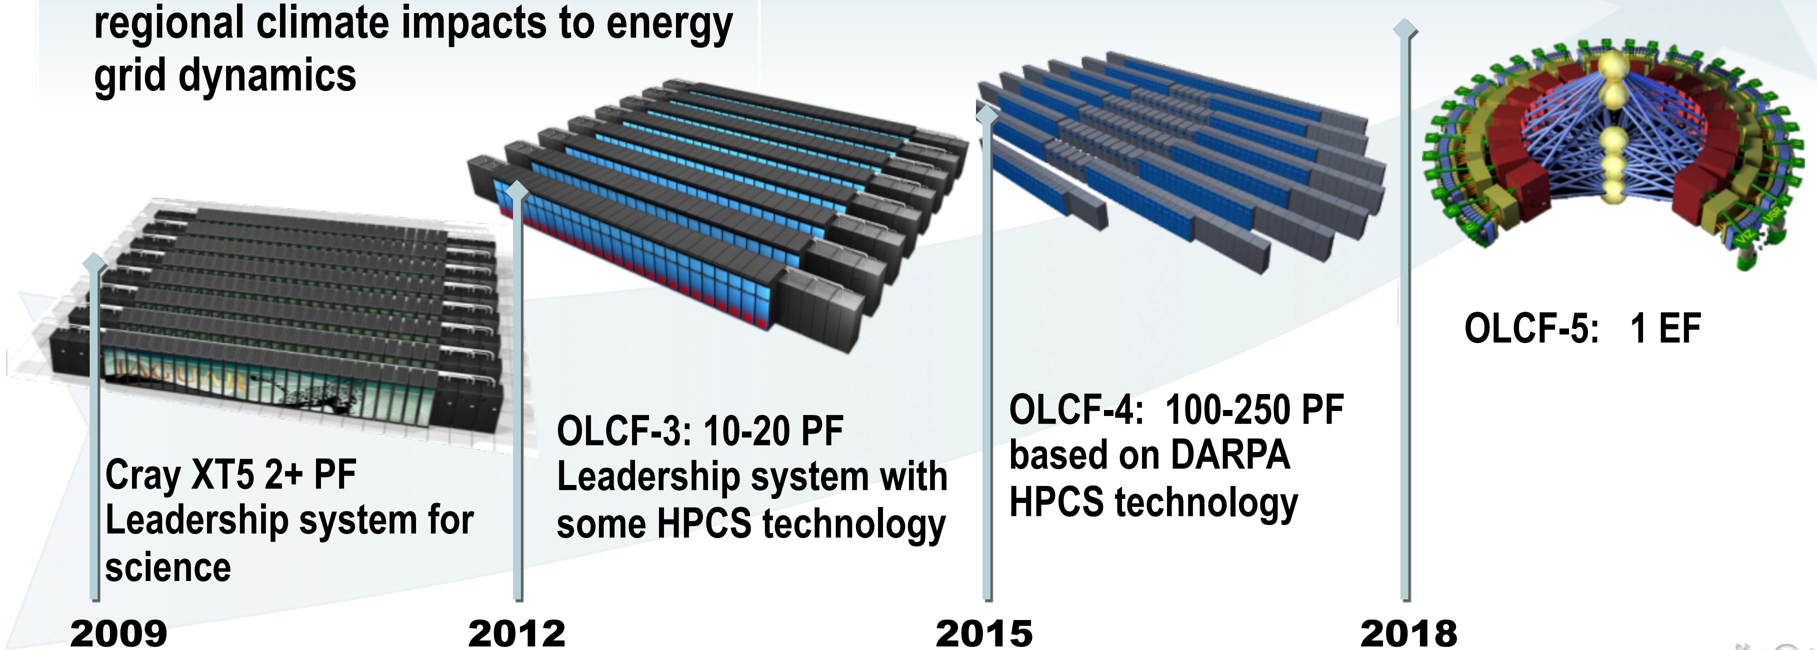

We are delivering Petascale Science Today!

