

Titan: ORNL's Next Step in Leadership Computing

Jack Wells

Director of Science

Oak Ridge Leadership Computing Facility

Electronic Structure on Accelerators Workshop

Oak Ridge

February, 6 2012



ORNL's mission

Deliver scientific discoveries and technical breakthroughs that will accelerate the development and deployment of solutions in clean energy and global security, and in doing so create economic opportunity for the nation



Translational research from fundamental science to application is a hallmark of ORNL

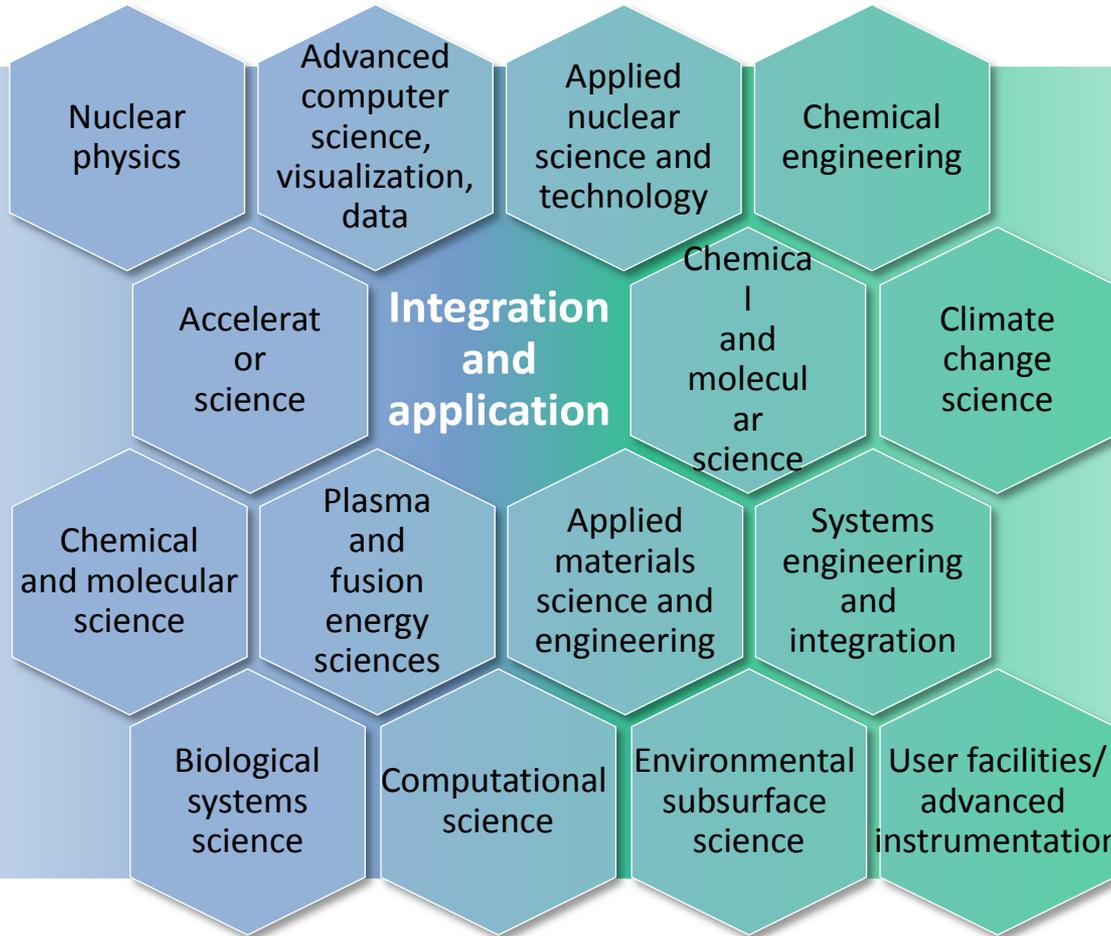
Signature strengths

Neutron science and technology

Materials science and engineering

Computational science and engineering

Nuclear science and technology



Critical mission outcomes

Scientific discovery and innovation

Clean energy

Global security

Integration and application of distinctive capabilities enables delivery of mission outcomes

Spallation
Neutron Source

Center for
Nanophase
Materials Sciences



Oak Ridge
Leadership
Computing
Facility



High Flux
Isotope
Reactor



National
Transportation
Research
Center

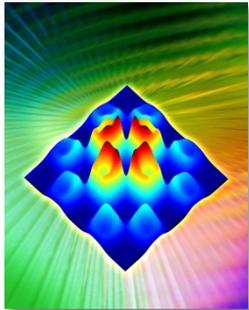


BioEnergy
Science
Center

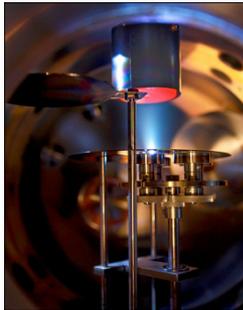


Carbon Fiber
Technology
Center

Science strategy for the future: Major initiatives



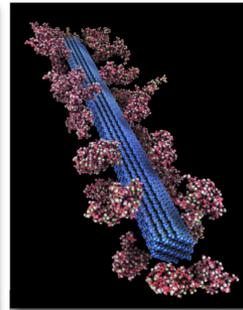
Deliver
forefront
science
using
neutrons



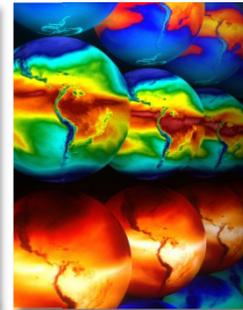
Develop and
demonstrate
advanced
materials
for energy
applications



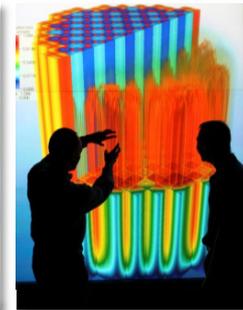
Enable
discovery
and innovation
with computing,
data infra-
structure,
and analytics



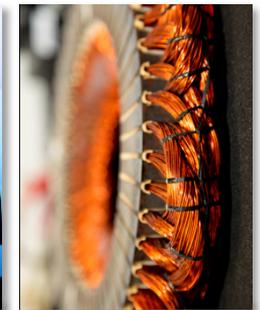
Accelerate
biomass
production
and
conversion
for energy
and materials



Advance
understanding
of climate
change
impacts



Develop,
virtualize,
and validate
advanced
nuclear
energy
systems



Deliver
sustainable
transportation
solutions

Enable discovery and innovation with computing, data infrastructure, analytics

Vision

- World's most capable complex for computational science: Infrastructure, staff, multiagency programs
- **Outcome:** Sustained world leadership in transformational research and scientific discovery through advanced computing

Strategy

Provide the nation's most capable computational site for advancing to the exascale and beyond

Leverage DOE-DOD partnership to lead U.S. development of exascale computing technology

Attract top talent; deliver outstanding user program; educate and train next generation of researchers

Leadership areas

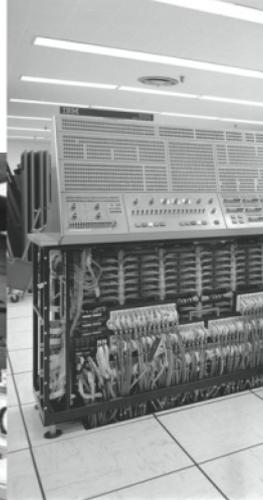
- Delivering leading-edge computational science for DOE missions
- Deploying and operating computational resources required to tackle national challenges in science, energy, and security
- Scaling applications to the exascale

