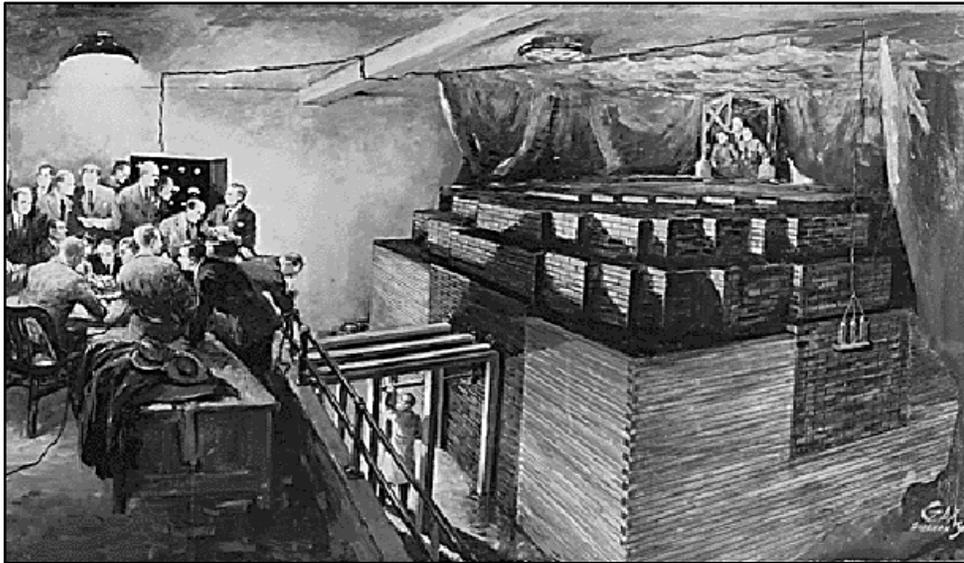


Allocation Programs



Katherine Riley
ALCF Director of Science

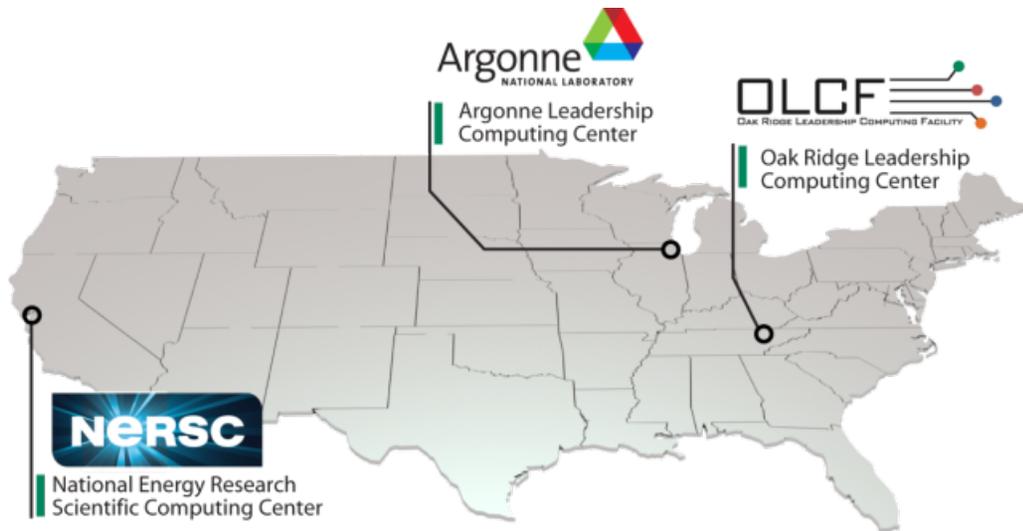
Argonne National Laboratory
24 May 2016



Advanced
Photon Source

Theory and Computing
Sciences Building

DOE'S OFFICE OF SCIENCE COMPUTATION USER FACILITIES



- DOE is leader in open High-Performance Computing
- Provide the world's most powerful computational tools for open science
- Access is free to researchers who publish
- Boost US competitiveness
- Attract the best and brightest researchers



