The Year in Review and Positioning for the Future

20 Years of Excellence in Computational Science

OLCF

Oak Ridge Leadership Computing Facility

1992–2012

Presented by:

Jack C. Wells

Director of Science

Oak Ridge Leadership Computing Facility (OLCF)

National Center for Computational Sciences (NCCS)

OLCF Users' Meeting

April 12, 2012

Oak Ridge

Oak Ridge Leadership Computing Facility 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.

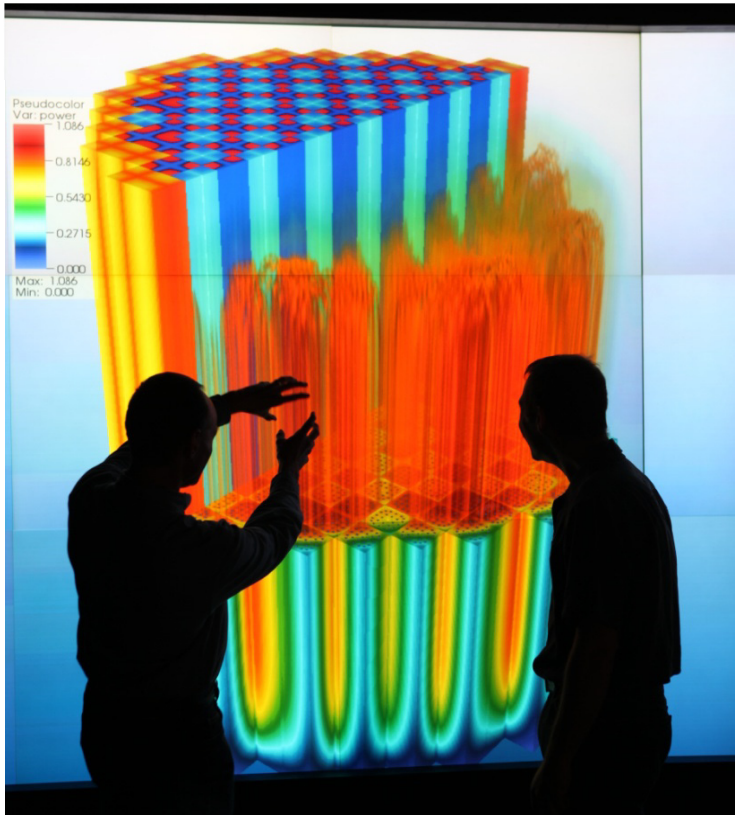

Breakthrough Science at Every Scale

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

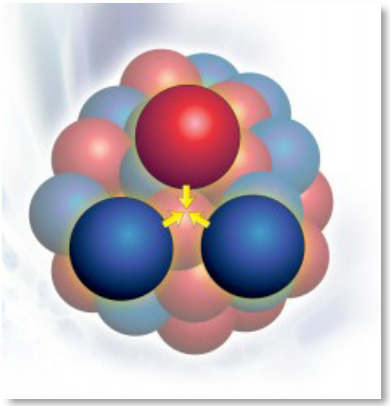

New Materials
Lopez-Bezanilla et al., discover that boron-nitride monolayers are an ideal dielectric substrate for future nanoelectronic devices constructed with graphene as the active layer

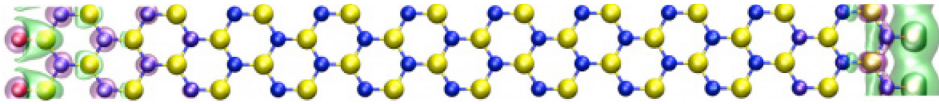

Biofuels
Smith et al., reveal the surface structure of lignin clumps down to 1 angstrom

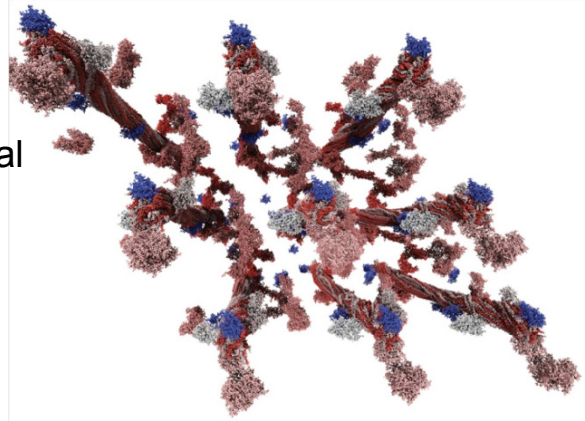

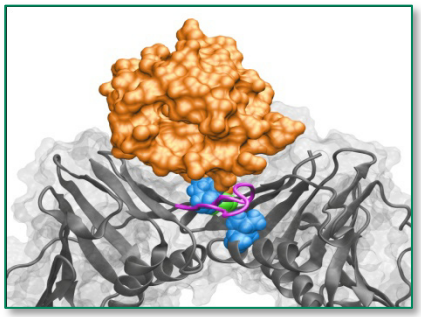

Biochemistry
Ivanov et al., illuminate how how DNA replication continues past a damaged site so a lesion can be repaired later

