



2023 INCITE Proposal Writing Webinar

Our presentation will begin at 2:00 pm EDT / 1:00 pm CDT

A few reminders:

- Please ensure your camera and microphone are both turned off.
- During the presentation, questions will be taken in the chat window. There will be a time for verbal questions near the end. You can temporarily unmute yourself either by pressing and holding the space bar while you speak.
- These slides and eventually a recording of this presentation can be downloaded from

<https://www.doeleadershipcomputing.org/proposal/informational-webinars/>



2023 INCITE Proposal Writing Webinar



Marta García
ALCF Catalyst Team
Argonne National Laboratory



Katherine Riley
INCITE Program Manager



Dmytro Bykov
OLCF Scientific Engagement Section
Oak Ridge National Laboratory

Objectives of INCITE Proposal Writing Webinar

- **Learn** to write a proposal to obtain significant allocations of compute time and resources at DOE's Leadership Computing Facilities
- **Describe** proposal sections, content, and tips for new and renewal authors.
 - Proposal narratives consist of up to 15 pages describing
 - Description of high-impact science challenge
 - Description of methods, codes, and dependencies
 - Demonstration that computational tasks cannot be done anywhere else
 - Description of team qualifications
- **Timeline**
 - **New** Proposals are due June 17, 2022
 - **Renewal** Proposals are due July 22, 2023 (only those with multi-year proposals)
 - PIs will be notified in November 2022
 - Awards begin on January 1, 2023

New and Renewal Proposals

New

Eligibility

- PI without current active proposal
- PI with current active proposal in Year 1, 2 or 3 to end in 2022
- PI with current active proposal and not eligible for renewal

Deadline

June 17, 2022

Renewal

Eligibility

- PI with current active proposal in Year 1 or 2 and eligible^(*) to submit a renewal proposal

^(*) PI will receive a notification from INCITE Program Manager with approval to submit a renewal.

Deadline

July 22, 2022

Today's Agenda

- Allocation programs [9]
- INCITE mission and recent stats [10 – 23]
- LCF computing systems [24 –32]
- Tips for applicants [33 – 54]
 - Common oversights
 - Requesting a startup account
 - Benchmarking data
- Q&A [55, open discussion]
- Conclusions [56 – 62]
 - Submittal, review, and awards decisions
 - Contact links



Changes to the Program / Proposal Submission

- **Able to request access to testbeds**
 - AI testbed at ALCF, and Quantum testbed at OLCF (see next slide)
- **New two years ago, but a reminder:**
 - **All** submission materials must be concatenated together into a single PDF file for submission. Individual files will not be permitted.
 - For new proposals, the required elements are Project Executive Summary, Project Narrative, Personnel Justification and Management Plan, a Milestone Table, Publications resulting from prior INCITE awards (if appropriate), and PI/Co-I biosketches



Changes to the Program / New Testbeds

- The ALCF and OLCF have testbed resources for new technologies, details below. If you would like access to these resources to support the work in this proposal, please provide the following information. Describe the experiments you would be interested in performing, resources required, and their relationship to the current proposal. Please note, these are smaller experimental resources and a large amount of resources is not available. Instead, these resources are to explore the possibilities for how these technologies might innovate future work.
- **The ALCF AI Testbed**
 - <https://www.alcf.anl.gov/alcf-ai-testbed>
- **OLCF Quantum Computing User Program**
 - <https://www.olcf.ornl.gov/olcf-resources/compute-systems/quantum-computing-user-program>

Changes to the Program / New Resources

- **ALCF Polaris** system is a hybrid CPU/GPU machine that will be available to 2023 INCITE projects on January 1, 2022.
- **OLCF Frontier** is an exascale hybrid CPU/GPU machine that will be available to 2023 INCITE projects.
- **ALCF Aurora** is an exascale hybrid CPU/GPU supercomputer that is leveraging several technological innovations to support machine learning and data science workloads alongside traditional modeling and simulation applications.

For the 2023 INCITE proposal submission, PIs are encouraged to express interest in Aurora for the 2023 cycle.

Changes to the Program / Early Career Track

- For the INCITE 2023 Call for Proposals, INCITE is committing **10%** of allocatable time to an **Early Career Track**.
- 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 (PhD on or after December 31st, 2012) and who have not been a previous INCITE PI 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.
- Project size: any size within the guidance given to INCITE projects.

Primary ways to gain access to LCF

Current distribution of allocable hours

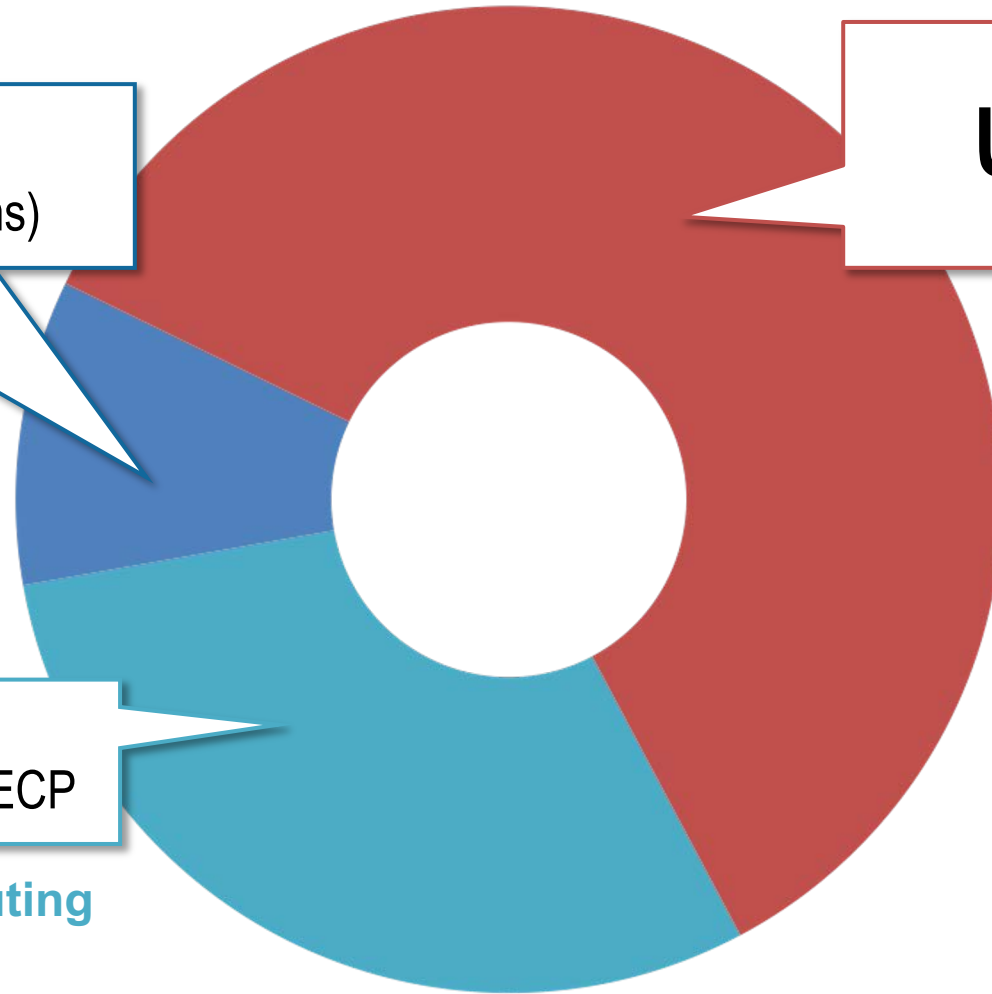
10% Director's Discretionary
(Includes LCF strategic programs)