ORNL has a long history in High Performance Computing

ORNL has had 20 systems

on the  **TOP 500**[®] lists
SUPERCOMPUTER SITES

1969
IBM 360/9



1954
ORACLE



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

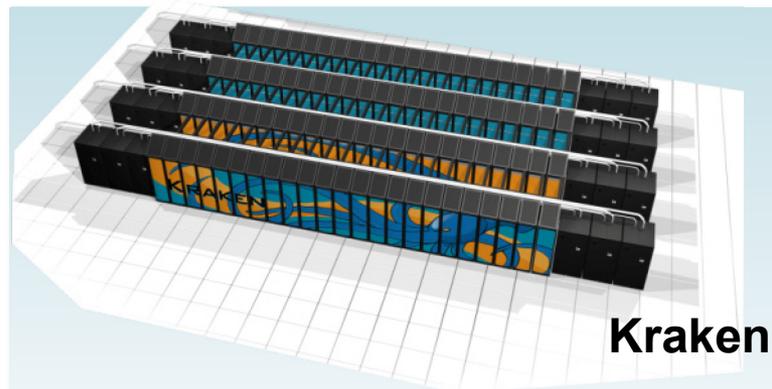


Today, we are among the world's most powerful computing facilities



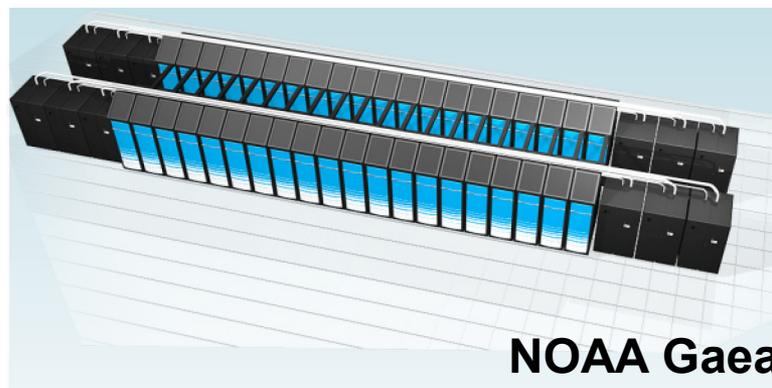
Jaguar

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



Kraken

Peak performance	1.2 PF/s
Memory	147 TB
Disk bandwidth	> 50 GB/s
Square feet	2,300
Power	3 MW



NOAA Gaea

Peak Performance	0.6 PF/s
Memory	248 TB
Disk Bandwidth	104 GB/s
Square feet	1,600
Power	1.0 MW



Dept. of Energy's most powerful computer

Nov. 2011
TOP 500[®]
SUPERCOMPUTER SITES

#3



National Science Foundation's most powerful computer

TOP 500[®]
SUPERCOMPUTER SITES

#11



National Oceanic and Atmospheric Administration's most powerful computer

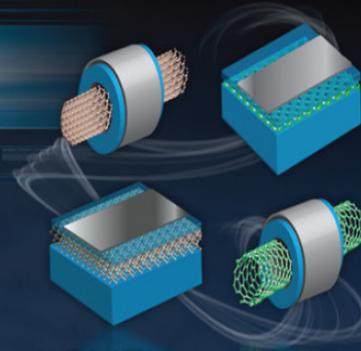
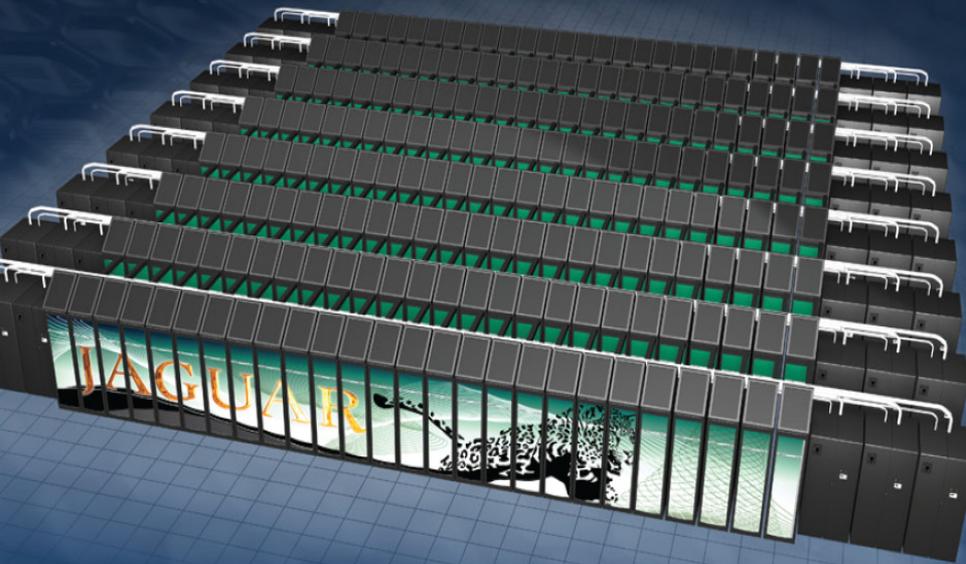
TOP 500[®]
SUPERCOMPUTER SITES

#20



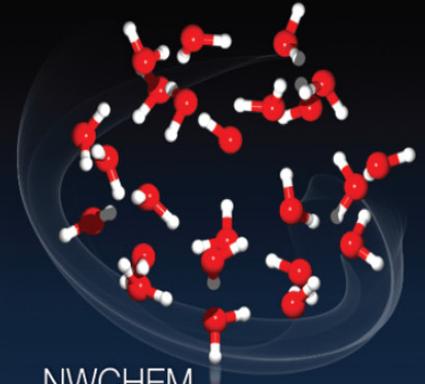
The research and activities described in this presentation were performed using the resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC0500OR22725.