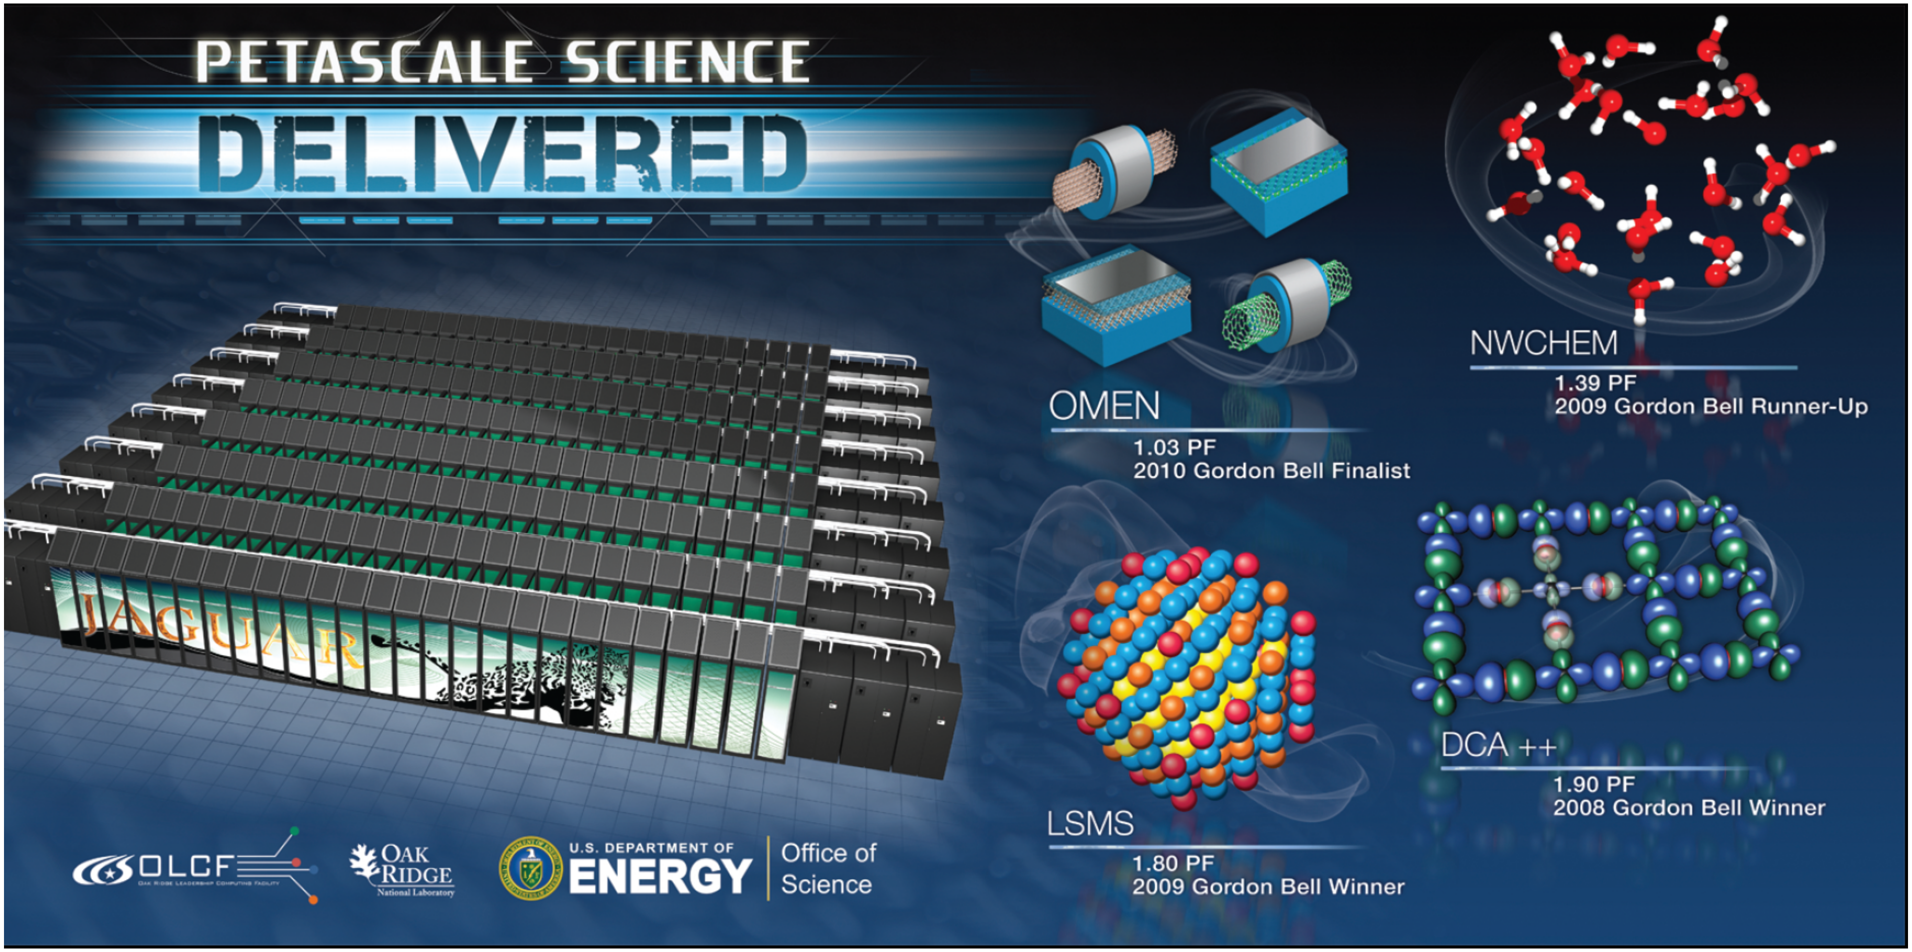

And our Department Of Energy missions require exascale computing tomorrow

Climate change	National nuclear security	Energy
Understanding and mitigating the effects of global warming – Sea level rise – Severe weather – Regional climate change – Geologic carbon sequestration	**Maintaining a safe, secure and reliable nuclear stockpile** – Stockpile certification – Predictive scientific challenges – Real-time evaluation of urban nuclear detonation	**Reducing U.S. reliance on foreign energy sources and reducing the carbon footprint of energy production** – Reducing time and cost of reactor design and deployment – Improving the efficiency of combustion energy sources

