Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility

Accelerating Discovery and Innovation

Katherine M Riley
Director of Science, Argonne Leadership Computing Facility

www.anl.gov

Supercomputing Resources

As a National User Facility we are focused on open science that cannot be readily tackled elsewhere.

Our supercomputers or allocations are 10 to 100 times more powerful than systems typically used for scientific research.

ALCF Resources - Theta

Many core architecture – Vendor: Intel (Prime) / Cray (Integrator)

Theta Specs	
Peak Performance	11.69 PF
Compute Nodes	4,392
Node	Intel “Knights Landing” (KNL) Xeon Phi 7230
Cores/node	64
Total Cores	281,088
Memory per node	16 GiB MCDRAM + 192 GiB DDR4 + 128 GiB SSD
Hardware threads per node	256
Total System Memory	72 TiB MCDRAM + 843 TiB DDR4 + 561 TiB SSD
Interconnect	Aries (Dragonfly)
File System Performance	200 PB (Grand+Eagle, Lustre)

Theta System Summary

Description	Theta	ThetaGPU
Total # Compute Nodes	4,392	24
Total # Cores	281,088	3,072
Total # HW Threads	1,124,352	
Total # GPUs		192 (8/node)
Total Memory	70 TB MCDRAM 843 TB DDR4	24 TB DDR 7,680 TB GPU
Interconnect Topology	Cray Aries Dragonfly	HDR200
Data Store (PB) (raw)	18 PB, 240 GB/s	
Total Peak DP FLOPS (PF)	11.7 PF	3.9 PF
Power (peak)	2.7 MW	
Power (average)	1.7 MW	
Total Compute Racks	24	7
Total Floor Space	1,000 sq. ft.	

- ThetaGPU is a COVID-19 focused extension to Theta
- https://www.alcf.anl.gov/support-center/theta/theta-thetagpu-overview

ALCF Resources - Polaris

Anticipated Polaris Specs	
Planned Installation / Production	Q2CY2021 / Q3CY2021
System Peak	35-45 PF
Peak Power	<2 MW
Total System Memory	>250 TB
System Memory Type	DDR, HBM
Node Performance	>70 TF
Node Processors	1 CPU; 4 GPUs
System Size (nodes)	>500
Node-to-Node Interconnect	200 Gb

Polar is expected to be in the INCITE 2022 allocation pool.

Polaris is planned to be a hybrid CPU/GPU machine to provide increased readiness for Aurora.

ALCF Resources - Polaris

- Programming Models
  - OpenMP 4.5/5
  - SYCL
  - Kokkos
  - Raja
  - HiP
- Performance / Debugging
  - GPU tools
  - PAPI, TAU, HPCToolkit
  - DDT
- Frameworks
  - Python / Numba
  - Tensorflow
  - Pytorch

Aurora : Nations First ≥1 Exaflop Supercomputer

Delivery in 2022

Intel Xeon scalable processors

X^e arch based GP-GPUs

Supporting the future of science

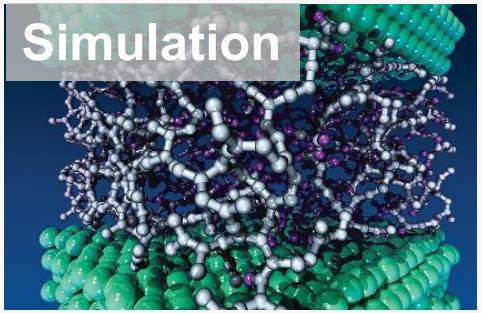

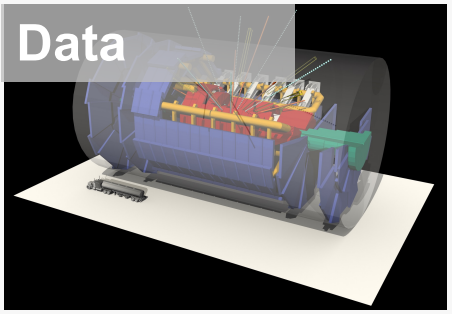

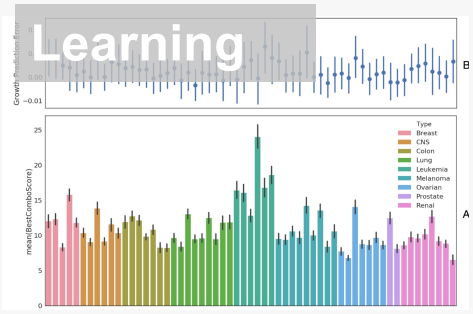

Aurora high-level configuration (public)