PETASCALE SCIENCE DELIVERED



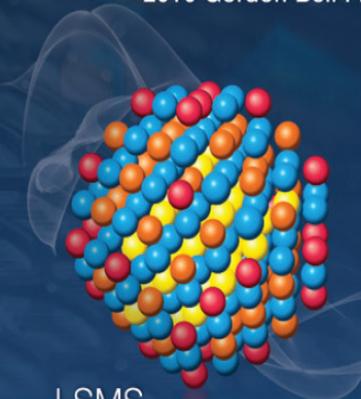
OMEN

1.03 PF
2010 Gordon Bell Finalist



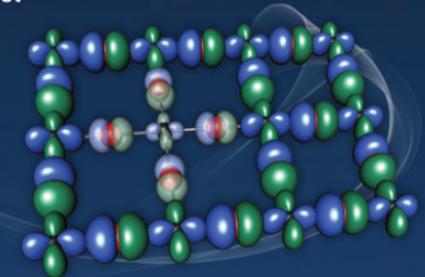
NWCHEM

1.39 PF
2009 Gordon Bell Runner-Up



LSMS

1.80 PF
2009 Gordon Bell Winner



DCA ++

1.90 PF
2008 Gordon Bell Winner

DRC

1.3 PF
2010 Gordon Bell
Honorable Mention



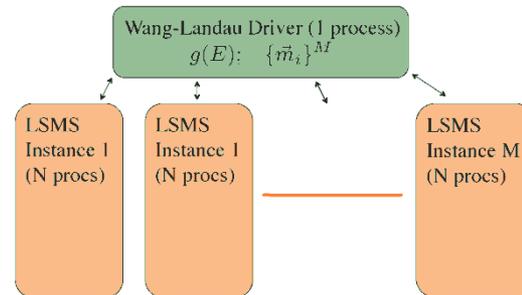
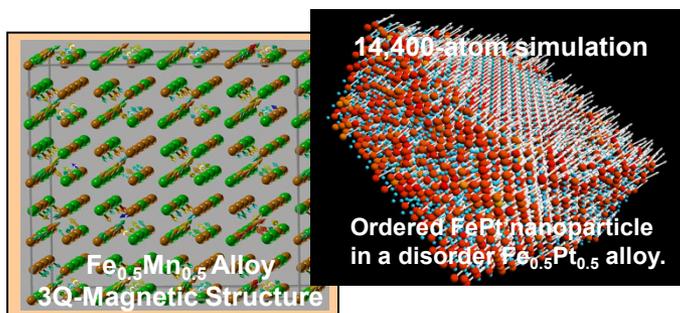
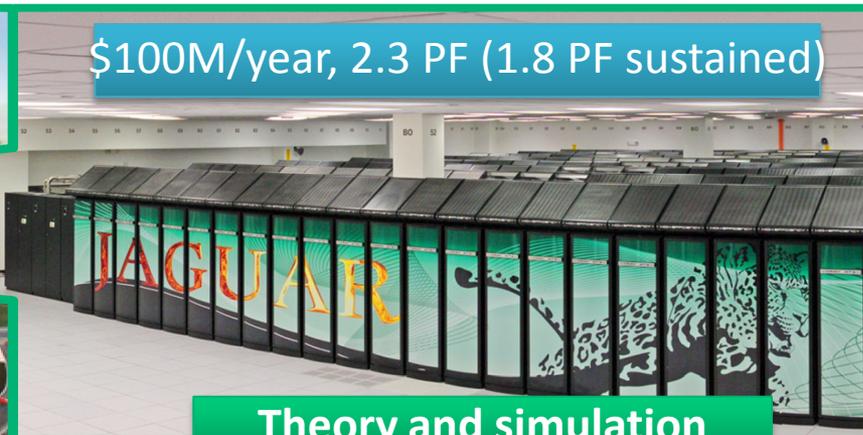
U.S. DEPARTMENT OF
ENERGY

Office of
Science

Opportunity: Extraordinary advances through integration of capabilities

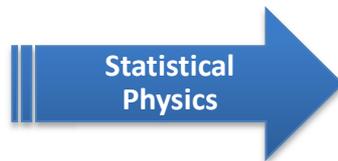
Spallation Neutron Source (SNS) and Center for Nanophase Materials Science (CNMS)

Center for Computational Sciences (CCS)



Ab Initio Electronic Structure
 Magnetic Nanoparticles
 Structural Alloys
 Simulation Cell: 10^3 - 10^4 atoms

LSMS
 10^0 - 10^2
 Tflops



Phase Stability
 Wang-Landau Statistical Sampling
 Simulation Cell: 10^3 - 10^4 atoms
 Number of Samples: 10^4 - 10^6

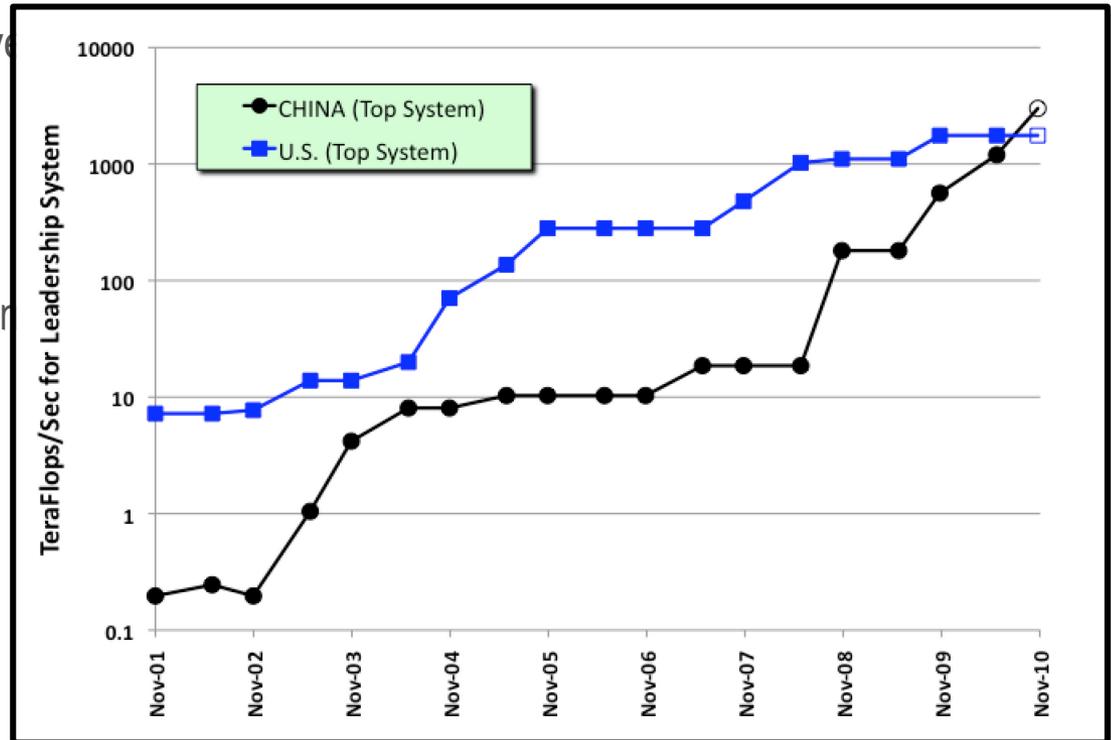
WL-LSMS
 10^0 - 10^2
 Pflops

Computational science and leadership in science and technology go hand-in-hand

- Future requires design and certification of complex engineered systems and analysis of climate mitigation alternatives with quantified levels of uncertainty
- Broader application of exascale computing can provide tremendous advantages for fundamental science and industrial competitiveness

International Competition in HPC

Chart shows total HPC resources available, from TOP500



“The United States led the world’s economies in the 20th century because we led the world in innovation. Today, the competition is keener; the challenge is tougher; and that is why innovation is more important than ever. It is the key to good, new jobs for the 21st century.”