Up to 60% INCITE

Leadership-class computing

30% ASCR Leadership
Computing Challenge & ECP

DOE/SC capability computing



What is 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 **compute intensive and/or data intensive** large-scale research projects.



INCITE Seeks High-impact Research Campaigns

Examples of previous successful INCITE applications that advance the state-of-the-art across a broad range of topics and different mission priorities

Astro-Physics

- Glimpse into dark matter
- Supernovae ignition

Chemistry

- Chemical catalyst design
- Batteries
- Solar Cells

Engineering

- Turbulent flow
- Propulsor systems
- Nanodevices

Earth Science

- Global climate
- Carbon sequestration

Biology

- Protein structure
- Creation of biofuels
- Replicating enzyme functions
- Membrane channels
- Protein folding

Additional topic areas

- Data and AI
- Computer Science
- Plasma Physics
- Quantum Physics
- Materials Science/Engineering
- Nuclear

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 and/or data intensive runs that cannot be done anywhere else: *capability, architectural needs, data and AI at scale*

3 Eligibility criterion

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

Twofold review process

	New proposal assessment	Renewal assessment
1 Peer review: INCITE science panels + (when appropriate) Data/AI panel	<ul style="list-style-type: none"> • Scientific and/or technical merit • Appropriateness of proposal method, milestones given • Team qualifications • Reasonableness of requested resources 	<ul style="list-style-type: none"> • Scientific and/or technical merit • Met milestones • Change in scope • On track to meet future milestones
2 Computational readiness review: LCF centers	<ul style="list-style-type: none"> • Technical readiness • Appropriateness for requested resources 	<ul style="list-style-type: none"> • Met technical/computational milestones • On track to meet future milestones
Award Decisions	<ul style="list-style-type: none"> • INCITE Awards Committee comprised of LCF directors, INCITE program manager, LCF directors of science, senior management 	

Some limitations on what can be done

- Federal law regulates what can be done on these systems
 - LCF systems have cyber security plans that bound the types of data that can be used and stored on them
- Some kinds of information we cannot have
 - Personally Identifiable Information (PII)
 - Classified Information or National Security Information
 - Unclassified Controlled Nuclear Information (UCNI)
 - Naval Nuclear Propulsion Information (NNPI)
 - Information about development of nuclear, biological or chemical weapons, or weapons of mass destruction

Proprietary Work

Proprietary use (e.g., using the resources to generate data or results that they wish to designate as proprietary) of the HPC resources is allowed and will include **full cost recovery through a proprietary user agreement**. The use of proprietary or export-controlled software applications or input data does not by itself constitute proprietary use of the facility. A single and standardized peer-review selection process and project reporting requirement, as described within the INCITE Overview and Policies, will be implemented for both proprietary and non-proprietary user proposals. Given the complexities associated with proprietary research, **individuals considering submittal of a proposal for proprietary research must contact the INCITE manager, INCITE@DOEleadershipcomputing.org**, before the call for proposals closes to discuss the policy on proprietary work.

[Reference: INCITE Overview and Policies.

See <http://www.doeleadershipcomputing.org/proposal/call-for-proposals/>]

TIP

Do not include proprietary and/or export controlled information in the proposal, whether or not the project itself is nonproprietary.

Community vs. Umbrella Proposals


The program encourages **community proposals** that reflect community consensus on, and participation in pursuing, a particular investigative strategy for addressing a grand-challenge-scale science problem. This consensus should be in large part reflected in the proposal's management plan, which describes how the problem will be tackled and the interrelationship of the activities of each of the proposal team investigators.

The program does not encourage **umbrella proposals**, which is defined to be a collection of individual proposals aimed at various aspects of a particular science problem but with no clearly articulated interrelationship of the activities of each of the proposal team investigators and/or that do not exploit the unique capabilities of the leadership facilities.

2022 INCITE Award Statistics

- Call for proposals closed June 18, 2021
- LCF resources were nearly 3-4.5X over-subscribed based on requests
- 60% of allocable time on the LCF production resources (Summit and Theta) were awarded for CY 2022
- 53 projects awarded of which 15 are renewals

Lead PI Affiliations (Awards, by Project)



Acceptance rates

35% of nonrenewal submittals were accepted.

Contact information

Katherine M. Riley, INCITE Manager
INCITE@DOEleadershipcomputing.org

2022 Award Statistics

	Summit	Theta	Polaris
Number of projects*	37	17	11
Average Project	508 K	1.24 M	102 K
Median Project	540 K	1.00 M	100 K
Total Awards (node-hrs in CY2022)	18.8 M	21.1 M	1.22 M

- Total of 53 INCITE projects (6 projects received time on both Theta and Summit)
 - 4 have time on both Theta and Summit
 - 11 have time on both Theta and Polaris
 - 1 project has time only on Polaris

* All reported in node-hours native to each resource.