System Spec	Aurora	System Spec	
System Performance	≥1EF DP sustained	**Node Performance**	> 130 TF
Compute Node	2 Intel Xeon scalable processors (Sapphire Rapids) 6 Intel Xᵉ arch based GP-GPUs (Ponte Vecchio)	**Number of nodes**	> 9,000
GPU Architecture	Xᵉ arch based GPU (Ponte Vecchio) Tile based, chiplets, HBM stack, Foveros 3D integration	**Number of Cabinets**	> 100
CPU-GPU interconnect	PCIe	**Node Memory Architecture**	Unified memory architecture, RAMBO
Aggregate System Memory	>10 PB	**Peak Power**	≤ 60 MW
System Interconnect	HPE Slingshot 11 Dragonfly topology with adaptive routing	**Programming languages and models**	**Fortran, C, C++ OpenMP 5.x (Intel, Cray, and possibly LLVM compilers), UPC (Cray), Coarray Fortran (Intel), Data Parallel C++ (Intel and LLVM compilers), OpenSHMEM, Python, Numba, MPI, OpenCL**
Network Switch	25.6 Tb/s per switch, from 64 - 200 Gbs ports (25GB/s per direction)	Compilers	Intel, LLVM, GCC
High-Performance Storage	≥230 PB, ≥25 TB/s (DAOS)	Programming tools	Open|Speedshop, TAU, HPCToolkit, Score-P, Darshan, Intel Trace Analyzer and Collector, Intel Vtune, Advisor, and Inspector, PAPI, GNU gprof
Programming Models	Intel oneAPI, OpenMP, DPC++/SYCL	Debugging and Correctness Tools	Stack Trace Analysis Tool, gdb, Cray Abnormal Termination Processing
Software stack	HPE Cray XE software stack + Intel enhancements + Data and Learning	Math Libraries	Intel MKL, Intel MKL-DNN, ScaLAPACK
Platform	HPE Cray XE	GUI and Viz APIs, I/O Libraries	X11, Motif, QT, NetCDF, Parallel NetCDF, HDF5
		Frameworks	TensorFlow, PyTorch, Scikit-learn, Spark Mllib, GraphX, Intel DAAL, Intel MKL-DNN

How Do Researchers Gain Access to ALCF?

We offer different pipelines based on your computational readiness. Apply to the allocation program that fits your needs.

- Getting Started (DD)
- Major Awards (INCITE, ALCC)
- Targeted Projects (ADSP, ESP)

Primary Allocation Programs for Access to the LCFs

Current distribution of allocatable hours

20% Director's Discretionary (Includes LCF strategic programs, ECP)