President Barack Obama, August 5, 2009

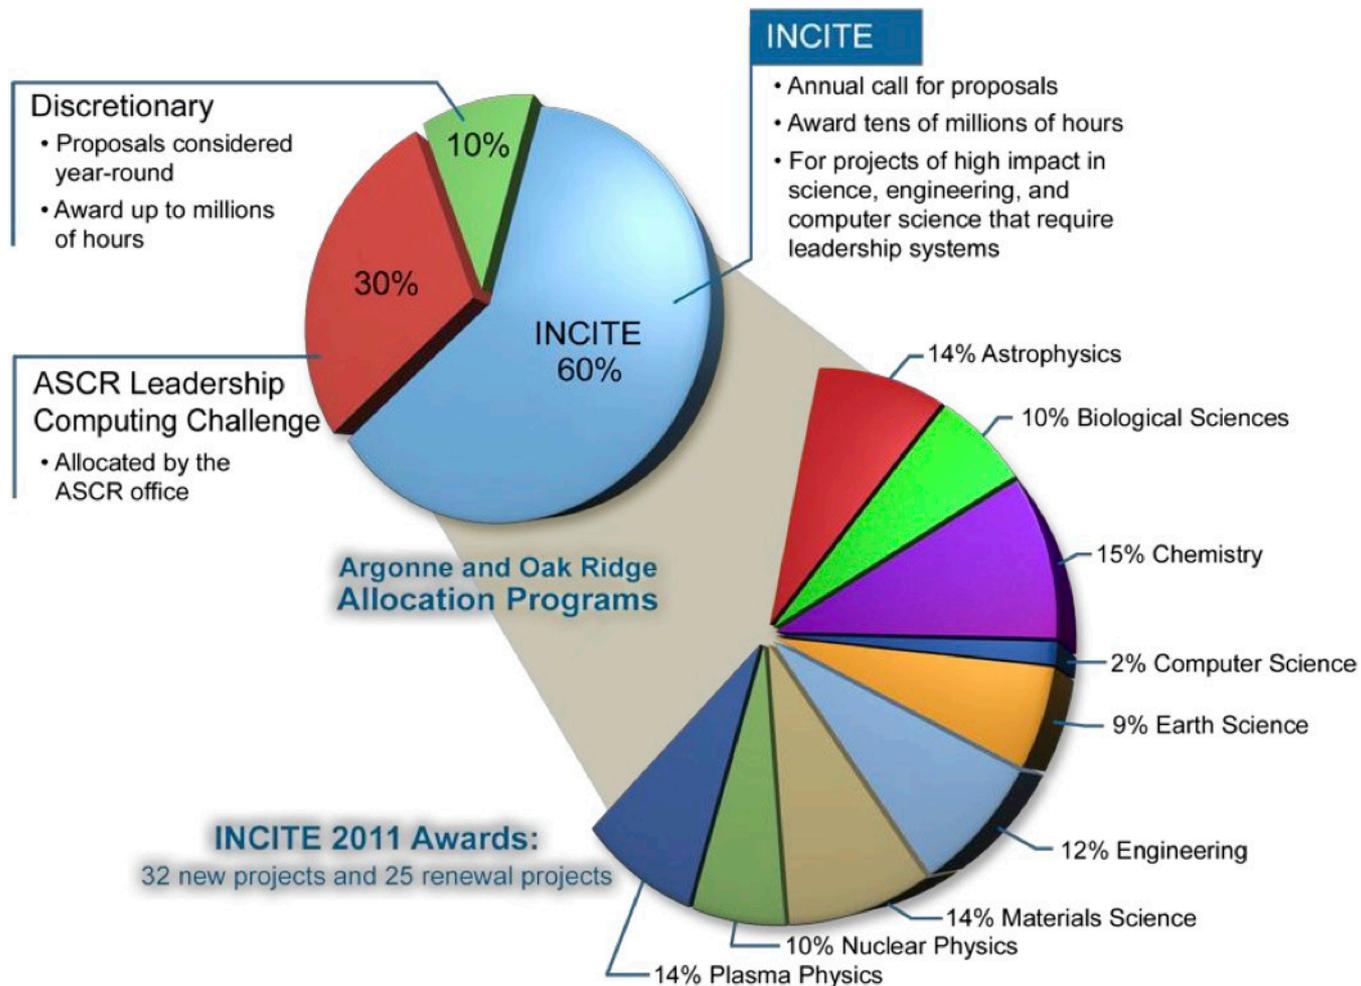
7 8 1 5 9 \$ & ' %) :

!" # !\$ % & ! (() * + , - * . ' + (/ ' ") * . 0 #) 1 2 3 , + , -) (+ 0 ! 1 2 + 4 , ') ('\$ 5 .) 6 7 ' + (/ ' % 8 2 . 6 1 . (,

./ ' . 0 1 & 5 ? @ ? + = & 5 / 3 = A 5 @ / 5 3 % # / < B / 3 8 , = 8 & 8 \$ 3 8 , 8 / 3 < & 8 * 3 ? # 4 C & 5 ? " 4 * 8 / 5 4 , & % # 7 8 % 5 3 = 2 A 8 ? @ " + 5 8 @ - & " > > ? 5 \$ 4 6 = ? " 5 5 , = 8 / 3 < & % / & = ? 5 / 4 , 8 / & * , 8 5 4 ? 3 3 , & / 3 < & % E F \$ 4 , 8 G , % , 5 * \$ & ? @ " 8 4 & / 8 8 \$ \$ = & G H = J & A 5 8 ? @ " + / 5 3 % # C 8 3 + , 3 = 8 , - & # / 5 4 , K < 8 % # 6 = , % 8 * & 5 ? 7 , 8 = (