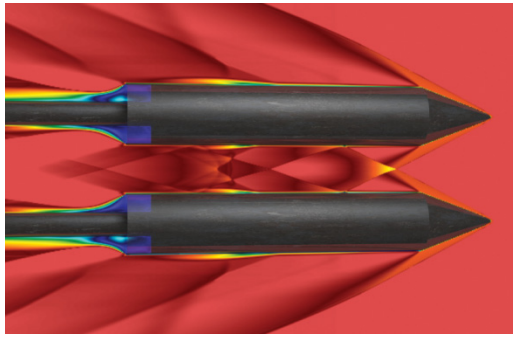

Design Innovation
Ramgen Power Systems accelerates their design of shock wave turbo compressors for carbon capture and sequestration

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

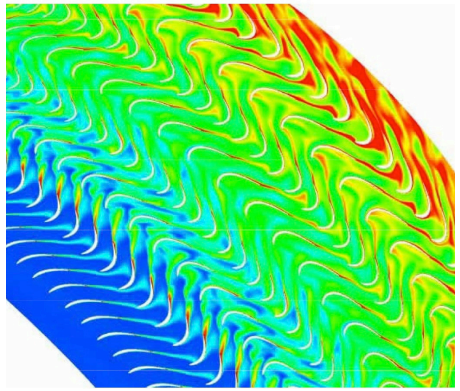

Scaling of Codes on 224,256 core Jaguar

Science Area	Code	Contact	Institution	Cores	Total Performance	Notes
CFD	MoBo	Rahimian	Georgia Tech	200,000	700 TF	2010 Gordon Bell Winner
Seismology	AWP-ODC	Cui	SDSC	223,074	220 TF	2010 Gordon Bell Finalist
Materials	DRC	Kozhevnikov	ETH Zurich	186,624	1.3 PF	2010 Gordon Bell Honorable Mention
Geoscience	P4est	Burstedde	UT-Austin	224,000		2010 Gordon Bell Finalist
Materials	DCA++	Schulthess	ETH & ORNL	213,120	1.9 PF*	2008 Gordon Bell Winner
Materials	WL-LSMS	Eisenbach	ORNL	223,232	1.8 PF	2009 Gordon Bell Winner
Chemistry	NWChem	Apra	ORNL PNNL	224,196	1.4 PF	2009 Gordon Bell Finalist
Nano Materials	OMEN	Klimeck	Purdue	221,400	1.28 PF	2011 Gordon Bell Finalist
Seismology	SPECFEM3D	Carrington	SCEC	149,784	165 TF	2008 Gordon Bell Finalist
Weather	WRF	Michalakes	NREL	150,000	50 TF	
Combustion	S3D	Chen	Sandia	144,000	83 TF	
Fusion	GTC	Chang	PPPL	102,000	20 billion Particles / sec	
Materials	LS3DF	Lin-Wang Wang	Berkeley	147,456	442 TF	2008 Gordon Bell Winner
Chemistry	MADNESS	Harrison	ORNL UTenn	140,000	550+ TF	