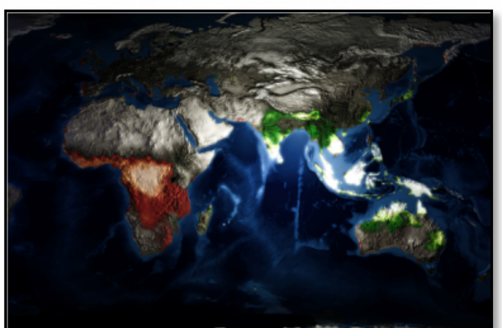

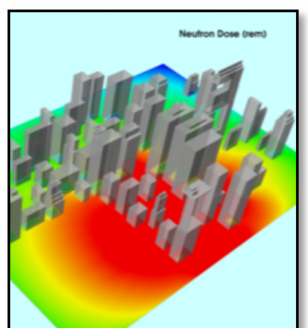

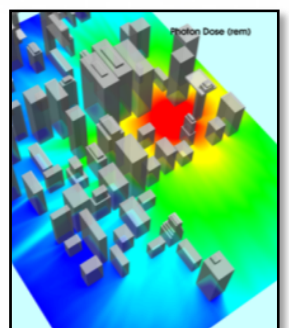

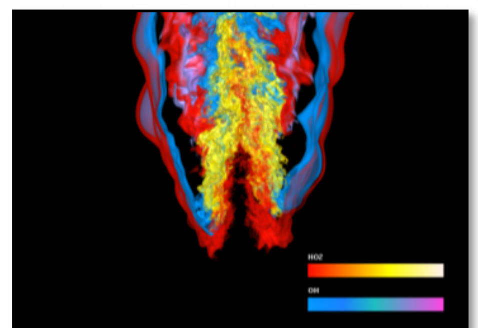

Accomplishing these missions requires exascale resources

Goals to Overcome the Barriers to Exascale

- **Power consumption goal:** Factor of 5 better than industry Business as Usual (BAU)
- **Memory and storage bandwidth goal:** Factor of 4 above industry BAU
- **Reliability and resiliency goal:** Factor of 10 better than industry BAU
- **Scalability of systems software goal:** Factor of 100 above industry BAU
- **Programming models and environments goal:** Factor of 10 productivity over today's mixed models while increasing parallelism in applications by a factor of 1,000