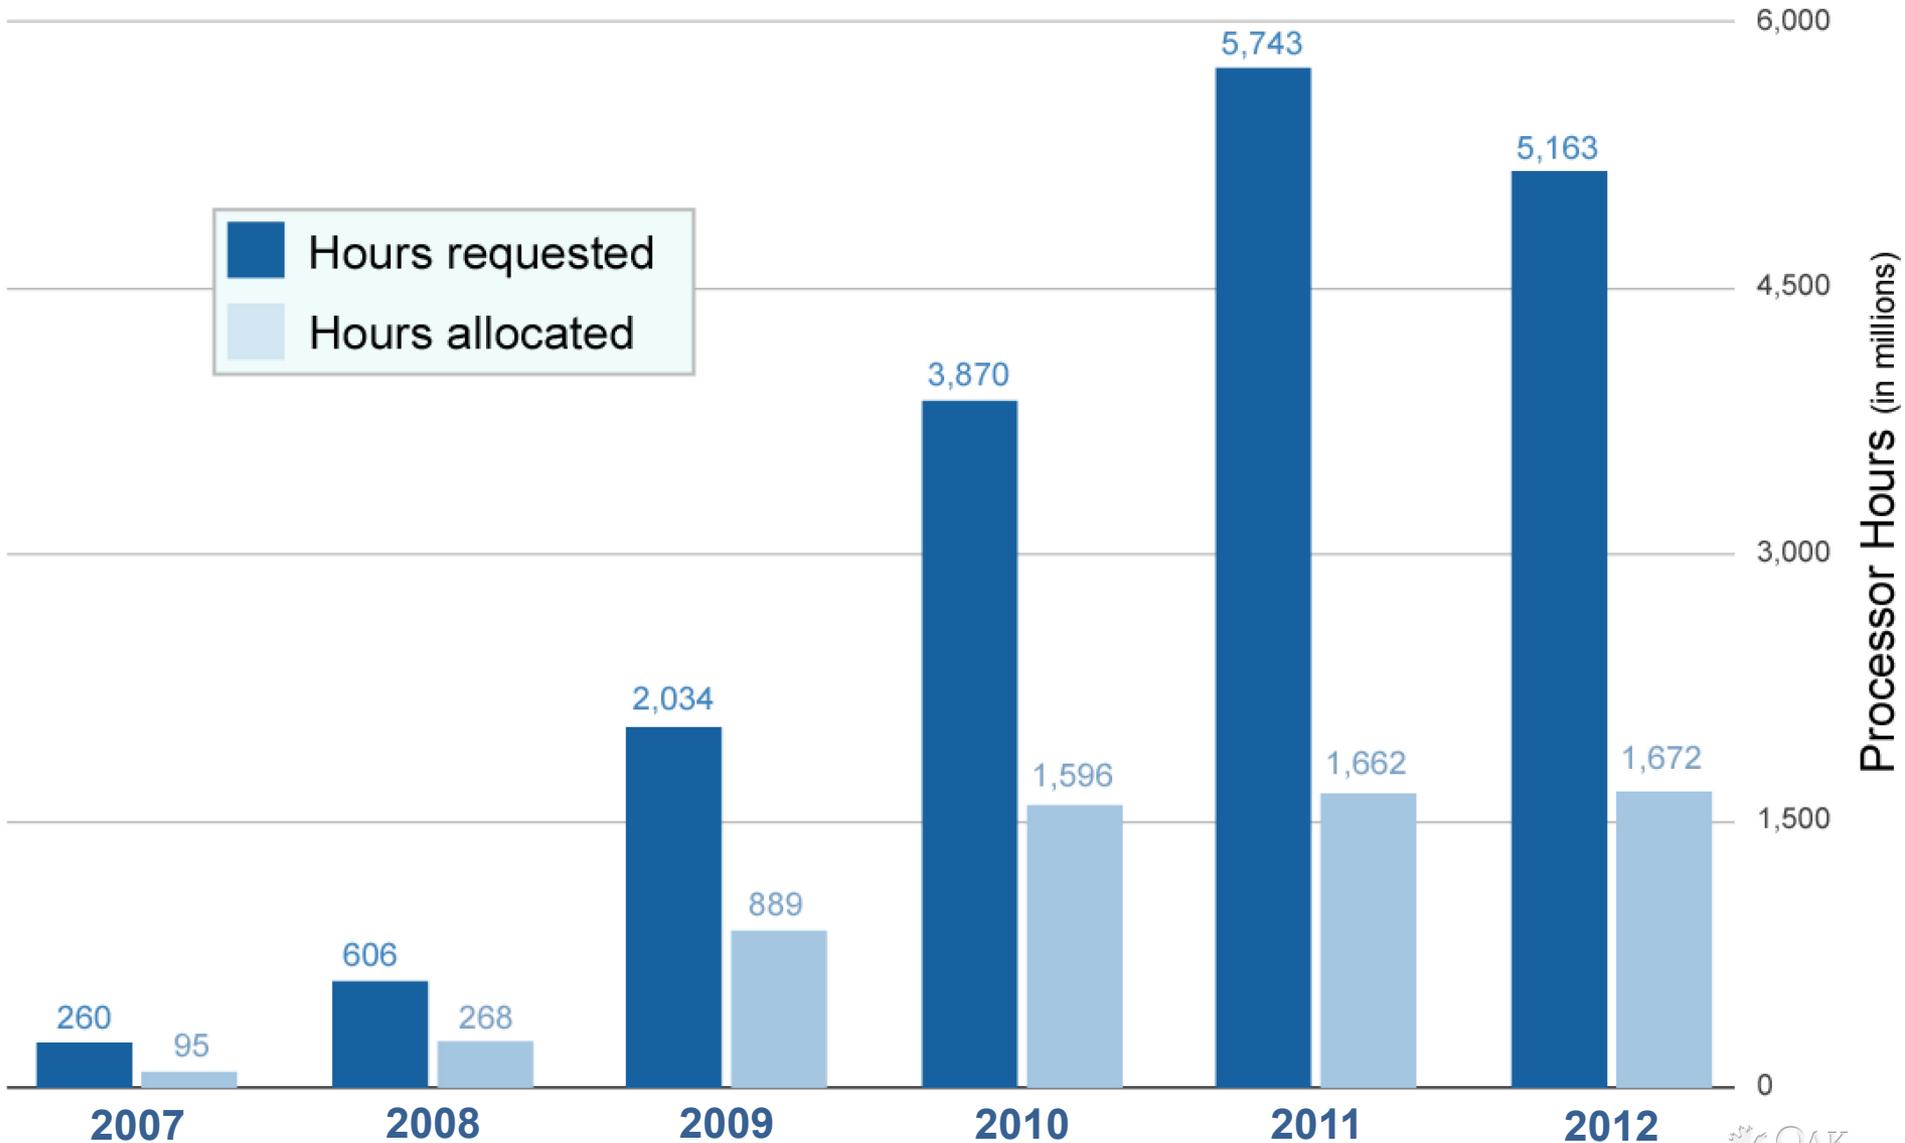


2011 INCITE awards: 1.7 billion hours out of 5.7 billion requested



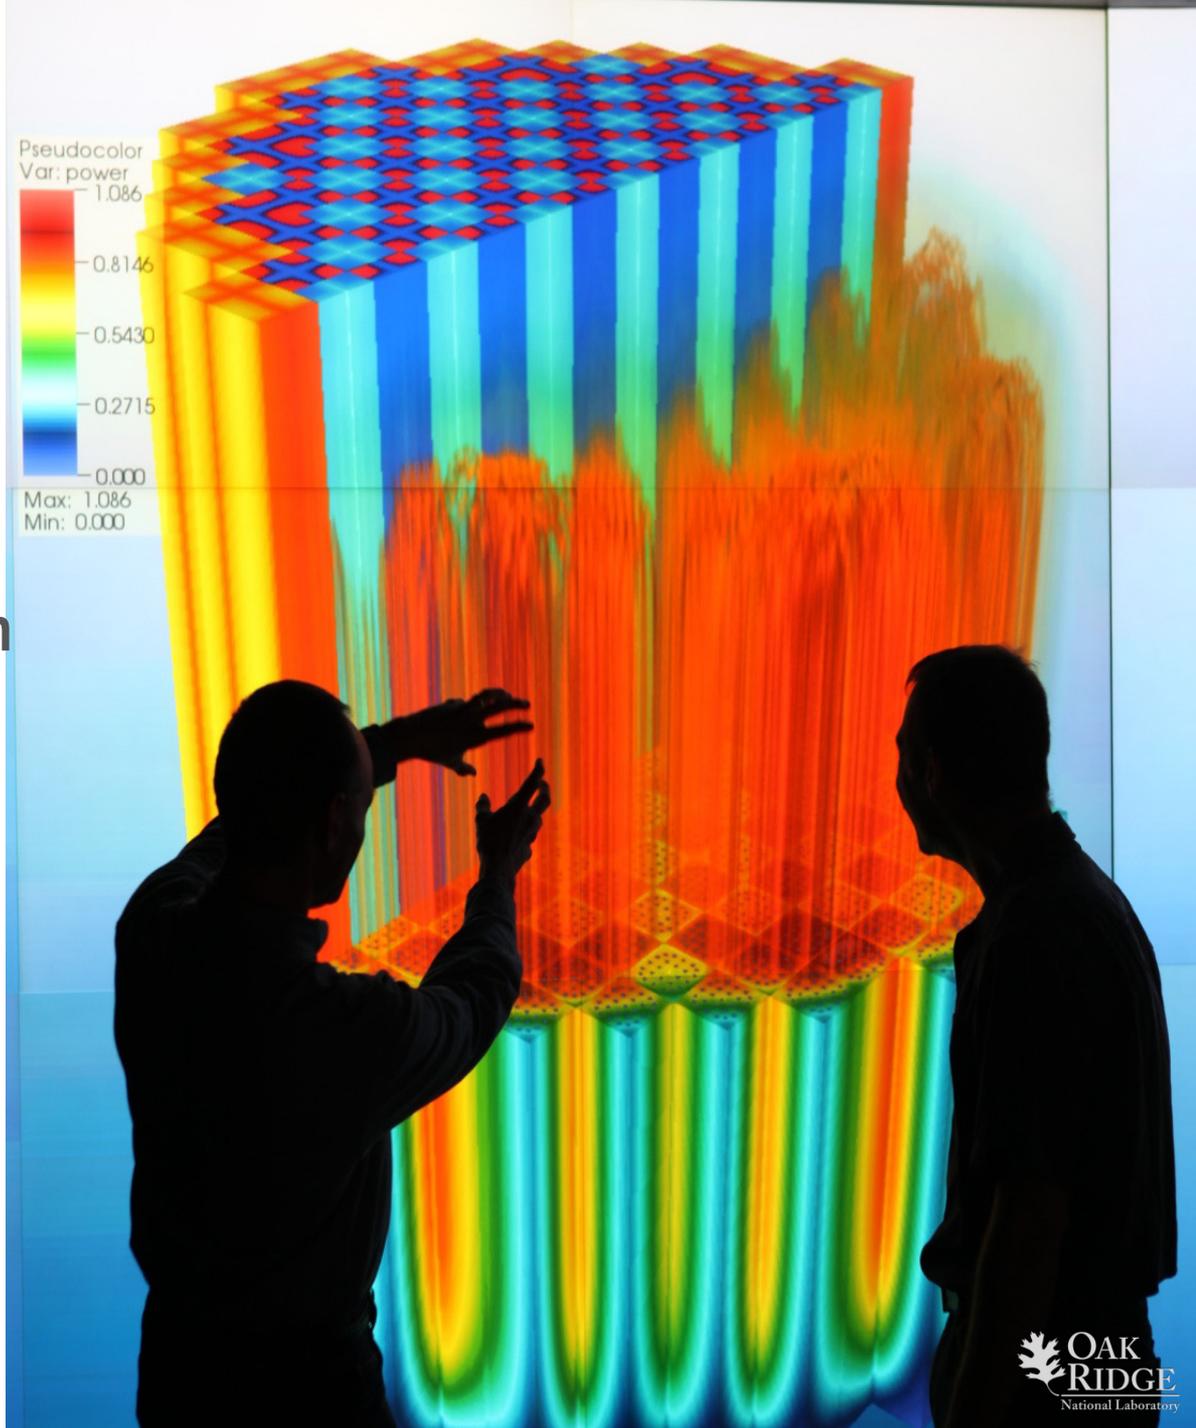
Percentage of total 2011 INCITE hours by domain
33% of new submittals and 89% of renewals accepted