OLCF User Demographics

- Number of Users: 857

Percentage of Users by Affiliation

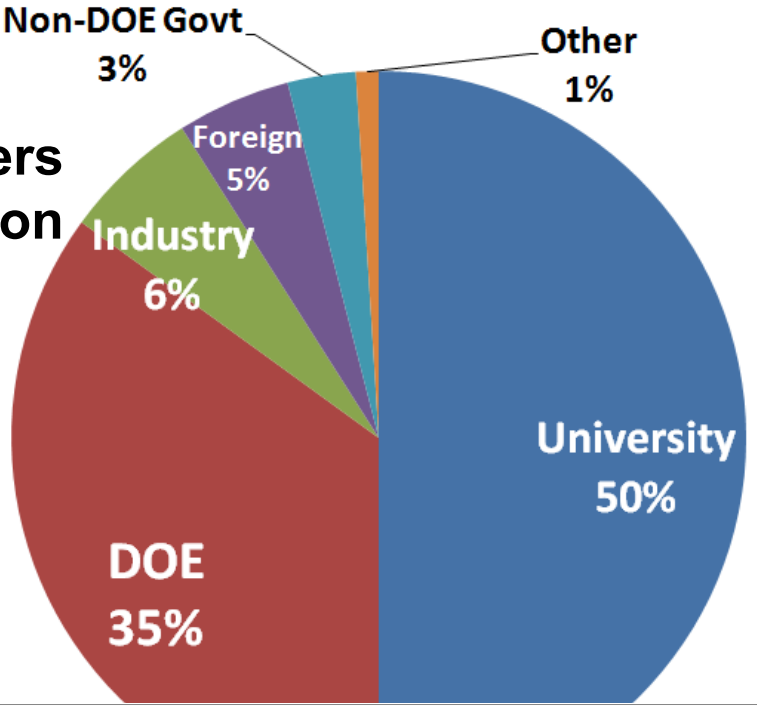

Usage by Science Area

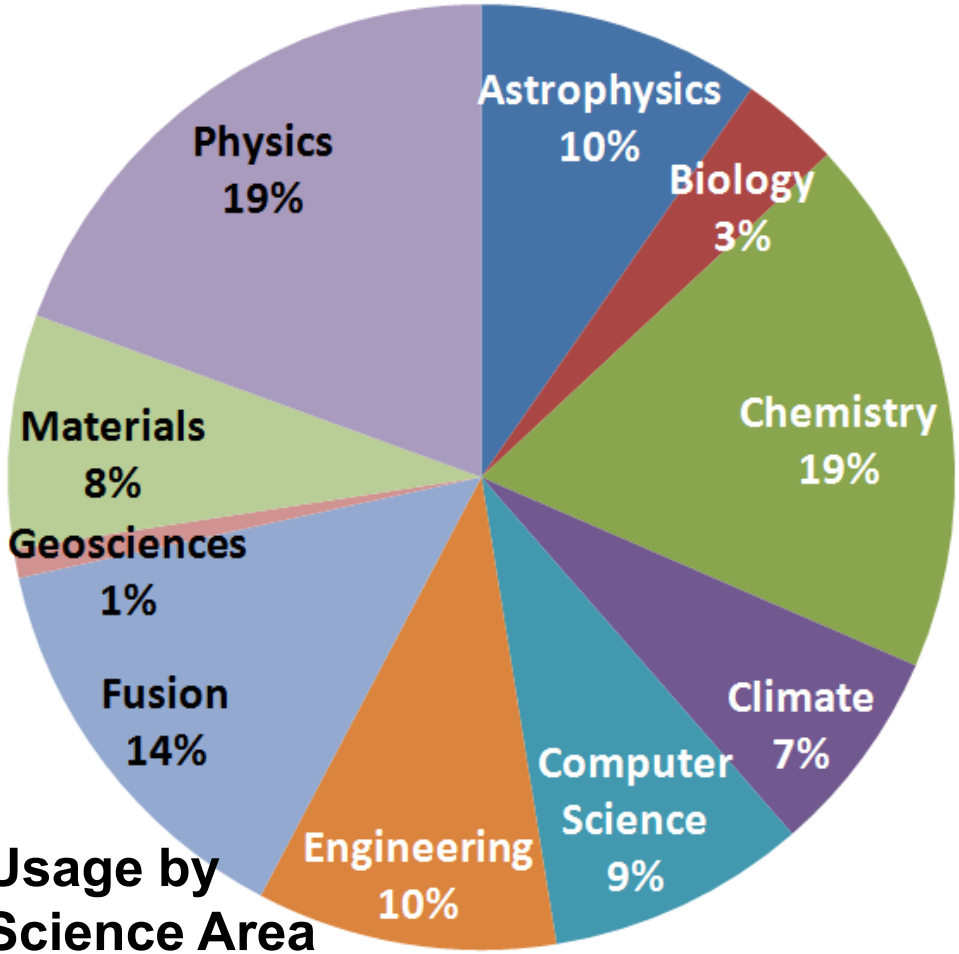

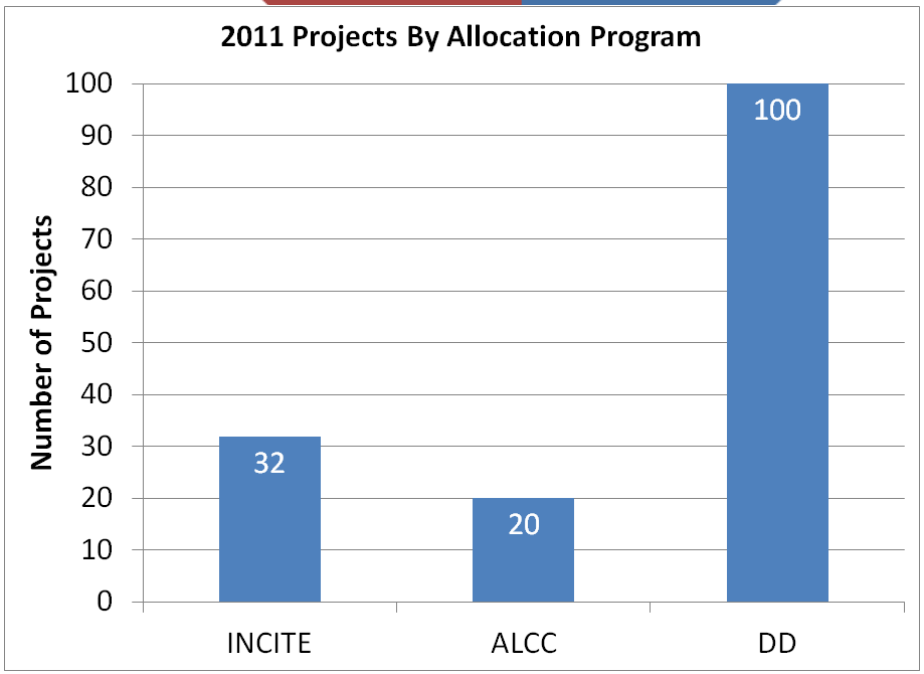

Provide access to LCFs:
More than 2.7 billion core hours awarded in 2012

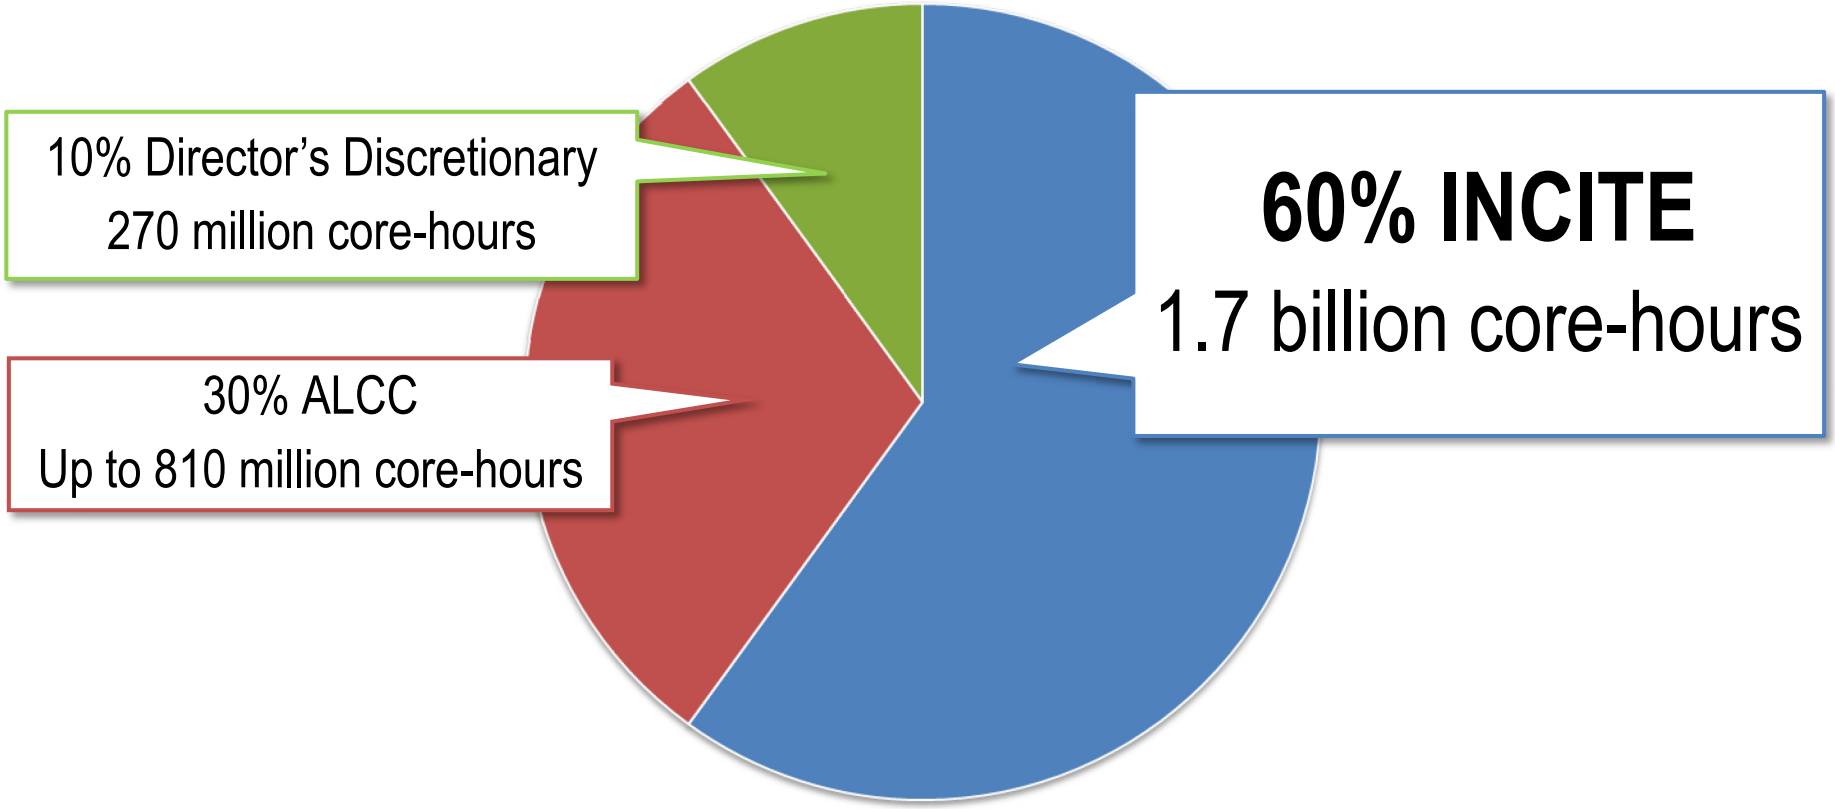

OLCF allocation programs
Selecting applications of national importance

	60% INCITE		30% ALCC		10% Director's Discretionary	
Mission	High-risk, high-payoff science that requires LCF-scale resources		High-risk, high-payoff science aligned with DOE mission		Strategic LCF goals	
Call	1x/year – (Closes June)		1x/year – (Closes February)		Rolling	
Duration	1-3 years, yearly renewal		1 year		3m,6m,1 year	
Typical Size	30 - 40 projects	20M - 100M core-hours/yr.	5 - 10 projects	1M – 75M core-hours/yr.	100s of projects	10K – 1M core-hours
Review Process	Scientific Peer-Review	Computational Readiness	Scientific Peer-Review	Computational Readiness	Strategic impact and feasibility	
Managed by	INCITE management committee (ALCF & OLCF)		DOE Office of Science		OLCF management	
Availability	Open to all scientific researchers and organizations including industry					

INCITE: Innovative and Novel Computational Impact on Theory and Experiment

INCITE promotes transformational advances in science and technology through large allocations of computer time, supporting resources, and data storage at the Argonne and Oak Ridge Leadership Computing Facilities (LCFs) for computationally intensive, large-scale research projects.