2022 New Proposal Statistics

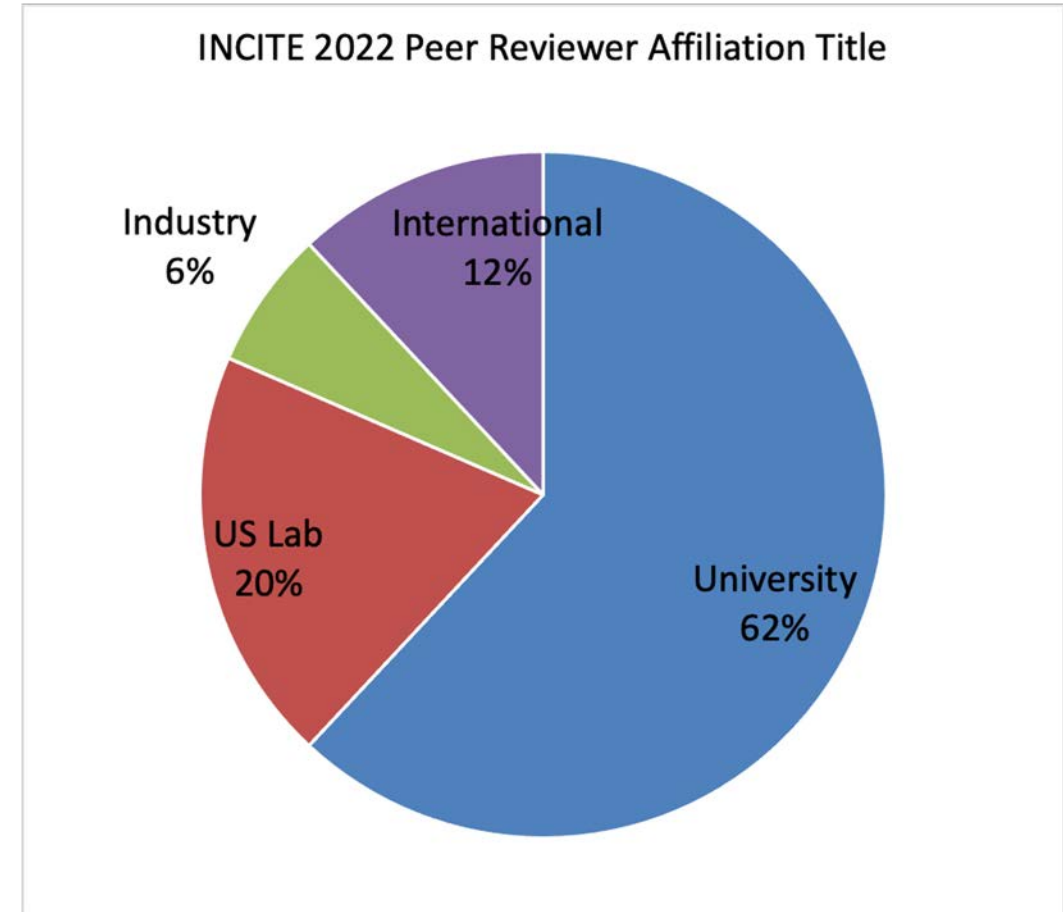
- 42% of the PI's had never before led an INCITE proposal
 - 106 new proposals, 45 led by new PI's
- 26% of non-renewal projects awarded to new PI's
 - 53 new projects awarded, 14 led by new PI's

INCITE actively engages with new research teams through outreach such as workshops, email distributions, and individual networking.



2022 INCITE Science Panels

- 92 science experts
- Diverse panels of reviewers including Society Fellows (AAAS, APS, IEEE, etc.), NSF or DOE Early Career scientists, Laboratory Fellows, National Academy members, Department Chairs or Full Professors
- Significant carryover from year to year to promote continuity
- When appropriate, proposals may also be assessed by a separate Data and AI Panel



2022 INCITE
Reviewer Affiliation

New Early Career Path

- Why?
 - INCITE looks to support and encourage the next generation of researchers by committing 10% or allocatable hours
- Who can apply to the Early Career Path?
 - Researchers less than 10 years out from their PhD who need LCF-level capabilities to advance their overall research plan, and who have not been a previous INCITE PI.
- How to apply to the Early Career Path?
 - The option to self-identify as Early Career will be available on the INCITE 2022 application.
- Will it increase my chances of an award?
 - If you qualify, it could. Overall, projects must still be meritorious in the peer review and computationally ready.

TIPS

Do examine the Frequently Asked Questions for these and other topics. (#20 - #23)

INCITE HPC Resources 1Y, 2Y and 3Y proposals

LCF

Year 1
2023



Summit
IBM/NVIDIA



Frontier
Cray/HPE

Year 2
2024



Frontier
Cray/HPE

Year 3
2025



Frontier
Cray/HPE



Theta
Intel-Cray XC40



Polaris
HPE



Polaris
HPE



Polaris
HPE



Aurora
Intel-HPE



Aurora
Intel-HPE



Aurora
Intel-HPE

Estimations of Project Award Sizes and node-hours

Oak Ridge LCF

System Peak Performance

200 PF

>1.5 EF

Awards size

Machine node-hours

300-800 K

~ 1M

Summit equivalent node-hours

300-800 K

3-4 M



Summit
IBM/NVIDIA

2023



Frontier
Cray/HPE

2023-2025

Argonne LCF

Node Performance compared to Theta

1x

System Peak Performance

11.7 PF

30x

44 PF

≥2 EF

Awards size

Machine node-hours

1-2.5 M

100-200 K

N/A

Theta equivalent node-hours

1-2.5 M

3-6 M

50-100 M



Theta
Intel-Cray XC40

2023



Polaris
HPE

2023-2025



Aurora
Intel-HPE

2023-2025

100x

OLCF Resources - Summit



Summit Specs	
Peak Performance	200 PF
Compute Nodes	4,608
Node	2 IBM Power9 Processors + 6 NVIDIA Volta GPUs
Node Performance	42 TF
Memory per node	512 GiB DDR4 + 96 GiB High Bandwidth Memory (HBM2)
NV Memory per node	1,600 GB
GPU Link	NVLink 2.0
Total System Memory	> 10 PB (DDR4 + HBM2 + Non-volatile)
Interconnect	Mellanox Non-blocking Fat Tree, Dual Rail EDR-IB (23 GB/s)
File System Performance	250 PB, 2.5 TB/s, GPFS