Argonne NATIONAL LABORATORY | OAK RIDGE National Laboratory

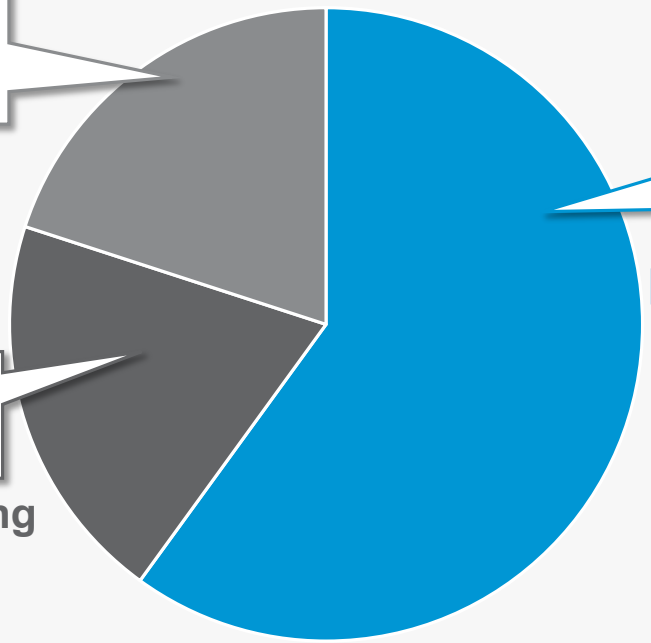

60% INCITE

Leadership-class computing

Argonne NATIONAL LABORATORY | OAK RIDGE National Laboratory

20% ASCR Leadership Computing Challenge

DOE/SC capability computing

LCF Allocation Programs	INCITE 60%	ALCC 20%	Director's Discretionary 20%
Mission	High-risk, high-payoff science that requires LCF-scale resources*	High-risk, high-payoff science aligned with DOE mission	50% Strategic LCF goals 50% ECP
Call	1x/year – Closes June *2022 Call Open*	1x/year – Closes February	Rolling
Duration	1-3 years, yearly renewal	1 year	3m,6m,1 year
Typical Size	**10-15 projects** **1-3M node-hours**	**5-15 projects** **0.5-2M node-hours**	**~100 of projects** **<0.5M node-hours**
Total Hours	**~17.8M Theta**	**~6M node-hours Theta**	**~6M node-hours Theta**
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 > 20% of resource***		

LCF Allocation Programs	INCITE 60%		ALCC 20%		Director's Discretionary 20%	
Mission	High-risk, high-payoff science that requires LCF-scale resources*		High-risk, high-payoff science aligned with DOE mission		50% Strategic LCF goals 50% ECP	
Call	1x/year – Closes June *2022 Call Open*		1x/year – Closes February		Rolling	
Duration	1-3 year		1 year		3m,6m,1 year	
Typical Size	**10-15 pro**		**5-15 projects**	**0.5-2M node-hours**	**~100 of projects**	**<0.5M node-hours**
Total Hours	**~17.8M Theta**		**~6M node-hours Theta**		**~6M node-hours Theta**	
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 > 20% of resource***					

New Early Career Track

Getting Started (DD)

Our Director's Discretionary (DD) allocation program provides researchers with small awards of computing time to "get started" on our computing resources while pursuing real scientific goals.

The DD allocation program allows users to prep their code so that it can take advantage of our massively parallel systems.