Trends in power efficiency

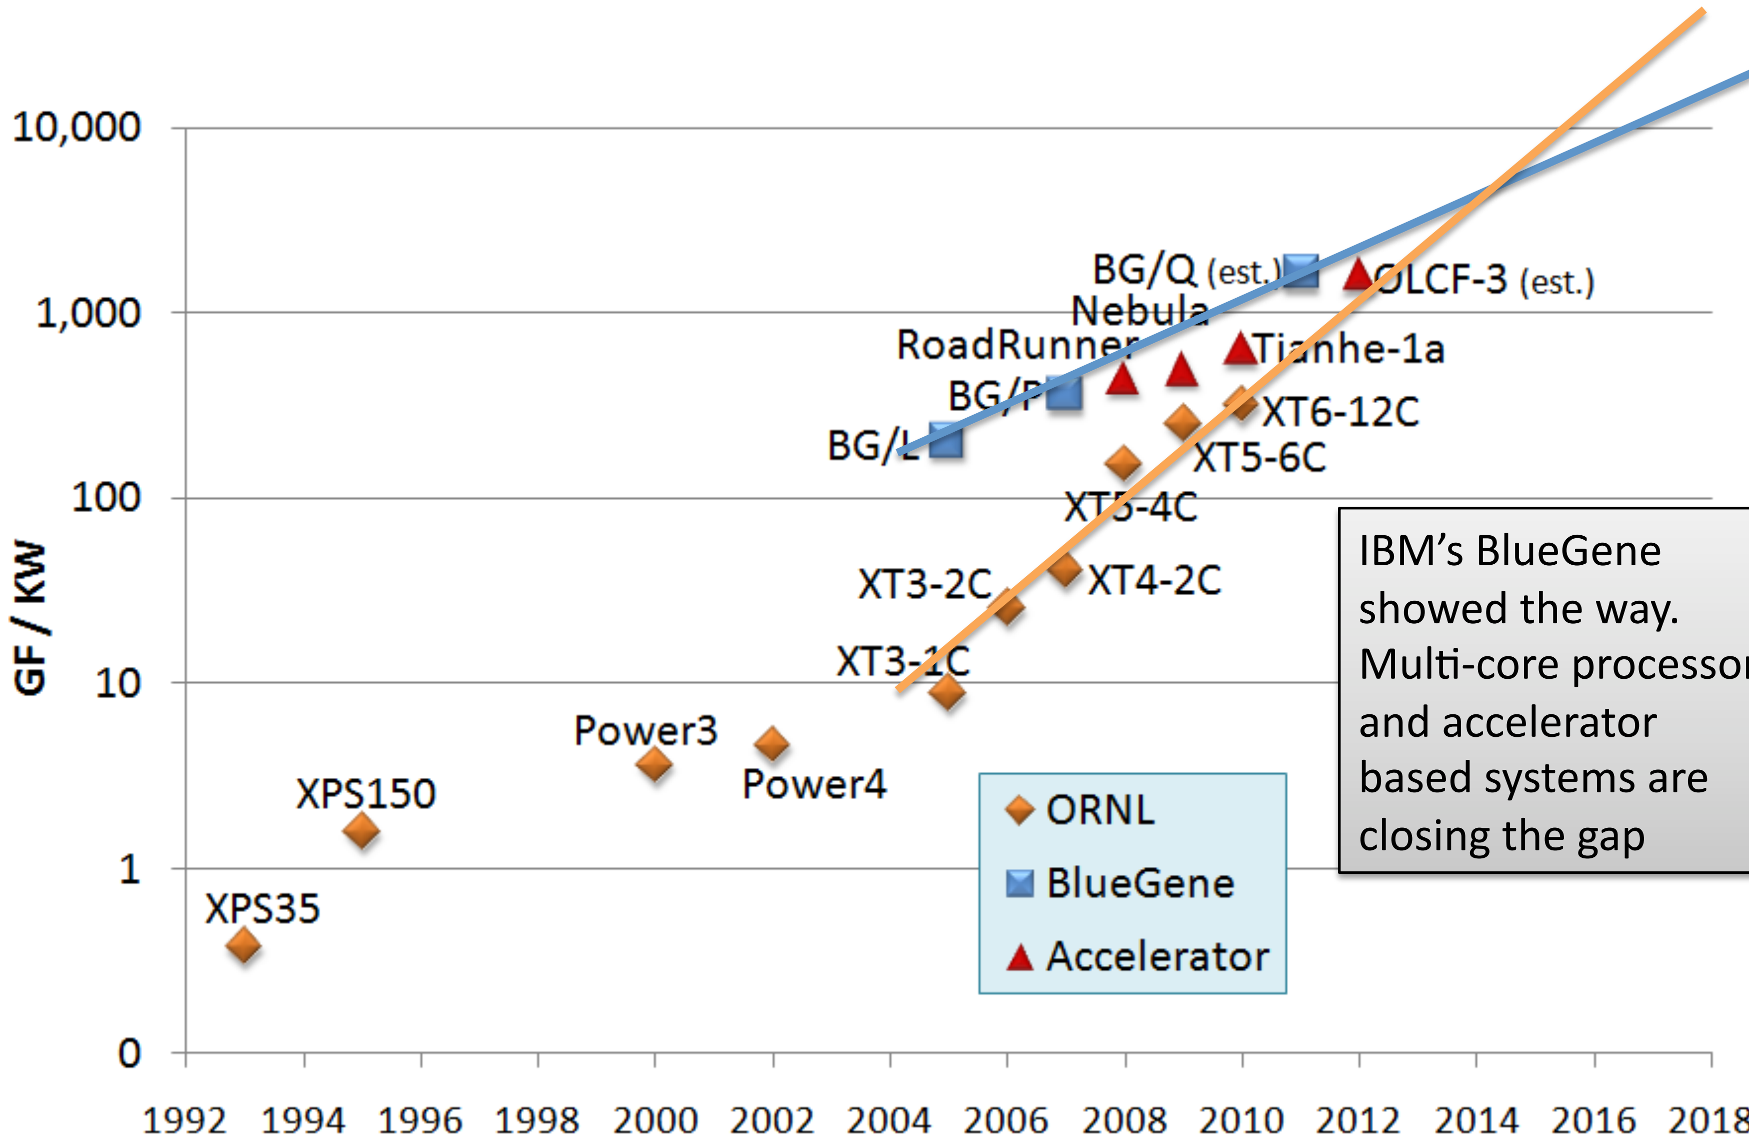

IBM's BlueGene showed the way. Multi-core processors and accelerator based systems are closing the gap