OLCF Resources - Summit

- Compilers supporting OpenMP or OpenACC
 - IBM XL, PGI, LLVM, GNU, NVIDIA
- Libraries
 - IBM Engineering and Scientific Subroutine Library (ESSL)
 - FFTW, ScaLAPACK, PETSc, Trilinos, BLAS-1,-2,-3, NVBLAS
 - cuFFT, cuSPARSE, cuRAND, NPP, Thrust
- Debugging
 - Allinea DDT, IBM Parallel Environment Runtime Edition (pdb)
 - Cuda-gdb, Cuda-memcheck, valgrind, memcheck, helgrind, stacktrace
- Profiling
 - IBM Parallel Environment Developer Edition (HPC Toolkit)
 - VAMPIR, Tau, Open|Speedshop, nvprof, gprof, Rice HPCToolkit
- Development Support
 - HIP, ROCm
- Machine/deep learning (ML/DL), workflow
 - Popular frameworks (TensorFlow, PyTorch, etc.) Python libraries, NVIDIA Rapids, Watson ML, etc.
 - Visualization: Paraview, Visit, SIGHT, etc.
 - OLCF Slate Workflow Cluster, Pegasus, Radical Toolkit, Flux, etc.;

https://docs.olcf.ornl.gov/systems/summit_user_guide.html

OLCF Resources - Frontier



The system will be based on Cray's new Shasta architecture and Slingshot interconnect with high-performance AMD EPYC CPU and Radeon Instinct GPU technology. The new accelerator-centric compute blades will support a 4:1 GPU-to-CPU ratio with high-speed links and coherent memory between them within the node. With Frontier, scientists will be able to pack in more calculations, identify new patterns in data, and develop innovative data analysis methods to accelerate the pace of scientific discovery.

Frontier Specs	
Peak Performance	> 1.5 EF
Node	1 HPC and AI Optimized AMD EPYC CPU (64-cores) 4 Purpose Built AMD Instinct MI250X GPUs
Number of Nodes	9408
CPU-GPU Interconnect	AMD Infinity Fabric Coherent memory across the node
System Interconnect	Multiple Slingshot NICs providing 100 GB/s network bandwidth Slingshot dragonfly network which provides adaptive routing, congestion management and quality of service.
Storage	2-4x performance and capacity of Summit's I/O subsystem. Frontier will have near node storage like Summit.

<https://www.olcf.ornl.gov/frontier>

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)



ALCF Resources - Theta

Cray Programming Environment

- Programming Models

- Distributed memory: MPI, SHMEM, GA
- Shared Memory: OpenMP 3.1, OpenACC 2.0

- Compilers

- Cray Compiling Environment (CCE), GNU, Intel, LLVM/Clang

- Optimized Libraries

- BLAS, LAPACK, ScaLAPACK, FFTW, Cray PETSc, Cray Trilinos, Iterative Refinement Toolkit
- I/O: NetCDF, HDF5

- Tools

- Debuggers & Tools: Allinea DDT, Igdb, ATP, STAT
- Performance Analysis: CrayPat, Cray Apprentice2
- Porting Tools: Reveal, CCDB

- Machine/Deep Learning & Workflow Software

- Deep learning frameworks including Intel Optimized Tensorflow and PyTorch, Keras, and Horovod
- Python distributions (Anaconda and Intel) as well as machine learning and analytics packages including Sci-kit learn, numpy, scipy, etc...
- Optimized with Intel MKL and MKL-DNN
- Support for Singularity containers, JupyterHub, Spark-based Apache BigData stack, Databases such as MongoDB, Postgres, R-based analytics, workflows including Balsam, DeepHyper and NAS
- Visualization: ParaView

<https://www.alcf.anl.gov/support-center/theta>

ALCF Resources - Polaris

Polaris Specs	
Platform	HPE Apollo 6500 Gen 10+
System Peak	44 PF DP
Peak Power	1.8 MW
Total System Memory	280 TB CPU DDR4 + 87.5 TB GPU HBM2
System Memory Type	DDR, HBM
Node Performance	78 TF DP
Node Processors	(1) AMD EPYC "Milan" 7543P CPU (4) NVIDIA HGX A100 GPUs
System Size (nodes)	560
Interconnect	HPE Slingshot



- Polaris is a hybrid CPU/GPU leading-edge testbed system that will give scientists and application developers a platform to test and optimize codes for Aurora
- Polaris is available to New and Renewal 2023 INCITE Proposals.

<https://www.alcf.anl.gov/polaris>

ALCF Resources - Polaris

- Programming Environment

- HPE Cray PE for Polaris
- NVIDIA HPC SDK
- SYCL/Data Parallel C++
 - CodePlay computecpp compiler
 - LLVM via intel DPC++ branch

- Programming Models

- OpenMP 4.5/5
- DPC++/SYCL
- Kokkos
- Raja
- HiP

- Performance / Profiling / Debugging

- GPU tools
- PAPI, TAU, HPCToolkit
- PAT (Performance Analysis Tool)
- NVIDIA® Nsight™
- DDT, STAT, gdb4hpc, CUDA-GDB, gdb

- Frameworks

- Python / Numba
- Tensorflow
- Pytorch
- Cray/HPE AI Stack (HPO, Jupyter, SmartSim)
- Spark BigData Stack

ALCF Resources - Aurora

Exascale architecture designed for Simulation, Data & Learning

Aurora Specs	
Peak Performance	≥2 EF DP sustained
Node	2 Intel Xeon scalable processors (Sapphire Rapids); 6 Intel Xe ^e arch based GPU (Ponte Vecchio)
GPU Architecture	Xe ^e arch based GPU (Ponte Vecchio); Tile based, chiplets, HBM stack, Foveros 3d integration
CPU-GPU Interconnect	PCIe
Aggregate System Memory	>10 PB
Interconnect	HPE Slingshot 11; Dragonfly topology with adaptive routing
Network Switch	25.6 TB/s per switch, from 64-200 GB ports (25GB/s per direction)
High-Performance Storage	≥230 PB, ≥25 TB/s (DAOS)



Aurora Specs	
Node Performance	> 130 TF
Compute nodes	> 9,000
Cabinets	> 100
Node Memory Architecture	Unified memory architecture, RAMBO