INCITE is 2.5 to 3.5 times oversubscribed



What is Titan?

- The next phase of the Leadership Computing Facility program at ORNL
- An upgrade of Jaguar from 2.3 Petaflops today to between 10 and 20 PF by the end of 2012 with operations in 2013
- Built with Cray's newest XK6 compute blades



Cray XK6 Compute Node

XK6 Compute Node Characteristics

AMD Opteron 6200 Interlagos
16 core processor @ 2.2GHz

Tesla M2090 @ 665 GF with
6GB GDDR5 memory

Host Memory
32GB
1600 MHz DDR3

Gemini High Speed Interconnect

Upgradeable to NVIDIA's
next generation KEPLER
processor in 2012

Four compute nodes per XK6
blade. 24 blades per rack



ORNL's "Titan" System

- Upgrade of existing Jaguar Cray XT5
- Cray Linux Environment operating system
- Gemini interconnect
 - 3-D Torus
 - Globally addressable memory
 - Advanced synchronization features
- AMD Opteron 6200 processor (Interlagos)
- New accelerated node design using NVIDIA multi-core accelerators
 - 2011: 960 NVIDIA M2090 "Fermi" GPUs
 - 2012: 10-20 PF NVIDIA "Kepler" GPUs
- 10-20 PF peak performance
 - Performance based on available funds
- 600 TB DDR3 memory (2x that of Jaguar)



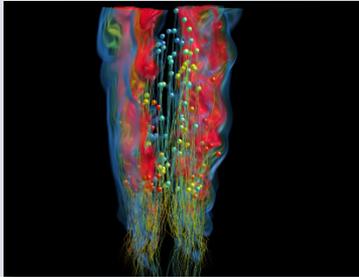
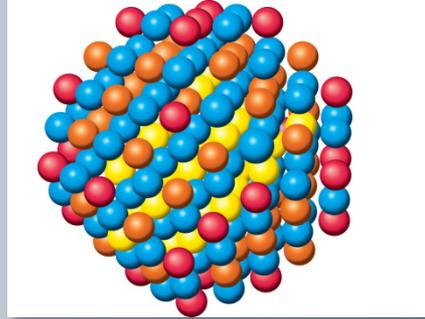
Titan Specs

Compute Nodes	18,688
Login & I/O Nodes	512
Memory per node	32 GB + 6 GB
NVIDIA "Fermi" (2011)	665 GFlops
# of Fermi chips	960
NVIDIA "Kepler" (2012)	>1 TFlops
Opteron	2.2 GHz
Opteron performance	141 Gflops
Total Opteron Flops	2.6 Pflops
Disk Bandwidth	~ 1 TB/s

Titan: Early Science Applications

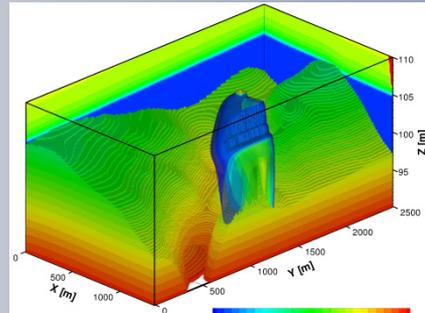
WL-LSMS

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



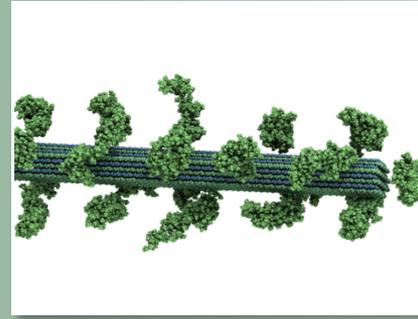
S3D

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



PFLOTRAN

Stability and viability of large scale CO₂ sequestration; predictive containment groundwater transport

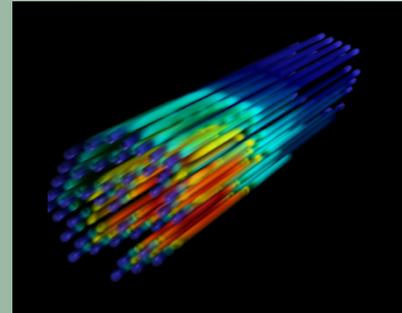


CAM / HOMME

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

LAMMPS

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



Denovo

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

OLCF-3 Project Timeline

Date	Event	Description
Jan 2009	CD-0	Statement of Mission Need for LCF program upgrade signed
July 2009	Lehman Review	Presented plan for hybrid system
Jan 2010	CD-1	Alternative selection and cost baseline approved
July 2010	Application Review	Present plan for application performance and portability across many platforms
Dec 2010	Lehman Review	Presented baseline and acquisition plan
Feb 2011	CD-3a	Approve long lead time acquisitions
Aug 2011	Lehman Review & CD-2/3b	Approved performance baseline and system acquisition
Oct 2011	Begin Upgrade	96 of Jaguar's 200 cabinets upgraded to new Cray XK6 blades
2H 2012	Kepler GPUs	Add in the 10-20 PF of NVIDIA Kepler GPUs
June 2013	Complete Acceptance Testing	Finish all acceptance testing and turn completed machine over to INCITE program

2011 Upgrade from XT5 to XK6

Oct: Segment system into 104 cabinets of existing Jaguar for users (1.2 PF)