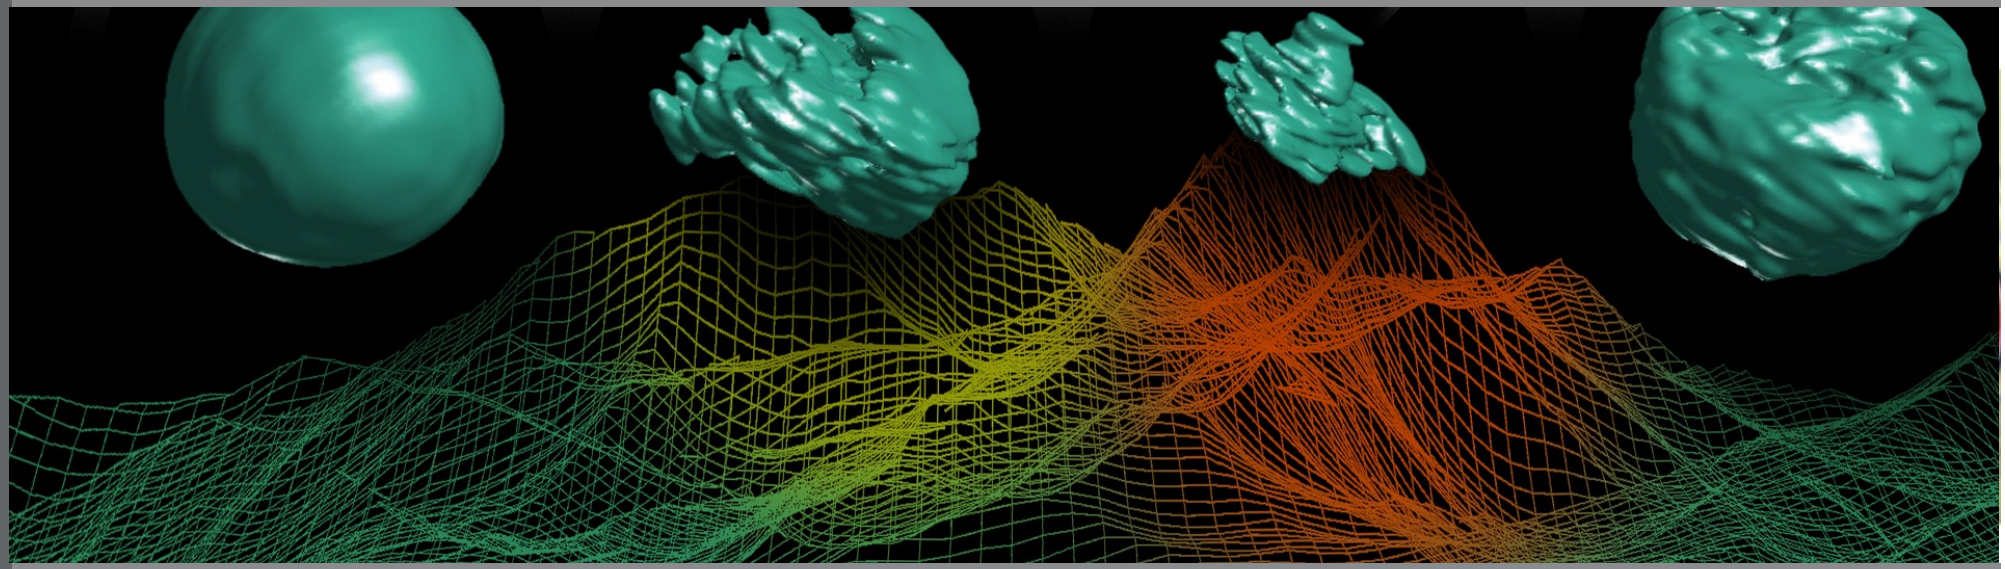

DD
Director's Discretionary

Purpose: A "first step" for projects working toward a major allocation

Eligibility: Available to all researchers in academia, industry, and other research institutions

Review Process: Projects must demonstrate a need for high-performance computing resources; reviewed by ALCF

Award Size: Low 10 thousand of node-hours

Award Duration: 3-6 months, renewable

Total percent of ALCF resources allocated: 20%

Award Cycle

Ongoing (available year-round)

ADSP
ALCF Data Science Program

Targeted at big data science problems, ADSP aims to explore and improve a variety of computational methods that will help enable data-driven discoveries across all scientific disciplines.

Eligibility: Available to researchers in academia, industry, and other research institutions

Review process: Applications undergo a review process to evaluate potential impact, data scale readiness, diversity of science domains and algorithms, and other criteria

Award size: ~Low hundred of thousand of node-hours

Award duration: 2 years

ESP
Early Science Program

As part of the process of bringing a new supercomputer into production, the ALCF hosts the Early Science Program (ESP) to ensure its next-generation systems are ready to hit the ground running.

The intent of the ESP is to use the critical pre-production time period to prepare key applications for the architecture and scale of a new supercomputer, and to solidify libraries and infrastructure to pave the way for other production applications to run on the system.

In addition to fostering application readiness, the ESP allows researchers to pursue innovative computational science projects not possible on today's leadership-class supercomputers.

Award Cycle
Determined by production timeline

Major Awards (INCITE, ALCC)

Our major allocations provides users with computationally intensive, large-scale research projects time on our machines.