<https://aurora.alcf.anl.gov>

ALCF Resources - Aurora

- Platform
 - HPE Cray XE
- Software Stack
 - HPE Cray XE software stack + Intel enhancements + Data and Learning
- Compilers
 - Intel, LLVM, GCC
- Programming Models
 - Intel oneAPI, OpenMP, DPC++/SYCL
- 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), Open SHMEM, Python, Numba, MPI, OpenCL
- Programming Tools
 - Open|Speedshop, TAU, HPCToolkit, Score-P, Darshan, Intel Trace Analyser and Collector, Intel Vtune, Advisor, and Inspector, PAPI, GNU gprof
- Debugging and Correctness Tools
 - Stack Trace Analysis Tool, gdb, Cray Abnormal Termination Processing
- Math Libraries
 - Intel MKL, Intel MKL-DNN, ScaLAPACK
- 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

<https://aurora.alcf.anl.gov>

Key Questions to Ask Yourself

- Is both the scale of the runs and the time demands of the problem of LCF scale?
 - Yes, I can't get the amount of time I need anywhere else.
 - Yes, my applications are too large to run on other systems.
- Do you need specific LCF hardware or systems?
 - Yes, the very large memory and I/O available here are necessary for my work.
 - Yes, my application requires mixed/reduced precision accelerator hardware.
- Do you have the people ready to do this work?
 - No, I'm waiting to hire a postdoc.
 - Yes, I have commitments from the major participants.

Key Questions to Ask Yourself (cont.)

- Do you have large data/AI needs?
 - Yes, my data-intensive needs require the LCF resources.
- Do you have a workflow solution?
- Do you have a post-processing strategy?
- Do you use ensemble runs and need LCF resources?
 - My ensembles can run under the direction of a large job or workflow manager, with I/O scaling on a parallel file system -> possible yes
 - My ensemble expects to run millions of serial batch jobs on nodes with local disk available -> probably no
- Do you understand the life cycle of your data?

Note: Some of these characteristics are negotiable, so make sure to discuss atypical requirements with the centers.

Proposal Form: Outline

0	Project title
1	Principal investigator and co-principal investigators
2	Project information (Research category, Project Summary)
3	INCITE allocation request; Other funding/computing support
4	Project narrative, other materials
	(A) Executive summary (1 page)
	(B) Project narrative including impact of the work, objectives, benchmarking (15 pages)
	(C) Personnel justification & management plan
	(D) Milestone table
	(E) Publications resulting from INCITE Awards
	(F) PI / Co-I Biographical Sketches
5	Software applications and dependent packages
6	Wrap-up Questions (Proprietary and sensitive information, Export Control)
7	Outreach and suggested reviewers

Getting Started: Know Your Audience

- Remember, INCITE is very broad in scope
 - Computational-science-savvy senior scientists/engineers drawn around the world from national labs, universities, and industry
 - They will be assessing potential impact of this work versus other proposals submitted

TIPS

Don't assume that your audience is familiar with your work through other review programs (ex. funding agencies). INCITE is very broad in scope and you may be competing against a diverse set of proposals.

Do look at the reviewer questionnaires, posted on the INCITE Call for Proposals site

Narrative: Impact of the Work

This is the principal determinant of a successful submittal!

- What is the scientific challenge and its significance
- Impact of a successful computational campaign — the big picture
- Reasons this work needs to be done now, on the resources requested

TIPS

Do give a compelling picture of the impact of this work, both in the context of your field and, where appropriate, beyond.

Do explain why this work cannot be done elsewhere. Reviewers scrutinize whether another allocation program may be a better fit.

Narrative: Objectives and Milestones

- Successful submittals must also very clearly
 - Describe approach to solving the problem, its challenging aspects, preliminary results
 - Tie to the resources requested your key objectives, key simulations, and project milestones in your milestone table

TIPS

Do clearly articulate your project's milestones for each year. Reviewers have downgraded proposals that don't show that the PI has a well thought out plan for using the allocation and LCF resources.

Do bear in mind that the average INCITE award of time for a single project is equivalent to several million dollars.

Narrative: Computational Approach

Provide the basic foundation

- Describe the underlying formulation
 - Don't assume reviewers know all the codes
 - Do show that the code you plan to employ is the correct tool for your research plan
 - Do explain the differences if you plan to use a private version of a well-known code
- List programming languages, libraries and tools used
 - Check that what you need is available on the system
- List required software (noting that users may need to build some software themselves)

Narrative: Computational Approach

Provide the basic foundation

- Data Management

- Requirements for bringing input data, storage, and movement of data
- Describe how long data needs to remain on spinning disk and archive and how those times were determined
- Describe the long term use of data (e.g. publicly available)
- Describe tools or infrastructure required to support data needs

- Data and AI, and Workflows

- Describe any AI, analytics and viz that are part of the project, and how you intend for that work to be done
- If appropriate, describe workflow tools that will be used to facilitate the volume of work

Narrative: Computational and/or Data Intensive Campaign

- Describe the kind of runs you plan with your allocation
 - L exploratory runs using M nodes for N hours
 - X big runs using Y nodes for Z hours
 - P analysis runs using Q nodes for R hours
- Big runs often have big output and/or big I/O
 - Show you can deal with it and understand the bottlenecks
 - Understand the size of results, where you will analyze them, and how you will get the data there

TIPS

Do clearly emphasize the relationship between the proposed runs and the major milestones. This helps the Awards Committee maximize your milestones, if they can't grant the full award requested.

The details are important!

Code Performance Overview

Performance data should support the required scale

- Use similar problems to what you will be running
- Show that you can get to the range of nodes required
- Demonstrate efficient use of nodes and node specific hardware (e.g. GPUs on Summit, HW threads on Theta)
- Best to run on the same machine, but similar size runs on other machines can be useful
- Be clear about the number of nodes, MPI ranks, threads and GPUs (if applicable) being used in runs
- Include production style I/O in benchmarks (checkpoint/restart, analysis)
- Describe how you will address any scaling deficiencies

TIPS

Do provide performance data in the requested format.

Do provide performance of the scaling baseline, not just scaling efficiency

Parallel Performance: Direct Evidence

WEAK SCALING,

which is defined as how the solution time varies with the number of nodes for a fixed problem size *per node*:

Increase problem size as resources are increased


STRONG SCALING,

which is defined as how the solution time varies with the number of nodes for a fixed *total* problem size:


Increase resources (nodes) while doing the same computation

Pick the approach(es) relevant to your work and show results

Weak Scaling Example



Strong Scaling Example



Note: the axes could be samples per second or throughput versus number of nodes. For Data and AI applications, show the scaling (e.g., for convergence) dependence on factors such as hyperparameters and discuss scientific impact.