Current Technology will require huge amounts of power for Exascale systems

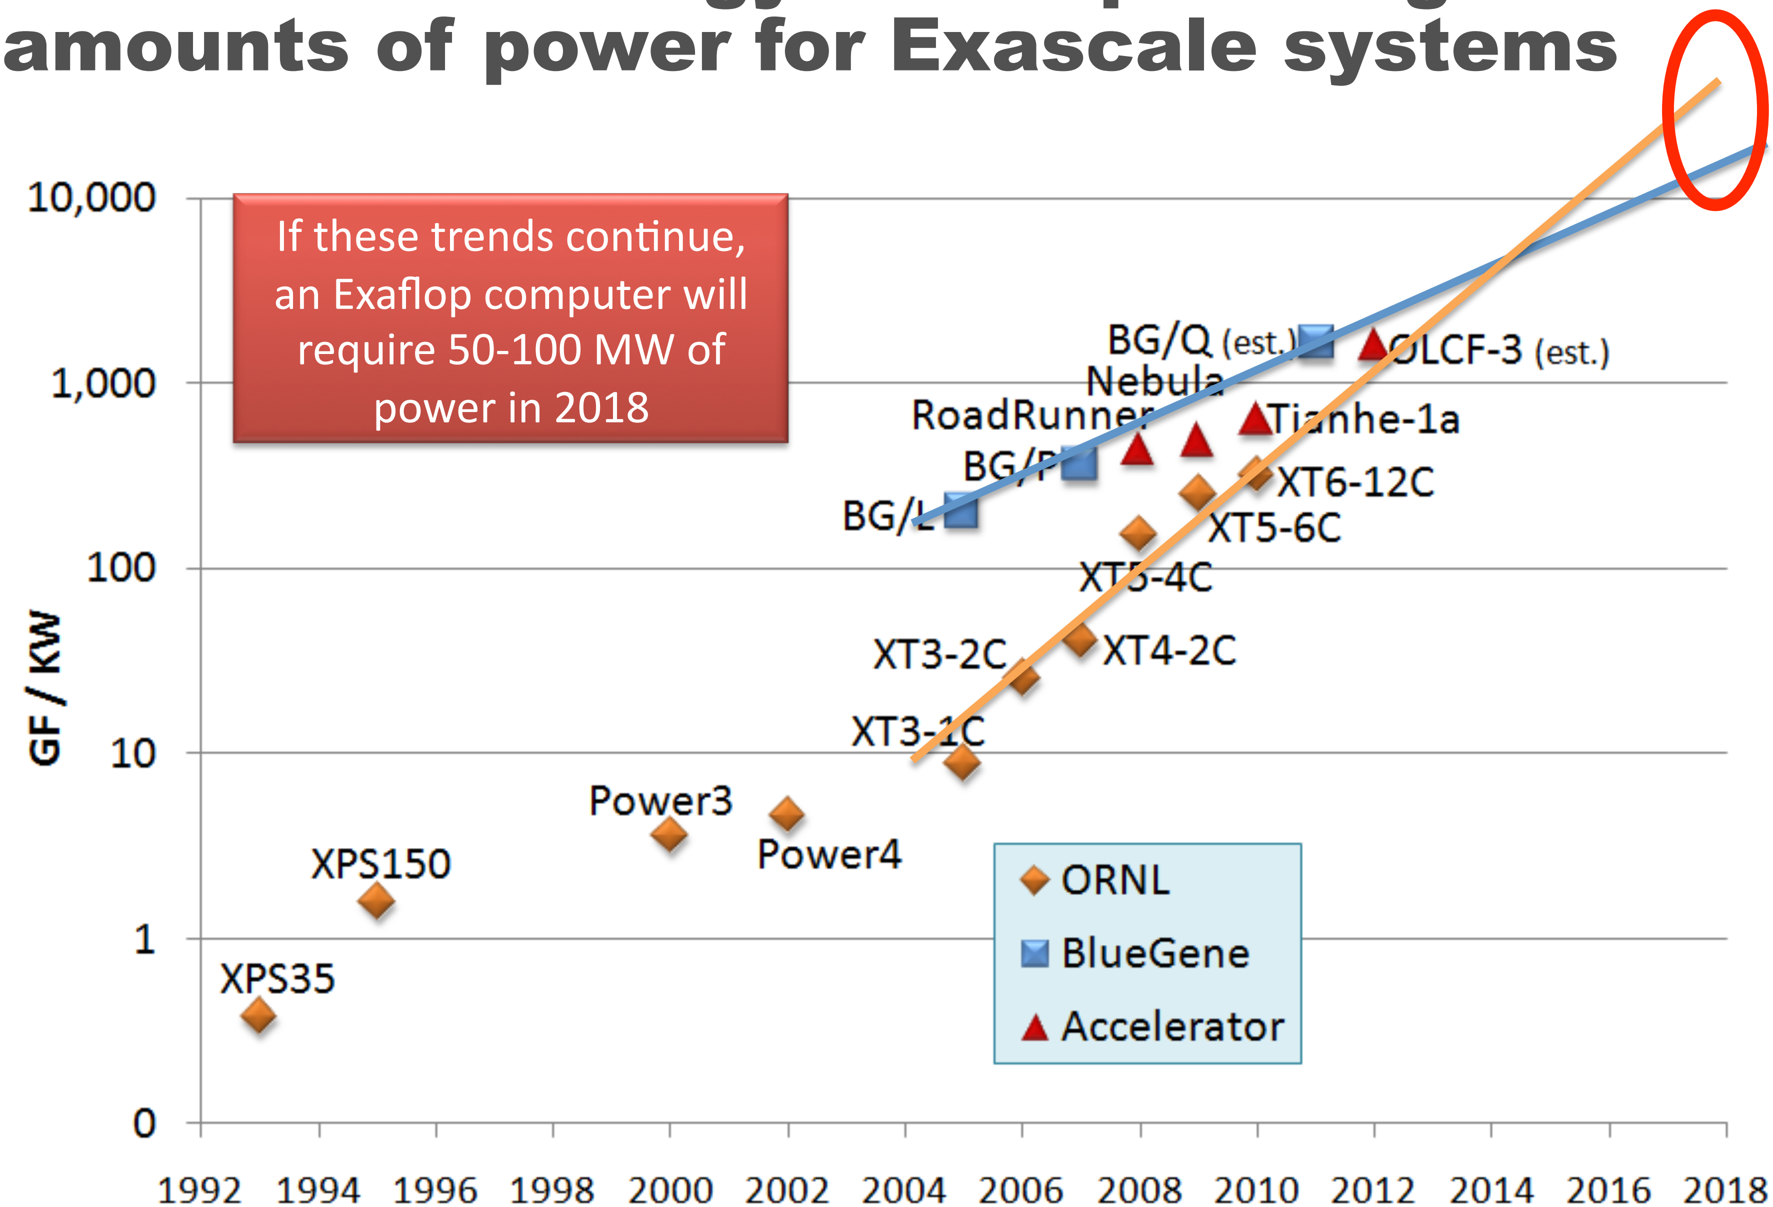

ORNL's "Titan" 20 PF System Goals

- Initial 1 PF delivery in 2011, final 20 PF system in 2012
- Designed for science from the ground up
- Similar number of cabinets, cabinet design, and cooling as Jaguar
- Operating system upgrade of today's Linux Operating System
- Gemini interconnect
  - 3-D Torus
  - Globally addressable memory
  - Advanced synchronization features
- New accelerated node design using GPUs
- 20 PF peak performance
  - 9x performance of today's XT5
- Larger memory
- 3x larger and 4x faster file system