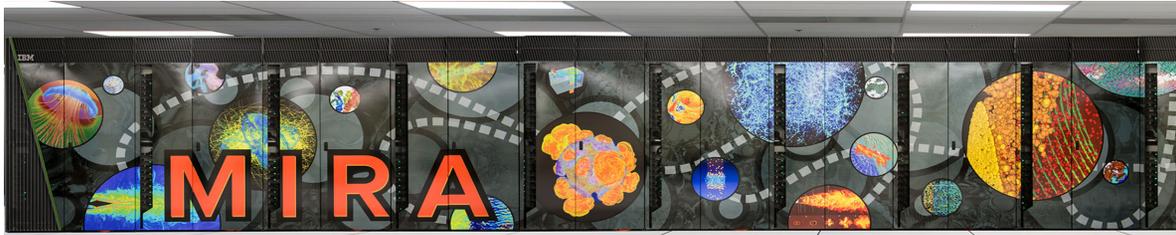
NERSC
Edison is 2.57 PF



ALCF
Mira is 10 PF



OLCF
Titan is 27 PF



The Largest Concentration of Scientific Computing in One Organization on the Planet



LEADERSHIP COMPUTING MISSION

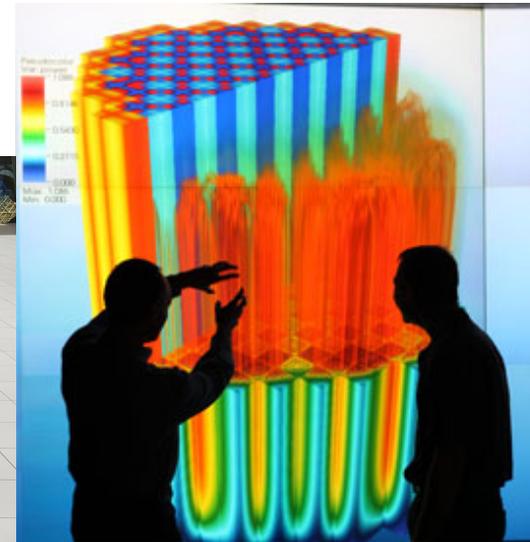
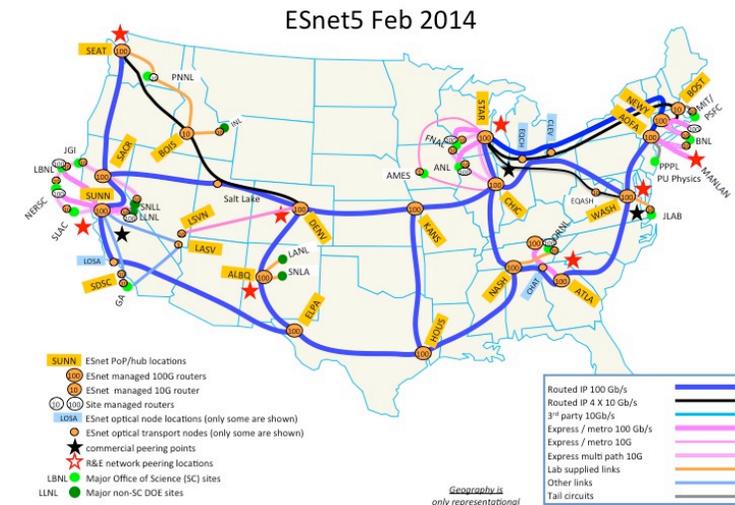


Leadership Class Resources

- Accelerate major scientific discoveries and engineering breakthroughs for humanity
- Dedicated to open science
- Two diverse architectures
 - ALCF & OLCF
 - Two of the worlds most powerful computers
- Supported by DOE's Advanced Scientific Computing Research Program (ASCR)

LEADERSHIP IS NOT JUST BIG COMPUTE

- Software Stack
- Large-scale data, data analysis and visualization
- World leading network interconnecting facilities
- Storage - Hundreds of Petabytes Storage Systems
- Science driven software architecture, performance, and facility design
- DOE invests more than \$1Billion/yr in the computing capabilities at the laboratories
- **People**
 - Computational Scientists
 - Performance Engineers
 - Visualization & Data Analysis
 - Data Science