More About our Ensemble Policy

“Can I meet the computationally intensive criterion by loosely coupling my jobs?”

- Possibly “yes”
 - If you require large numbers of discrete or loosely coupled simulations where time-to-solution is an untenable pacing issue, **and**
 - If a software workflow solution (e.g., pre- and post-processing scripts that automate run management and analysis) is provided to facilitate this volume of work.
- Probably “no”
 - If by decoupling the simulations the work could be effectively carried out on a smaller resource within a reasonable time-to-solution.

TIPS

Do examine the Frequently Asked Questions (#33)

Data and AI Applications

In addition to traditional computationally intensive simulation campaigns, INCITE encourages Data and AI projects.

- Strong aspects of scalable data processing and/or AI.
- Example areas
 - Machine and/or Deep learning
 - Data-intensive computing
 - Experimental/observational/simulation data analytics
 - Complex and interactive workflows
 - Streaming/real-time data analysis
 - Statistical methods
 - Graph analytics
 - Uncertainty quantification

12% of proposals in 2022 reviewed by Data/AI panel

Data and AI Applications

In addition to traditional computationally intensive simulation campaigns, INCITE encourages Data and AI projects.

- Successful proposals must
 - Clearly articulate the objectives and dependencies of the end-to-end research campaign e.g. training data requirements and generation; deep learning/data-driven model selection and validation; model embedding to augment simulations etc.
 - List the application requirements, including databases, machine learning/deep learning frameworks, workflow software, containers, etc.
 - Demonstrate that the software can run efficiently on the resources requested.
- Proposals which target the convergence or interleaving of simulation, data analytics and AI are also encouraged

TIPS

Do examine the Frequently Asked Questions for these and other topics.
(#34 & #35)

Narrative: Development Work –

- Developmental Work

- Describe any additional development required to execute any of your milestones. Identify any dependencies and how you will validate.
- Estimate the computational resources required for this work and when it will be completed.



Narrative: Management Plan

- Experience and credibility
 - List the scientific and technical members and their experience as related to the proposed scientific or technical goals
 - Successful proposal teams demonstrate a clear understanding of petascale computing and can optimally use these resources to accomplish the stated scientific/technical goals
 - Team members can be part of different proposals
- Transparent use of time
 - Projects involving multiple teams or different thrust areas should clearly state how the allocation will be distributed and managed

TIPS

Do include in “Personnel Justification” a brief description of the role of each team member. Although not a requirement, proposals with application developers or clear connections to development teams are favorably viewed by readiness reviewers.

Narrative: Milestone Table

- Clearly state the scientific and technical milestones for each year of your proposed work.
 - Milestones should be appropriate for the size and duration of the requested award. Future modifications in scope will be tracked through this table.
 - As appropriate, details to provide for each milestone:
 - Computing resource and allocation request (node-hours).
 - Production size runs (number of nodes).
 - Filesystem and archival storage, and duration.
 - Software application employed.
 - Summary of the computational tasks (e.g., computational runs, data analysis, data reduction).
 - Dependencies on or between other milestones.
 - Expected start, target completion date and status.

Note:
The Milestone Table is a required document. Proposals without it will not be accepted for review.

Proposal Title Lead PI

[Refer to the guidelines for instructions in preparing the proposal. Table does not count toward project narrative page limit.]