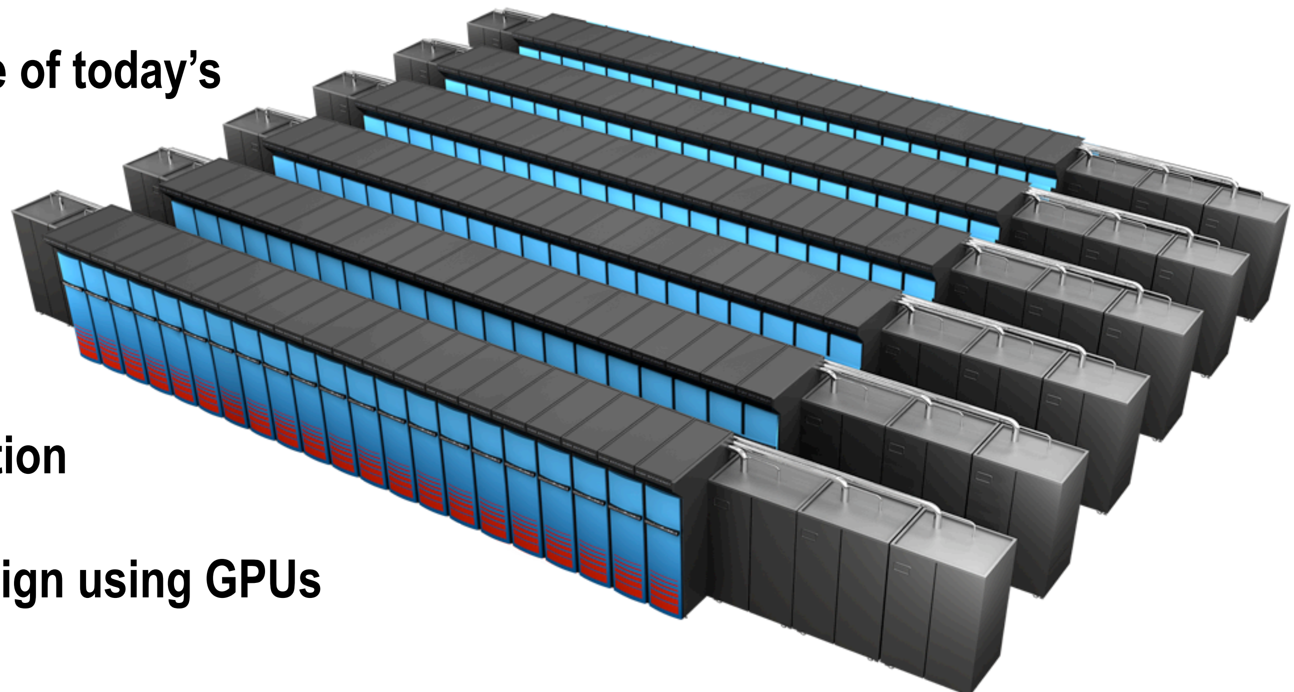

Early Science Applications

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

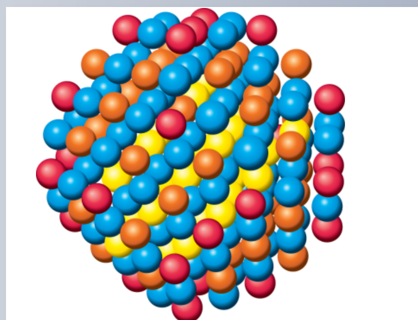

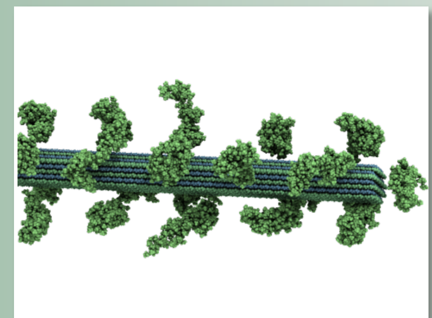

LAMMPS
Simulated time evolution of the atmospheric CO_2 concentration originating from the land's surface

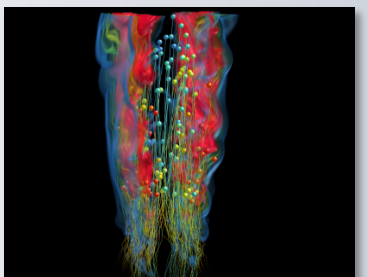

S3D
How are going to efficiently burn next generation diesel/bio fuels?
.

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

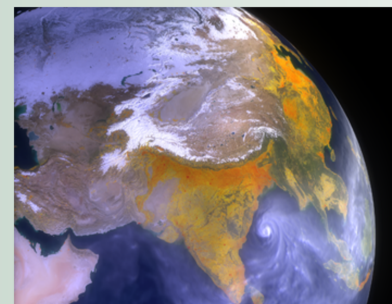

PFLOTRAN
Stability and viability of large scale CO_2 sequestration; predictive containment groundwater transport

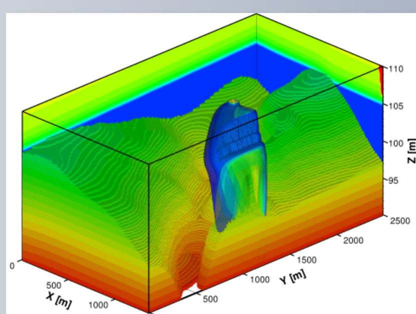

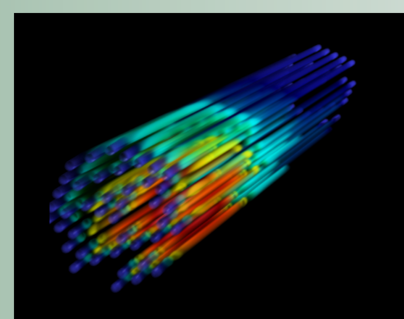

Denovo
Unprecedented high -fidelity radiation transport calculations that can be used in a variety of nuclear energy and technology applications.

Titan Project: Sustained Productivity

- **Code team for each project**
  - ☑Science team, performance engineer, applied mathematician, library specialist
- **Working with vendors on tools**
  - ☑CAPS (HMPP) – Compiler mods for accelerators, C++
  - ☑Allinea – Scale DDT to 250K cores; support for accelerators
  - ☑Vampir – Support for profiling accelerator code
  - ☑Cray – Compilers, Performance tools, unified tool set
- **Application Readiness Review of our preparation**
  - ☑Spent 6 months analyzing and porting 6 applications to a hybrid CPU/GPU platform
  - ☑Review panel was asked to assess our analysis of the challenges, level of effort, and potential for performance gains of science applications on a hybrid architecture

Results of Application Readiness Review of Titan Accelerator-Based system

– "Use of the GPU did lead to a performance relative to power cost improvement in almost all cases."
– "There is significant upside potential in GPU performance as we learn how to effectively use manycore architectures and develop new algorithms."
– "GPUs are a harbinger of all future processors to come and there is ample evidence that designing applications for today's GPUs will positively impact the performance of all multicore and manycore processors both today and in the future."
– "Giving OLCF users access to a machine that is competitive as both a CPU and GPU system will provide an excellent transition vehicle for manycore applications development."

Questions?

Buddy Bland
Email: BlandAS@ORNL.GOV

The research and activities described in this presentation were performed using the resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC0500OR22725.

OAK RIDGE
National Laboratory