The programs conduct a two-part review of all proposals: a peer review by a panel of experts and a computational readiness review.

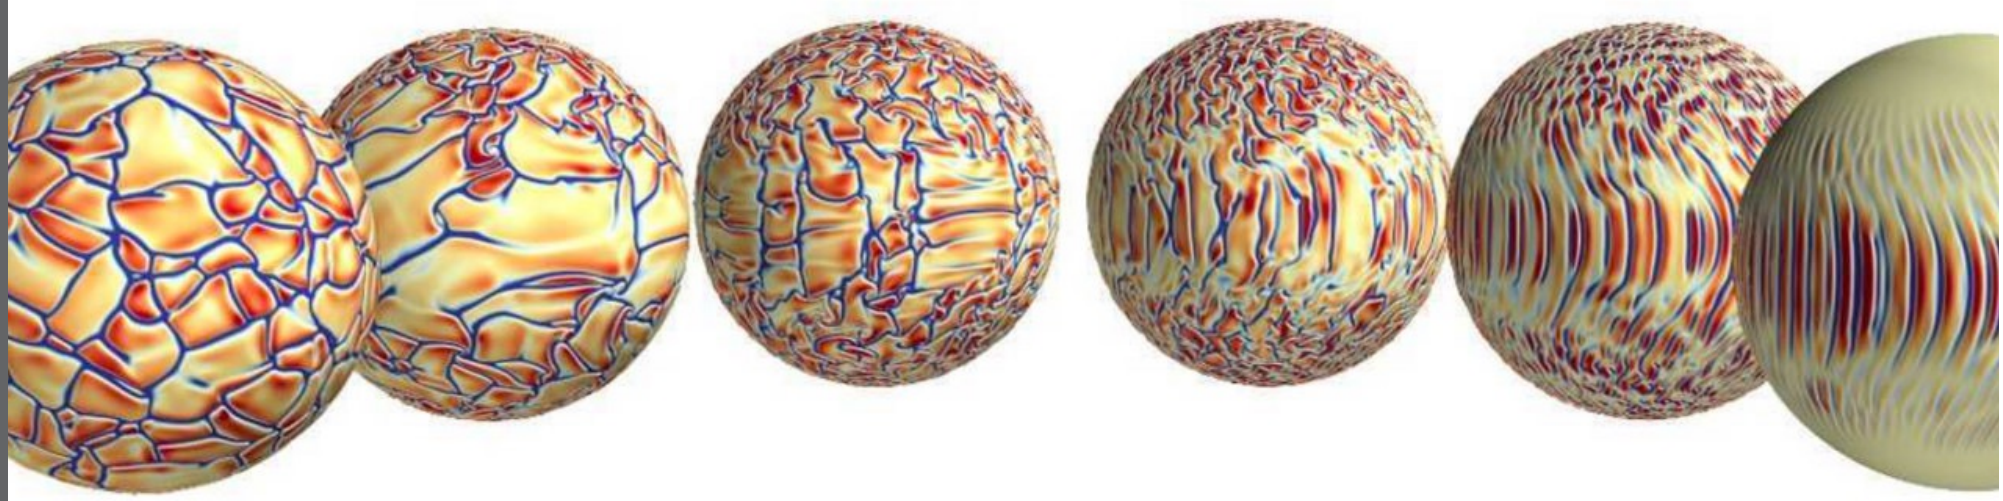

INCITE

Innovative & Novel Computational Impact on Theory and Experiment

The DOE's INCITE program provides allocations to computationally intensive, large-scale research projects that aim to address "grand challenges" in science and engineering.