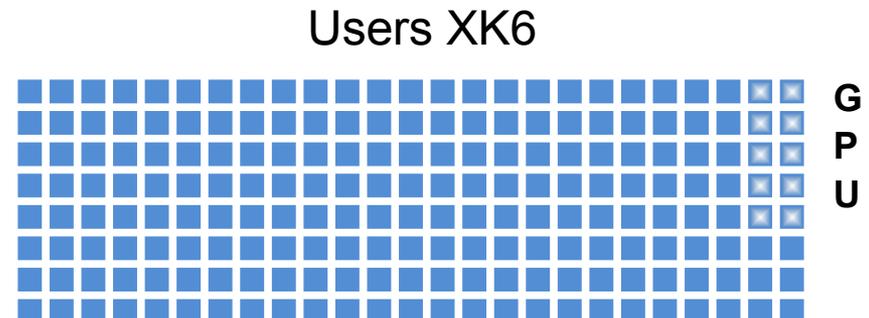
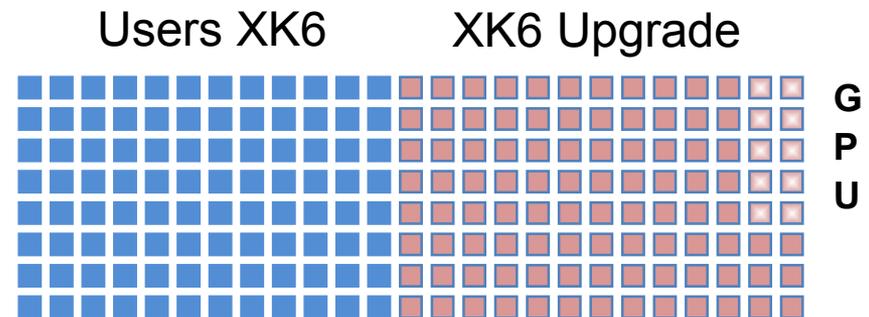
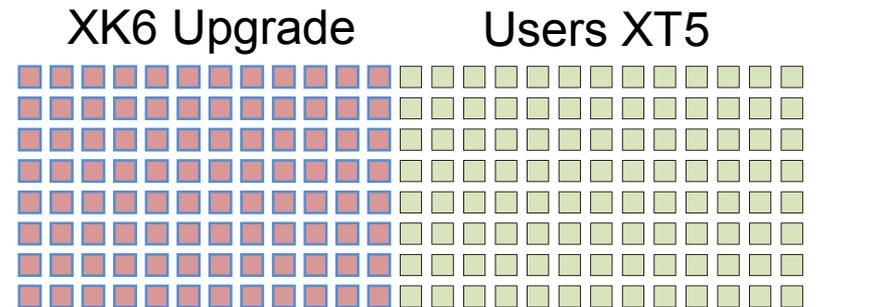
- Upgrade and test 96 cabinets to XK6 (16-core) nodes

Nov: Move users to upgraded 96 cabinets (1.3 PF)

- Upgrade and test 104 cabinets to XK6 nodes including 10 cabinets with GPUs

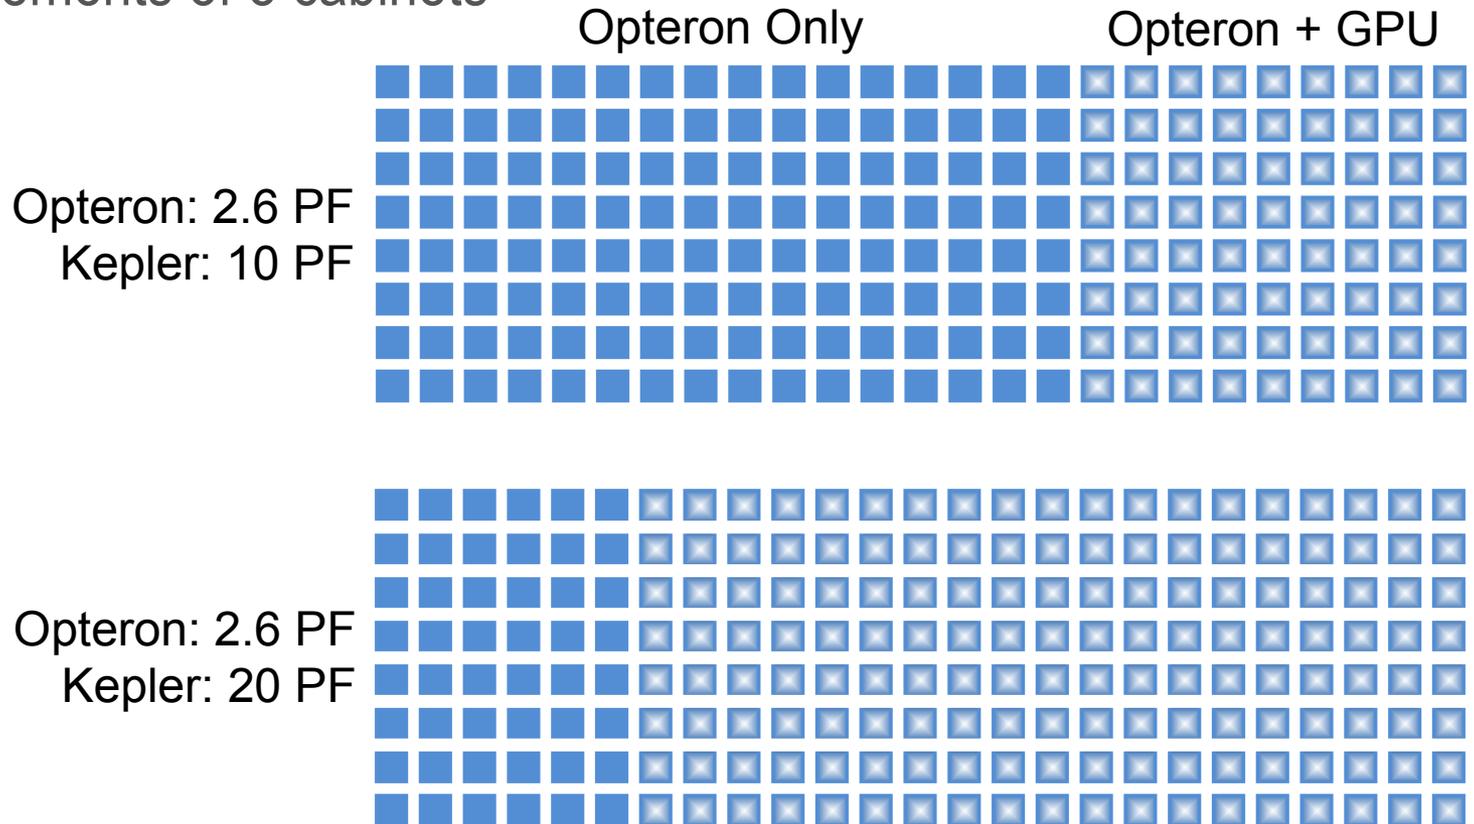
Dec: Combine two halves and run acceptance test (3.3 PF)

- Users will be running during acceptance test period



2012 Upgrade: Add Kepler GPUs

- When NVIDIA Kepler GPUs are available we will replace Fermi GPUs and increase system to 10 to 20 petaflops based on available funding
- The contract with Cray has options to add GPUs in increments of 8 cabinets



TITANdev – What is it?