LCF Allocation Programs	INCITE	60%	ALCC	30%	Director's Discretionary	10%
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) <i>2017 Call Open</i>		1x/year – Closes February		Rolling	
Duration	1-3 years, yearly renewal		1 year		3m,6m,1 year	
Typical Size	30 - 40 projects	75M - 500M core-hours/yr.	10-20 projects	10M – 300+M core-hours/yr.	~100 of projects	.5M – 10M core-hours
Total Hours	~5 billion core-hours (~3.5B ALCF)		~2.5 billion core-hours (~1.75 ALCF)		~590 million ALCF	
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		LCF management	
Readiness	High		Medium to High		Low to High	
Availability	Open to all scientific researchers and organizations Capability > 131,072 cores (16.7% of Mira)					

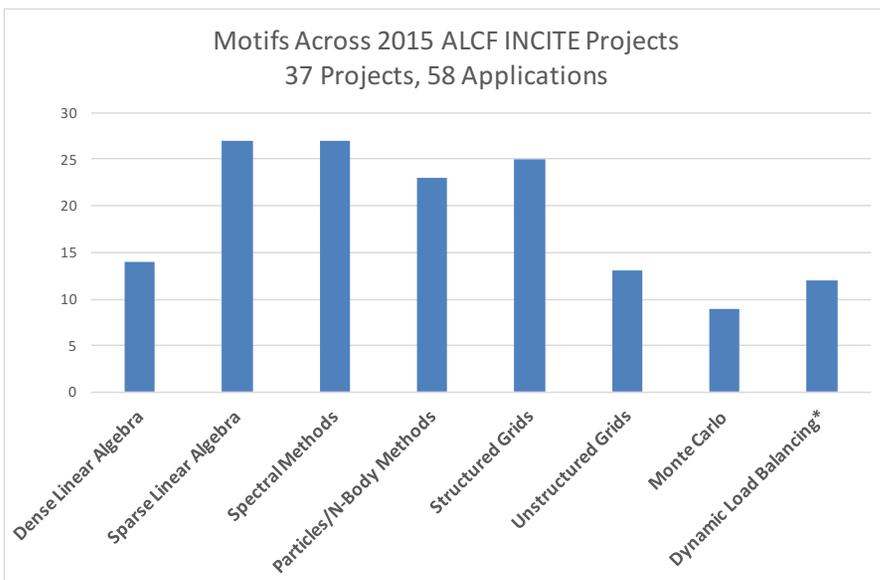
ALLOCATION PROGRAMS AT THE LCFS

	60%	30%	10%
	INCITE		Director's Discretionary
Mission	High-risk, high-payoff science requires LCF-scale resources	Aligned	Strategic LCF goals
Call	1x/year – (Closes June) <i>2016 Call Opens in</i>)	Rolling
Duration	1-3 years, yearly renewal		3m, 6m, 1 year
Typical Size	30 - 40 projects 75M - 50M core-hours/yr		~100 of projects .5M – 10M core-hours
Review Process	Scientific Peer-Review Computer Readiness		Strategic impact and feasibility
Managed By	INCITE management committee (ALCF & OLCF)		LCF management
Readiness	High		Low to High
Availability	Open to all scientific researchers and organizations Capability > 131,072 cores (16.7% of Mira)		

- Development of INCITE & ALCC proposals (short term)
- Development of applications for extreme scale (longer term)
- Computer Science
- Training
- Proof of concept
- <http://www.alcf.anl.gov/getting-started/apply-for-dd>

ALCF WORKLOAD: DIVERSE IN METHOD AND DOMAIN

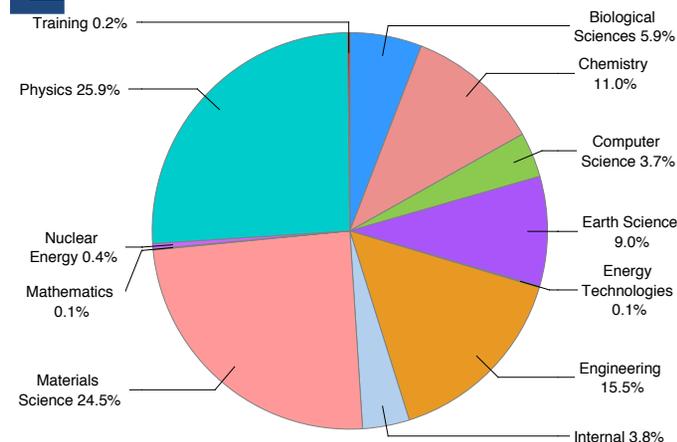
Motifs Across 2015 ALCF INCITE Projects
37 Projects, 58 Applications