Eligibility: Available to researchers in academia, industry, and other research institutions

Review process: INCITE program conducts a two-part review of all proposals including a peer review by an international panel of experts, and a computational-readiness review

Award size: ~1.0-2.5M node-hours

Award duration: 1-3 years, renewable

Total percent of ALCF resources allocated: 60%

Award Cycle

January 1 to December 31

Call Open
Due June 18th

INCITE criteria

Access on a competitive, merit-reviewed basis*

1	Merit criterion
	Research campaign with the potential for significant domain and/or community impact

2	Computational leadership criterion
	Computationally demanding runs that cannot be done anywhere else: capability, architectural needs

3	Eligibility criterion
	• Grant allocations regardless of funding source* • Non-US-based researchers are welcome to apply

*DOE High-End Computing Revitalization Act of 2004: Public Law 108-423

Twofold review process

	New proposal assessment	Renewal assessment
1 Peer review: INCITE panels	• Scientific and/or technical merit • Appropriateness of proposal method, milestones given • Team qualifications • Reasonableness of requested resources	• Change in scope • Met milestones • On track to meet future milestones • Scientific and/or technical merit
2 Computational readiness review: LCF centers	• Technical readiness • Appropriateness for requested resources	• Met technical/ computational milestones • On track to meet future milestones
Award Decisions	• INCITE Awards Committee comprised of LCF directors, INCITE program manager, LCF directors of science, sr. management	

Recent Trends in INCITE

Data, Learning and Nontraditional Uses of the Architecture

- In addition to traditional computationally intensive simulation campaigns, INCITE encourages Data and/or Learning projects with unique data requirements (e.g. large scale data analytics) or workflow needs that can only be enabled by the LCFs.
  - A “Learning” panel evaluated proposals that had significant machine / deep learning component to their campaign
  - When appropriate, these proposals were also assessed by their scientific discipline peers as well

Early Career Track in INCITE

For the INCITE 2022 Call for Proposals, INCITE is committing 10% of allocatable time to an Early Career Track in INCITE.

The goal of the early career track is to encourage the next generation of high-performance computing researchers.

Researchers within 10 years from earning their PhD (on or after December 31st 2011) may choose to apply.

Projects will go through the regular INCITE Computational Readiness and Peer Review process, but the INCITE Management Committee will consider meritorious projects in the Early Career Track separately.

2020 and 2021 award statistics, by system

2020	Summit	Theta
Number of projects*	39	14
Average Project	482 K	1.41 M
Median Project	500 K	1.50 M
Total Awards (node-hrs in CY2020)	18.8 M	19.7 M

* Total of 47 INCITE projects (6 projects received time on both Theta and Summit)
* All reported in node-hours native to each resource.

2021	Summit	Theta
Number of projects*	41	16
Average Project	468 K	1.2 M
Median Project	470 K	1.2 M
Total Awards (node-hrs in CY2020)	19.2 M	18.5 M

* Total of 51 INCITE projects (6 projects received time on both Theta and Summit)
* All reported in node-hours native to each resource.

ALCC
ASCR Leadership Computing Challenge

The DOE's ALCC program allocates resources to projects directly related to the DOE's energy mission, as well as national emergencies, and for broadening the community of researchers capable of using leadership computing resources.

Eligibility: Available to researchers in academia, industry, and other research institutions

Review process: DOE peer reviews all proposals for scientific/technical merit; appropriateness of approach; and adequacy of personnel and proposed resources

Award size: ~1M node-hours

Award duration: 1 year

Total percent of ALCF resources allocated: 20%

Award Cycle
July 1 to June 30

LOIs often due in January with follow-up invitation to write a proposal

LCF Allocation Programs	INCITE 60%	ALCC 20%	Director's Discretionary 20%
Mission	High-risk, high-payoff science that requires LCF-scale resources*	High-risk, high-payoff science aligned with DOE mission	50% Strategic LCF goals 50% ECP
Call	1x/year – Closes June *2022 Call Open*	1x/year – Closes February	Rolling
Duration	1-3 years, yearly renewal	1 year	3m,6m,1 year
Typical Size	**10-15 projects** **1-3M node-hours**	**5-15 projects** **0.5-2M node-hours**	**~100 of projects** **<0.5M node-hours**
Total Hours	**~17.8M Theta**	**~6M node-hours Theta**	**~6M node-hours Theta**
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 > 20% of resource***		


Thank You!
Learn more at: alcf.anl.gov