Proposal Title (exactly as it appears on submission): Insert Text

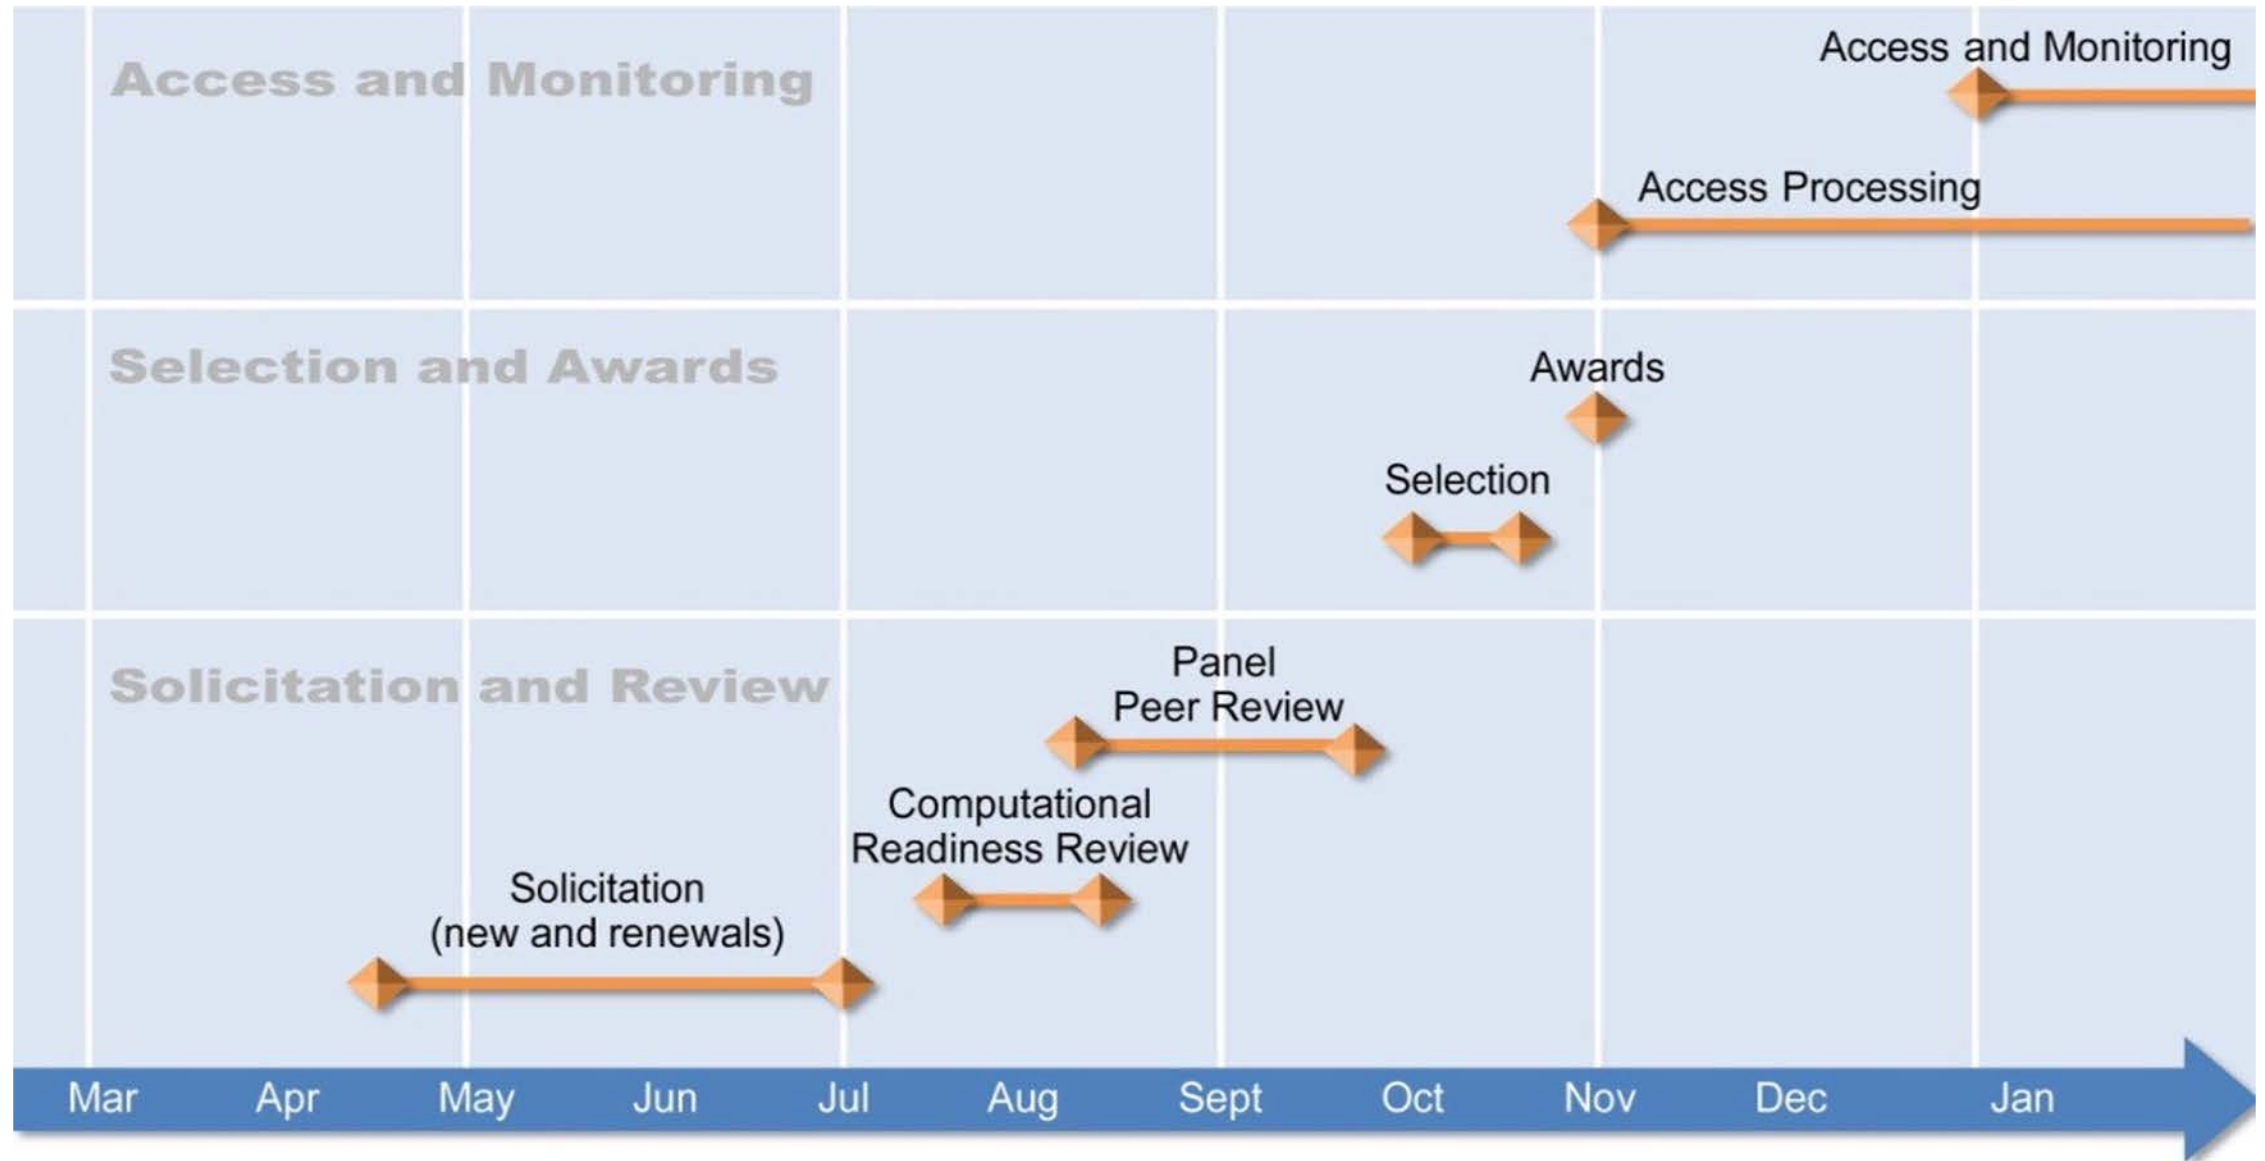
Year 1		Total number of node-hours for Year 1: <i>Insert Text</i>	
Milestone:	Details (as appropriate):	Dates:	Status: (renewals only)
Insert Text	Resource: Insert Text Node-hours: Insert Text Production size runs (number of nodes): Insert Text Filesystem storage (TB and dates): Insert Text Archival storage (TB and dates): Insert Text Software Application: Insert Text Tasks: Insert Text Dependencies: Insert Text	Insert Text	Insert Text
Insert Text	Resource: Insert Text Node-hours: Insert Text Production size runs (number of nodes): Insert Text Filesystem storage (TB and dates): Insert Text Archival storage (TB and dates): Insert Text Software Application: Insert Text Tasks: Insert Text Dependencies: Insert Text	Insert Text	Insert Text
Year 2 (if appropriate)		Total number of node-hours for Year 2: <i>Insert Text</i>	
Insert Text	Resource: Insert Text Node-hours: Insert Text Production size runs (number of nodes): Insert Text Filesystem storage (TB and dates): Insert Text Archival storage (TB and dates): Insert Text Software Application: Insert Text Tasks: Insert Text Dependencies: Insert Text	Insert Text	Insert Text
Year 3 (if appropriate)		Total number of node-hours for Year 3: <i>Insert Text</i>	
Insert Text	Resource: Insert Text Node-hours: Insert Text Production size runs (number of nodes): Insert Text Filesystem storage (TB and dates): Insert Text Archival storage (TB and dates): Insert Text Software Application: Insert Text Tasks: Insert Text Dependencies: Insert Text	Insert Text	Insert Text