Is INCITE right for me?

INCITE seeks computationally intensive, large-scale research projects with the potential to significantly advance key areas in science and engineering.

1 Impact criterion

High-impact science and engineering

2 Computational leadership criterion

Computationally intensive runs that cannot be done anywhere else

3 Eligibility criterion

- INCITE grants allocations **regardless of funding source** (ex. DOE, NSF, private, etc)
- **Non-US-based researchers** are welcome to apply

Call for 2013 proposals opens in Spring 2012

The INCITE program seeks proposals for high-impact science and technology research challenges that require the power of the leadership-class systems

2012 INCITE allocations

28 new projects, 32 renewals
Acceptance rate: 33% for new proposals, 91% for renewals

Distribution of INCITE time by science domain

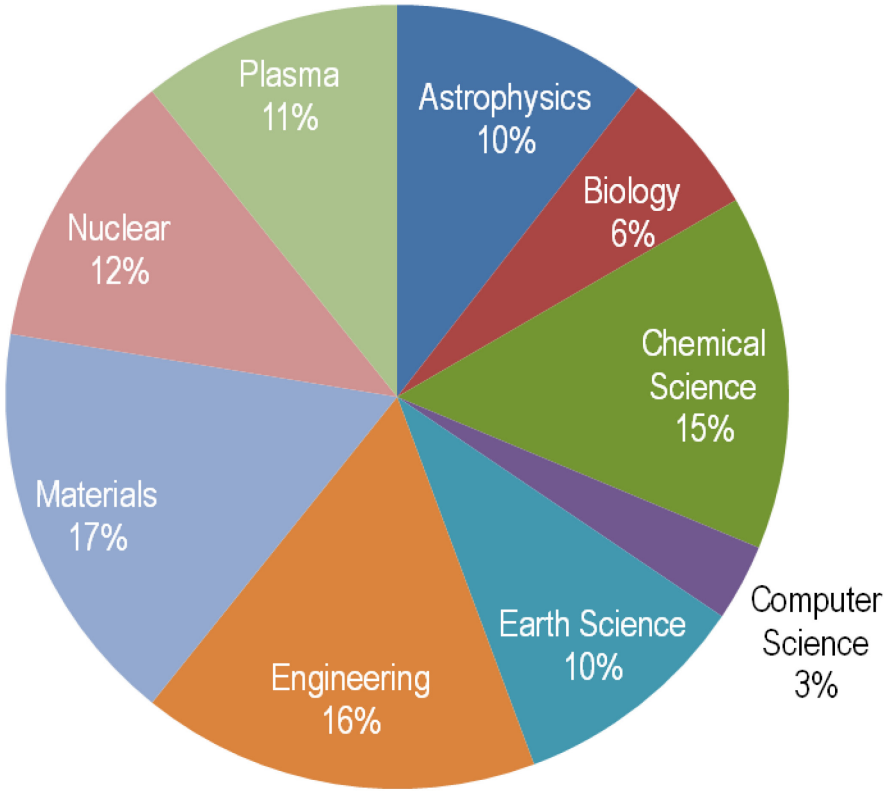

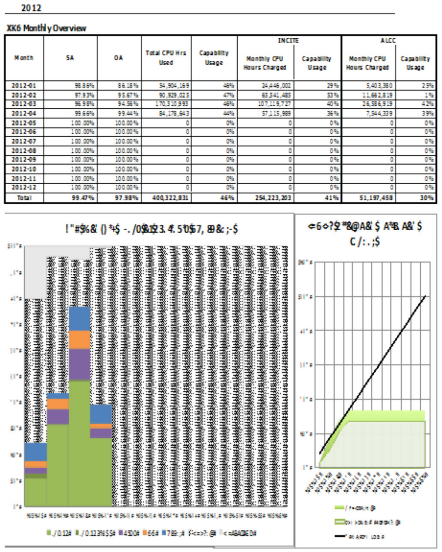

2012

XK6 Monthly Overview

Month	SA	OA	Total CPU Hrs Used	Capability Usage	INCITE Monthly CPU Hours Charged	INCITE Capability Usage	ALCC Monthly CPU Hours Charged	ALCC Capability Usage
2012-01	98.84%	86.18%	54,904,169	48%	24,448,000	28%	5,403,380	29%
2012-02	97.93%	93.67%	90,829,023	47%	60,241,483	33%	11,662,838	1%
2012-03	96.99%	94.36%	170,312,993	48%	107,118,727	42%	24,384,638	42%
2012-04	99.64%	99.44%	84,178,843	44%	37,215,988	38%	7,344,329	39%
2012-05	100.00%	100.00%	0	0%	0	0%	0	0%
2012-06	100.00%	100.00%	0	0%	0	0%	0	0%
2012-07	100.00%	100.00%	0	0%	0	0%	0	0%
2012-08	100.00%	100.00%	0	0%	0	0%	0	0%
2012-09	100.00%	100.00%	0	0%	0	0%	0	0%
2012-10	100.00%	100.00%	0	0%	0	0%	0	0%
2012-11	100.00%	100.00%	0	0%	0	0%	0	0%
2012-12	100.00%	100.00%	0	0%	0	0%	0	0%
Total	99.47%	97.98%	400,322,831	46%	256,223,203	41%	51,197,458	30%

Awarded 1.67 billion hours for CY 2012

Leadership Metric and Scheduling Policy

As a DOE Leadership Computing Facility, the OLCF has a mandate to be used for large, *leadership-class* (aka *capability*) jobs.