- 10 cabinets of XK6
 - 1st glimpse of Titan
- Program for early access
 - Limited priority access for early porting, debugging, performance analysis, benchmarking, etc
 - Can submit jobs to GPU nodes
 - Increased priority for job execution

TITANdev – the machine

- 10 cabinets of Jaguar XK6 with GPUs
 - Approximately 1 week away depending on acceptance test progress
- 960 nodes
 - Current generation Tesla (Fermi +)

TITANdev Program – Why?

Generate significant scientific results on latest hybrid architecture

Provide an opportunity for application teams to explore new architecture and produce exciting, forward looking scientific results

TITANdev Program – How?

- Allow users priority access to GPUs
 - Lower competition from non-GPU jobs
 - Limited support
 - Reporting requirements
- Advantages to users
 - Time to develop apps on latest architecture
 - Experiences can strengthen future proposals (INCITE, DD, etc)
- Advantages to OLCF
 - Lessons-learned
 - More visibility into user experience
 - Better able to address issues early

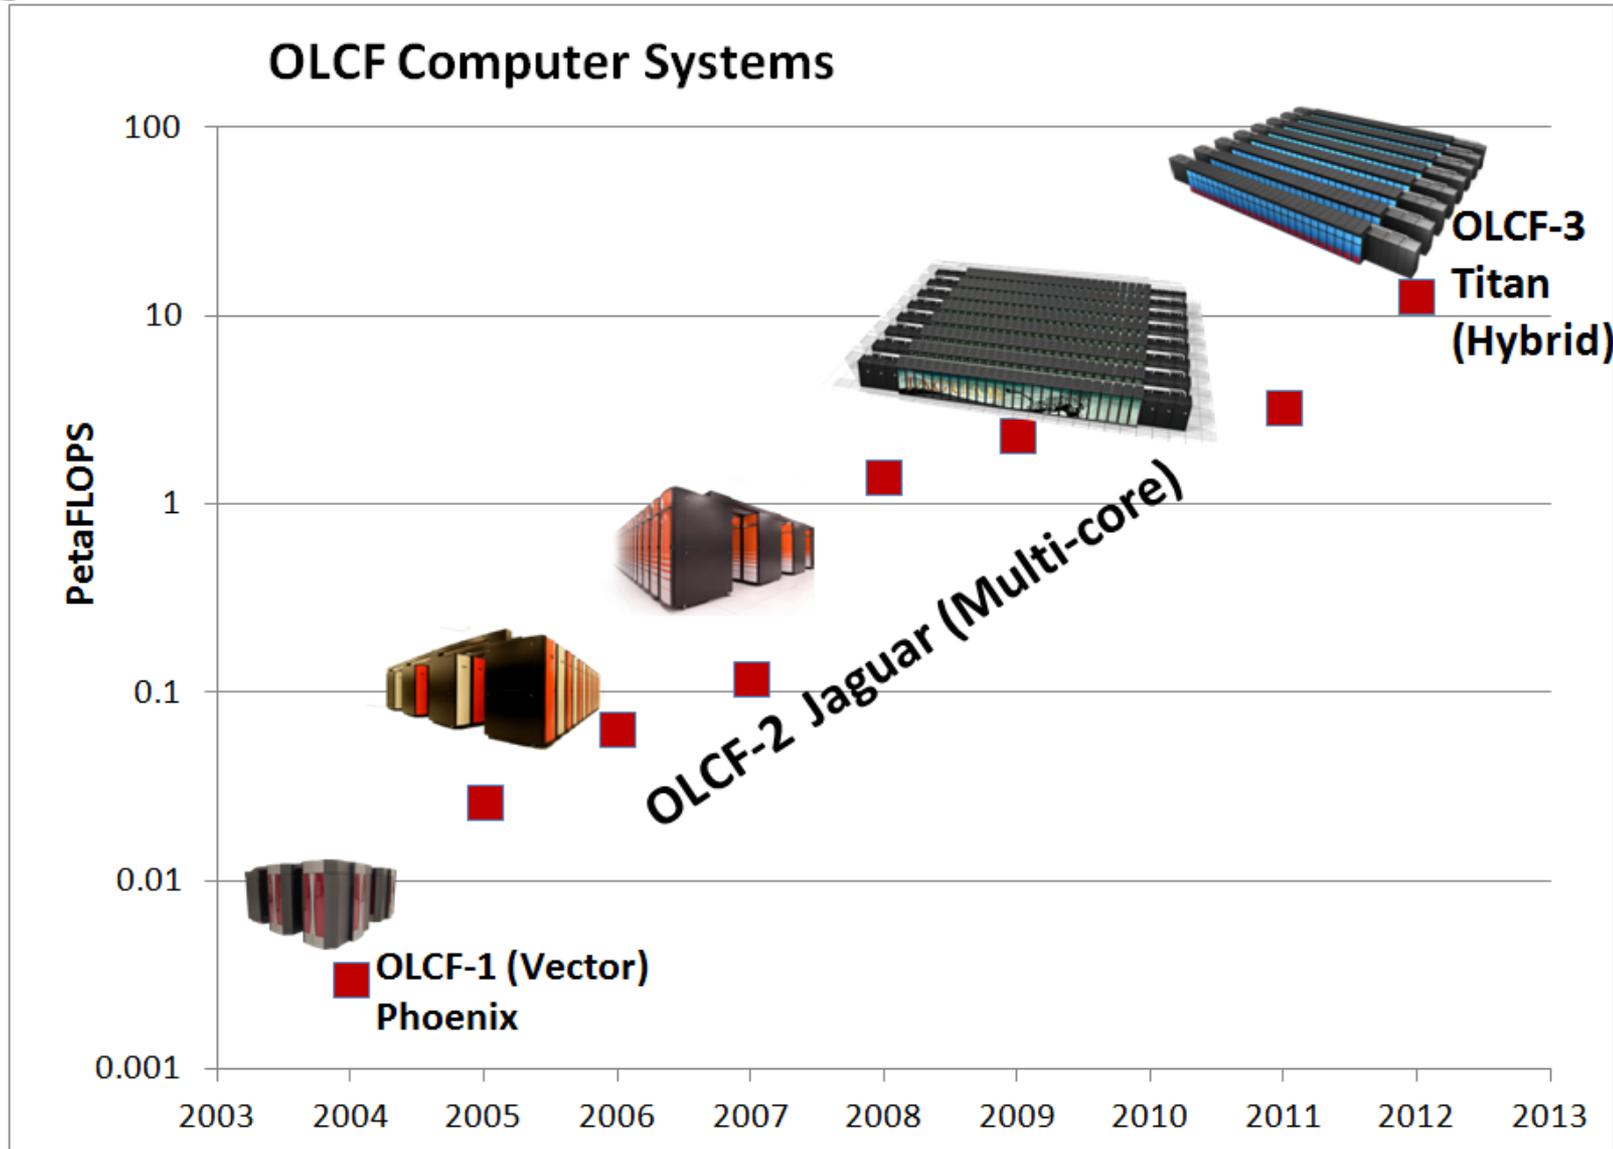
TITANdev Program – How?

- Targeted at users (not projects)
- Users complete streamlined request
 - Username, application, strategy (tool use, programming model, etc)
 - Agree to use agreement for TITANdev
 - Development resource
 - Limited support
 - Streamlined bi-weekly status reports
 - Small number of questions designed to gauge progress, issues/concerns
 - Used to better support for TITANdev

Program Details

- Limited to approximately 50 users
- Limited to 2-3 users per project
- Must already have OLCF project / account
 - If not, apply for DD account
- Will be ready soon, after acceptance is complete

Oak Ridge Leadership Computing Systems



Questions & Discussion