Top Methods in 2015 by INCITE Allocated Time

- Climate
- DFT
- MD (Reactive)
- Engineering – fluid flow, combustion
- Quantum Monte Carlo
- QCD

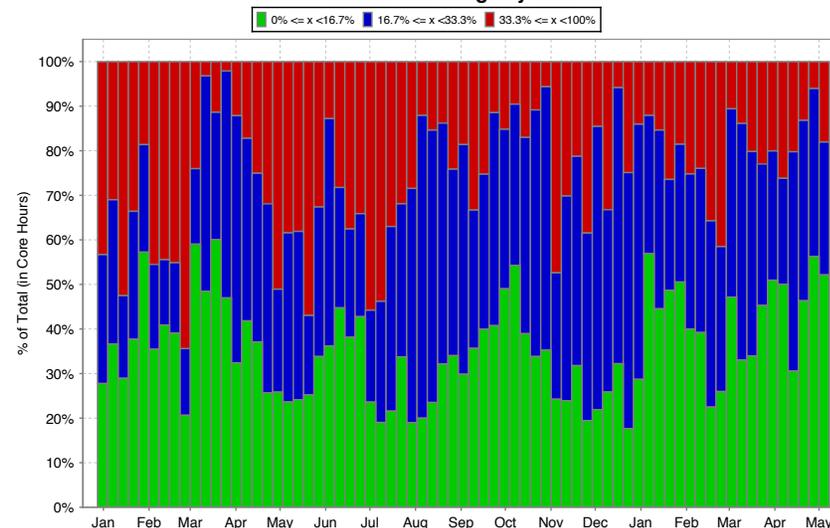
Mira Allocations Active Between 2015-01-01 and 2016-05-02
324 Projects, 10.12B core-hours



Generated on 2016-05-02

Start	End	Allocated Time
2015-01-01	2016-05-01	7.8b
0% <= X < 16.7%		2.8B
16.7% <= X < 33.3%		3.0B
33.3% <= X <= 100%		2.0B
Total		7.8b

Mira Overall Job Usage by Size



ALCF RESOURCES

Mira - compute

- 10 PF IBM BG/Q
- 48K nodes/786K cores
- 786 TB memory
- 5D Torus interconnect
- 26 PB GPFS, 400 GB/s



INCITE
ALCC
DD

Cooley – data analytics

- 223 TF
- 126 nodes/1512 Xeon cores/126 Tesla K80 GPUs
- 384 TB (CPU)/3 TB (GPU) memory
- FDR InfiniBand interconnect
- Connected to Mira file systems



Cetus – app T&D

- 840 TF IBM BG/Q
- 4K nodes/64K cores – **8.3% of Mira**
- 64 TB memory
- 5D Torus interconnect
- Connected to Mira file systems



DD

Vesta – system SW T&D

- 420 TF IBM BG/Q
- 2K nodes/32K cores – **4.2% of Mira**
- 32 TB memory
- 5D Torus interconnect
- 1 PB GPFS



DD

TWO NEW ALCF SYSTEMS THETA AND AURORA

- Intel Xeon Phi compute architecture
- Deep memory architecture – very fast memory, slower capacity memory, burst buffers