Narrative: INCITE Publications

- Publications resulting from **INCITE** awards
 - To show impact of the INCITE program, we ask authors to list the publications, **including DOIs when available**, resulting from previous INCITE awards to this project team for work related to the proposal under consideration.
 - All publications that were enabled by an INCITE project should include an acknowledgement of the INCITE program and/or the LCF.
 - Include only publications with INCITE acknowledgements.



INCITE Annual Timeline



Are You Ready to Apply Now?

✔ Port your code before submitting the proposal

- Check to see if someone else has already ported it
- Request a startup account if needed (see next slide)

✔ Provide compelling benchmark data

- Prove application scalability in your proposal
- Run example cases at proposed production scale
- If you cannot show proof of runs at the proposed production scale, then provide a very tight story about how you will succeed

TIPS

Do make the benchmark examples as similar to your production runs as possible, or, make it clear why another benchmark example is valid for your proposed work.

Request a Start-up Account Now

- Director's Discretionary (DD) requests can be submitted anytime
- DD may be used for porting, tuning, scaling in preparation for an INCITE submittal
- Submit applications at least 2 months before INCITE Call for Proposals closes

Argonne DD Program:

<https://www.alcf.anl.gov/science/directors-discretionary-allocation-program>


Oak Ridge DD Program:

<https://my.olcf.ornl.gov/project-application-new>

- Director's Discretionary Proposals considered year-round
- Awards up to thousands of node-hours
- Allocated by LCF center directors

30% ASCR Leadership Computing Challenge, ECP

60% INCITE



Proposal Form: Final Check

0	Project title
1	Principal investigator and co-principal investigators
2	Project information (Research category, Project Summary)
3	INCITE allocation request; Other funding/computing support
4	Project narrative, other materials
	(A) Executive summary (1 page)
	(B) Project narrative including impact of the work, objectives, benchmarking (15 pages)
	(C) Personnel justification & management plan
	(D) Milestone table
	(E) Publications resulting from INCITE Awards
	(F) PI / Co-I Biographical Sketches
5	Software applications and dependent packages
6	Wrap-up Questions (Proprietary and sensitive information, Export Control)
7	Outreach and suggested reviewers

Open discussion on what authors should include in their proposal

Reminder: The INCITE program has significantly changed its proposal submission site. Please allow extra time to create an account on this site and submit your proposal.

Submitting Your Proposal or Renewal

- You may save your proposal at any time without having the entire form complete
- Required fields must be completed for the form to be successfully submitted
 - An incomplete form may be saved for later revisions
- After submitting your proposal, you **will** be able to edit and resubmit it up until the submission deadline

Submit

INCITE Awards Committee Decisions

- The INCITE Awards Committee is comprised of the LCF center directors, INCITE program manager, LCF directors of science and senior management.
- The committee identifies the top-ranked proposals by a) peer-review panel ratings, rankings, and reports; and b) additional considerations, such as the desire to promote use of HPC resources by underrepresented communities.
- Computational readiness review is used to identify whether the top-ranked proposals are “ready” for the requested system.

INCITE Awards Committee Decisions

- A balance is struck to ensure
 - each awarded project has sufficient allocation to enable all or part of the proposed scientific or technical achievements
 - a robust support model for each INCITE project
- When the centers are oversubscribed, each potential project is assessed to determine the amount of time that may be awarded to allow the researchers to accomplish significant scientific goals.
- Requests for appeals can be submitted to the INCITE manager or LCF center directors. If an error has occurred in the decision-making process (e.g. procedural, clerical), consideration is given by the INCITE management and an award may be granted. Assessment of the relative scientific merit is not considered in appeals.

2023 INCITE Award Announcements

- Awards will be announced by INCITE Manager, Katherine Riley, in November 2022
 - Welcome and startup information from centers
 - Agreements to sign: Start this process as soon as possible!
 - Getting started materials: Work closely with the center
- Centers provide expert-to-expert assistance to help you get the most from your allocation
 - Scientific “Liaisons” (OLCF) and “Catalysts” (ALCF)

PI Responsibilities

- Provide quarterly status updates (on supplied template)
 - Milestone reports
 - Publications, awards, journal covers, presentations, etc., related to the work
- Provide highlights on significant science/engineering accomplishments as they occur
- Submit annual renewal request
- Complete annual surveys
- Encourage your team to be good citizens on the computers
 - Use the resources for the proposed work

Let us know your achievements and challenges

PI Responsibilities (cont.)

- Encourage your team to attend the center's user meeting
- All INCITE users are expected to acknowledge the center and the program in publications resulting from their award:
 - An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program.

and either

- This research used resources of the Oak Ridge Leadership Computing Facility located in the Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725.
- This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357.

It is a Small World...

- Let the science agency that funds your work know how significant the INCITE program and the Leadership Computing Facilities will be to your work
- Contact us if you have questions; we want to hear from you



Contacts

For details about the INCITE program:

www.doeleadershipcomputing.org
INCITE@DOEleadershipcomputing.org



For details about the centers:

www.olcf.ornl.gov
help@olcf.ornl.gov, 865-241-6536
<https://my.olcf.ornl.gov/project-application-new/>



www.alcf.anl.gov
support@alcf.anl.gov
<https://www.alcf.anl.gov/science/directors-discretionary-allocation-program>