To that end, the OLCF implements queue policies that enable large jobs to run in a timely fashion.

- Basic queue priority is set by the time a job has been waiting relative to other jobs in the queue.
- However, we use several factors to modify the *apparent* time a job has been waiting. These factors include:
  - The job's processor core request size.
  - The queue to which the job is submitted.
  - The 8-week history of usage for the project associated with the job.
  - The 8-week history of usage for the user associated with the job.

Leadership Usage Metric:

35% of the CPU time used on the system will be accumulated by jobs using 20% or more of the available processors (60,000 cores)

OLCF Scheduling Policy

Bin	Min Cores	Max Cores	Max Walltime (Hours)	Aging Boost (Days)
1	180,000	- - - - -	24	15
2	60,000	179,999	24	5
3	5,008	59,999	12	0
4	2,004	5,007	6	0
5	1	2,003	2	0

Bin 2 is the leadership mark.

Leadership Metric Status
April 2012

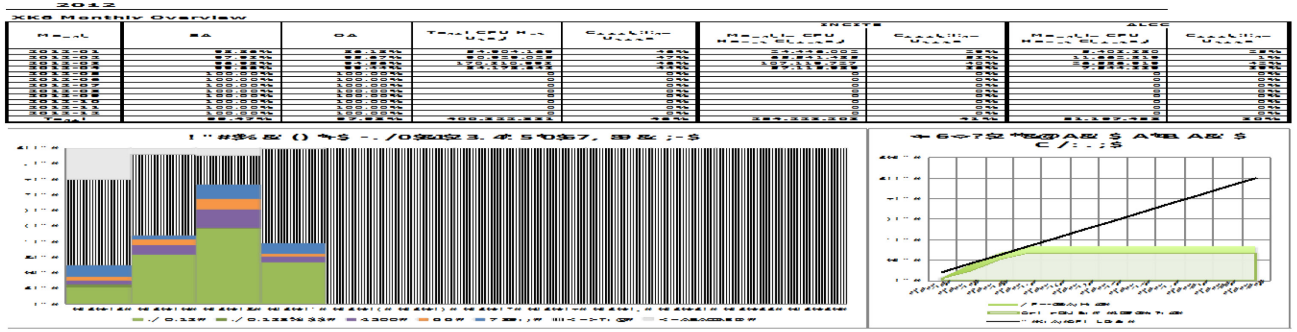

OLCF Allocation Overuse Policy

Projects that overrun their allocation are still allowed to run on LCF systems, although at a reduced priority.

- For projects that have used between 100% and 125% of their allocations, the following rules apply:
  - Jobs have their priority reduced by 30 days.
- For projects that have used greater than 125% of their allocation, the following rules apply:
  - Jobs have their priority reduced by 365 days.

To view the entire scheduling policy please see:
http://www.olcf.ornl.gov/kb_articles/scheduling-policy-olcf/

OLCF Pull Back Policy for INCITE Projects

On May 1 of the current INCITE calendar year:

- if usage is less than 15% remove up to 15% of the unused balance
- if usage is less than 10% remove up to 30% of the unused balance

On September 1 of the current INCITE calendar year:

- if usage is less than 50% remove up to 33% of the unused balance
- if usage is less than 33% remove up to 50% of the unused balance
- if usage is less than 10% remove up to 75% of the unused balance

OLCF Quarterly Reports

Principal Investigators of current OLCF projects must submit quarterly progress reports. The quarterly reports are essential as the OLCF must track and report the use of the center's resources.

The OLCF (and DOE Leadership Computing Facilities in general) imposes the following penalties for late submission:

- One Month Late:
  - Job submissions against offending project will be suspended.
- Three Months Late:
  - Login privileges will be suspended for all OLCF resources for all users associated with offending project.

OLCF Storage Policy

Users are provided with several storage areas, each of which serve different purposes and have different quotas, backup policies, and retention times.

Area	Nickname	Type	Quota	Backups	Purge	Rentention
User Home	-	NFS	5 GB	Yes	Not purged	1 month after project deactivation
User Work	“Spider”	Lustre	None	**No**	Files > 14 days are subject to deletion	Not retained
User Archive	“HPSS”	HPSS	2 TB or 2,000 Files	Yes	Not purged	3 months after project deactivation
Project Home	-	NFS	50GB	Yes	Not purged	1 month after project deactivation
Project Work	“Spider”	Lustre	1 TB	**No**	Not purged	Not retained
Project Archive	“HPSS”	HPSS	45TB (or 4,500 files)	Yes	Not purged	3 months after project deactivation

OLCF Publications

	2011
Number of 2011 publications reported to OLCF or identified by OLCF	670
Number of refereed publications reportable within OA guidance	300

Selected High-Impact

- 1 *Science*
- 25 *Phys. Rev. Lett*
- 4 *Geophy. Rev. Lett*
- 3 *NanoLetters*

OLCF Acknowledgement:

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Procedure implemented for collecting publications:

- Collect publications from users' quarterly reports
- Manual search to identify non-reported publications

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

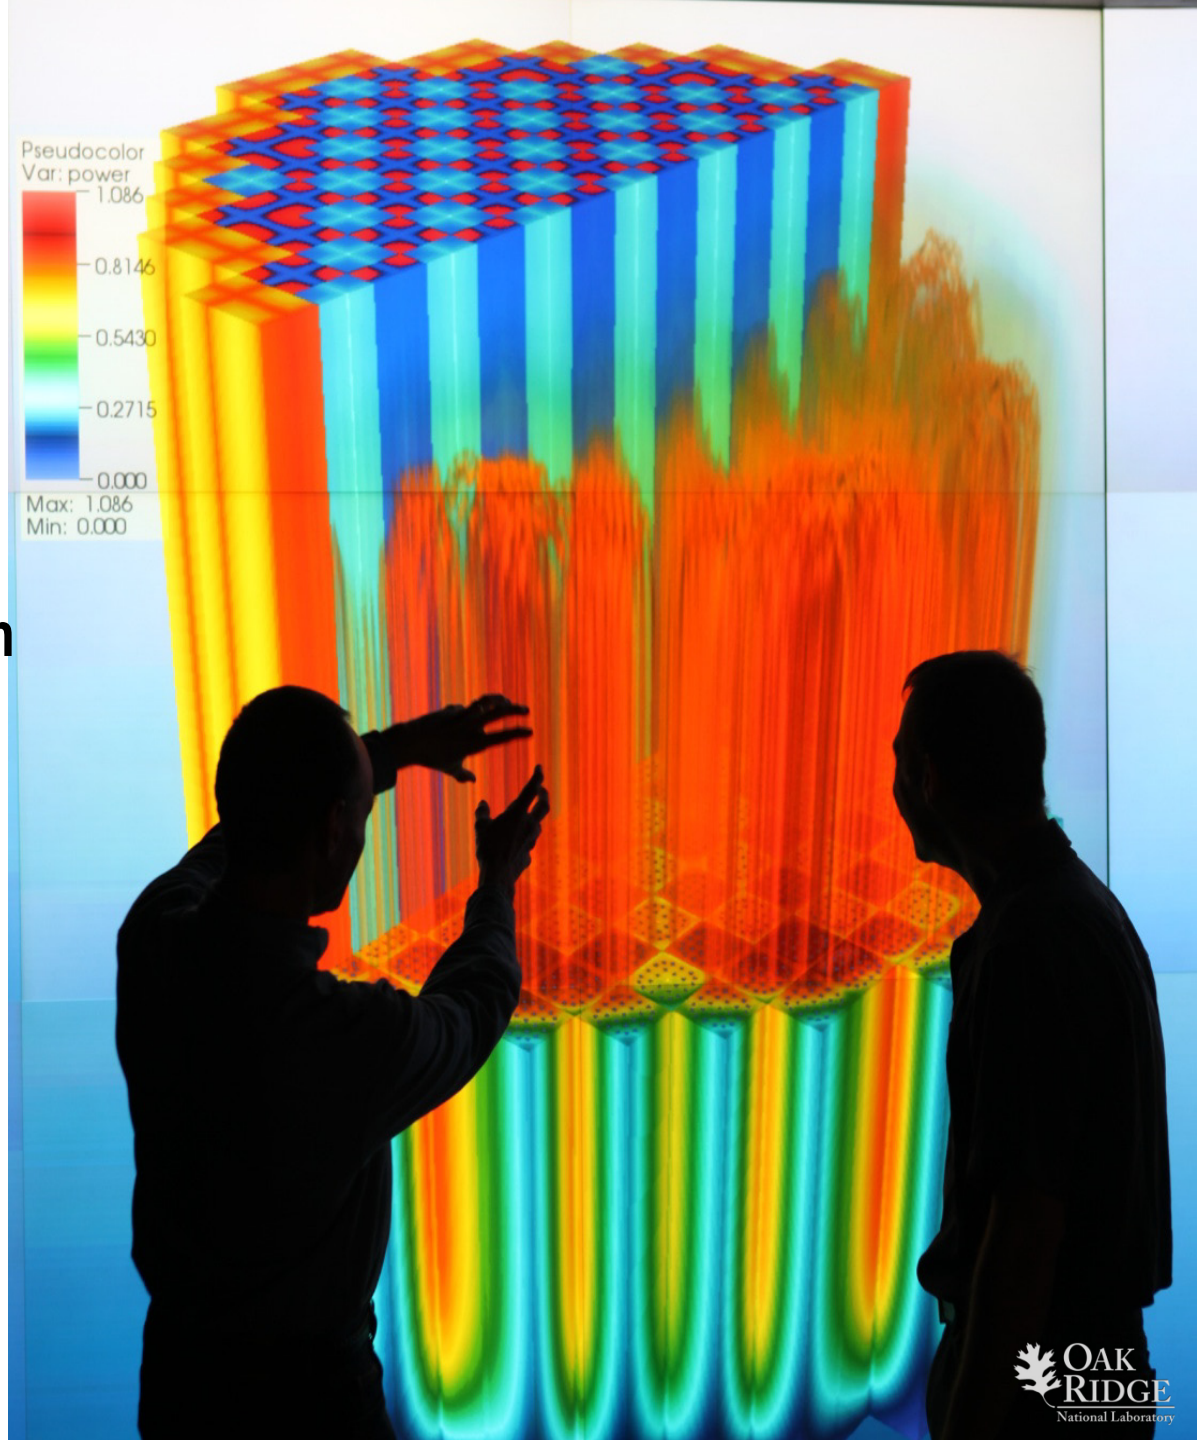

Cray XK6 Compute Node

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)