THE RIGHT PATH FOR OUR USERS

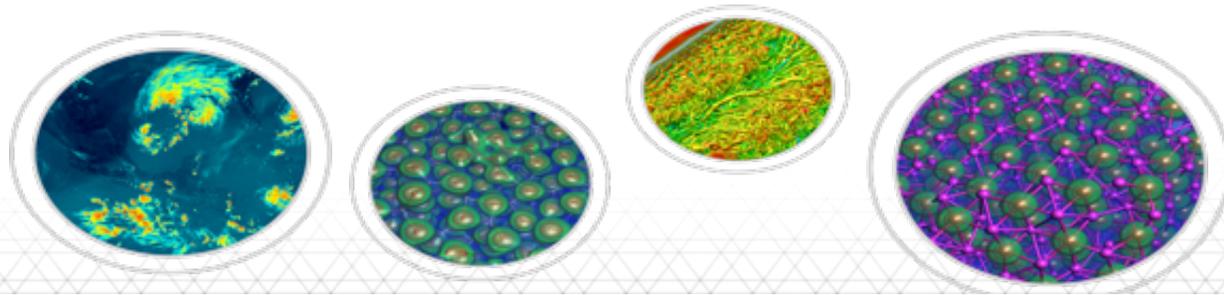
- Many core evolution
- Easy to port codes
- Well-balanced between compute, memory, network, and storage
- Robust and well-known Cray user environment combined with Intel innovations

THETA: COMING TO ALCF IN 2016



Stepping stone to Aurora

- Over 8.5 PF peak performance
- 2nd Generation Intel® Xeon Phi™ processors (KNL)
 - 192GB DDR4 memory and up to 16GB HBM on each node
 - 128GB SSD on each node
- Cray Aries high speed interconnect in dragonfly topology
- Lustre file system
- Cray software stack



TOWARD EXASCALE: ALCF'S NEW AURORA SYSTEM

A 3-way Collaboration That Accelerates Discovery And Innovation

180+ PFLOPS

(can increase to 450+ PF)

2018 delivery

**13X higher APPLICATION
performance**

6.5X+ more energy efficient

- Intel HPC Scalable System Framework
- 3rd Generation Intel® Xeon Phi™ processor
- 50,000+ nodes
- 7+ PB total system memory
- 2nd Generation Intel® Omni-Path Architecture silicon photonics in a dragonfly topology
- 150+ PB Lustre file system, 1+ TB/s I/O



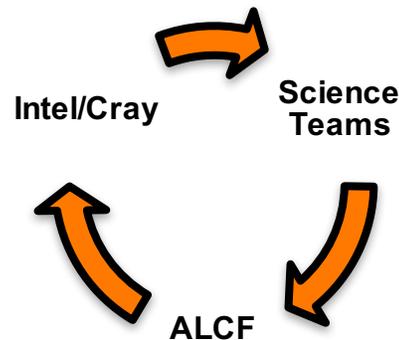
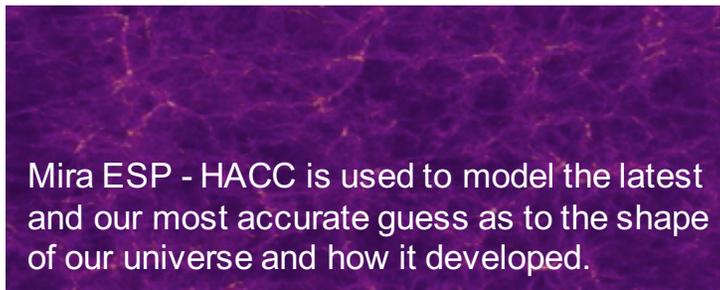
Prime Contractor



Integrator

ALCF EARLY SCIENCE PROGRAM

GOAL: SCIENCE ON DAY ONE



- Based on successful Mira ESP pioneering program
- Theta ESP (2015 – 2017)
- Aurora ESP (2016 – 2019)



Call for Aurora ESP proposals 3Q 2016

ALCF DATA SCIENCE PROGRAM

- “Big Data” science problems that require the scale and performance of leadership computing resource
- Wide variety of application domains that span computational, experimental and observational sciences
- Focus on data science techniques Projects will target science and software technology scaling for data science
- Proposal Deadline: May 31, 2016 (Expected yearly call for proposals)
<https://www.alcf.anl.gov/alcf-data-science-program>

SUPPORT

- Funded postdoctoral appointee
- ALCF staff support

COMPUTE RESOURCES

- Theta
- Big-Data Cluster
- Cooley and Mira
- Aurora (Future)



QUESTIONS