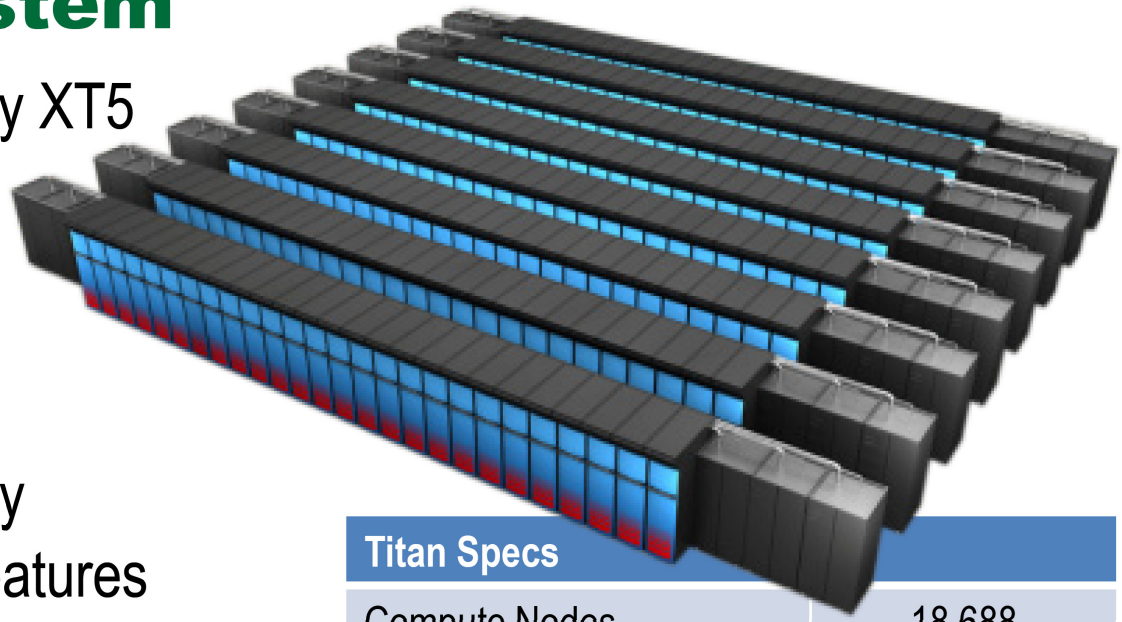

Titan Specs	
Compute Nodes	18,688
Login & I/O Nodes	512
Memory per node	32 GB + 6 GB
NVIDIA "Fermi" (2011)	665 GFlops
# of Fermi chips	960
NVIDIA "Kepler" (2012)	>1 TFlops
Opteron	2.2 GHz
Opteron performance	141 Gflops
Total Opteron Flops	2.6 Pflops
Disk Bandwidth	~ 1 TB/s

Final Upgrade Phase

Q4 2012

- The final upgrade from Jaguar to Titan will begin; the OLCF will outfit 152 of the 200 cabinets in Jaguar with accelerators
- Users will be able to run on the 46 non-upgraded opteron only cabinets, except during periods when full system reservations are needed
  - 2 cabinets are for IO (152 + 46 + 2 = 200)
- Much like last fall, the OLCF will perform rolling upgrades to minimize user impact
- Therefore, users should anticipate having access to ~5/8 of Jaguar during Q4 except during periods when full system reservations are needed

Final Upgrade Phase

Q1 2013

- The system will undergo stability and acceptance testing
- Users will be able to run on the 46 non-upgraded opteron only cabinets, except during periods when full system reservations are needed and during the final acceptance test phase
- Users will have some access to the 152 upgraded cabinets except during periods when full system reservations are needed and during the final acceptance test phase
- Targeted completion of the project is Q1, 2013
- At the completion of the project, the name will officially change from Jaguar to Titan
- Allocation/charging algorithm will change in 2013
  - See the INCITE Call for Proposals

NCCS Division Org Chart


Jim Hack
Division Director
Buddy Bland
OLCF Project Director






Jim Rogers
Director of Operations
Jack Wells
Director of Science
Kathlyn Boudwin
Deputy Project Director

Al Geist
Chief Technology Officer
Julia White
INCITE Program Manager
Suzy Tichenor
Industrial Outreach

Application Performance Tools Group
Rich Graham
FTE: 4.7

Scientific Computing Group
Bronson Messer (A)
FTE: 16.7

User Assistance & Outreach Group
Ashley Barker
FTE: 16.6

HPC Operations Group
Ann Baker
FTE: 17.6

Technology Integration Group
Galen Shipman
FTE: 9.7





Total OLCF Supported FTEs: 72

Questions & Discussion