OAK RIDGE National Laboratory

August 6th, 2025 | Oak Ridge National Laboratory

OLCF User Meeting: INCITE Proposal Workshop

Dmytro Bykov

U.S. DEPARTMENT OF ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY

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.

- **Proposal Development Support**
  - https://www.doeleadershipcomputing.org/proposal/proposal-development-support/

Primary ways to gain access to LCF

Current distribution of allocable hours

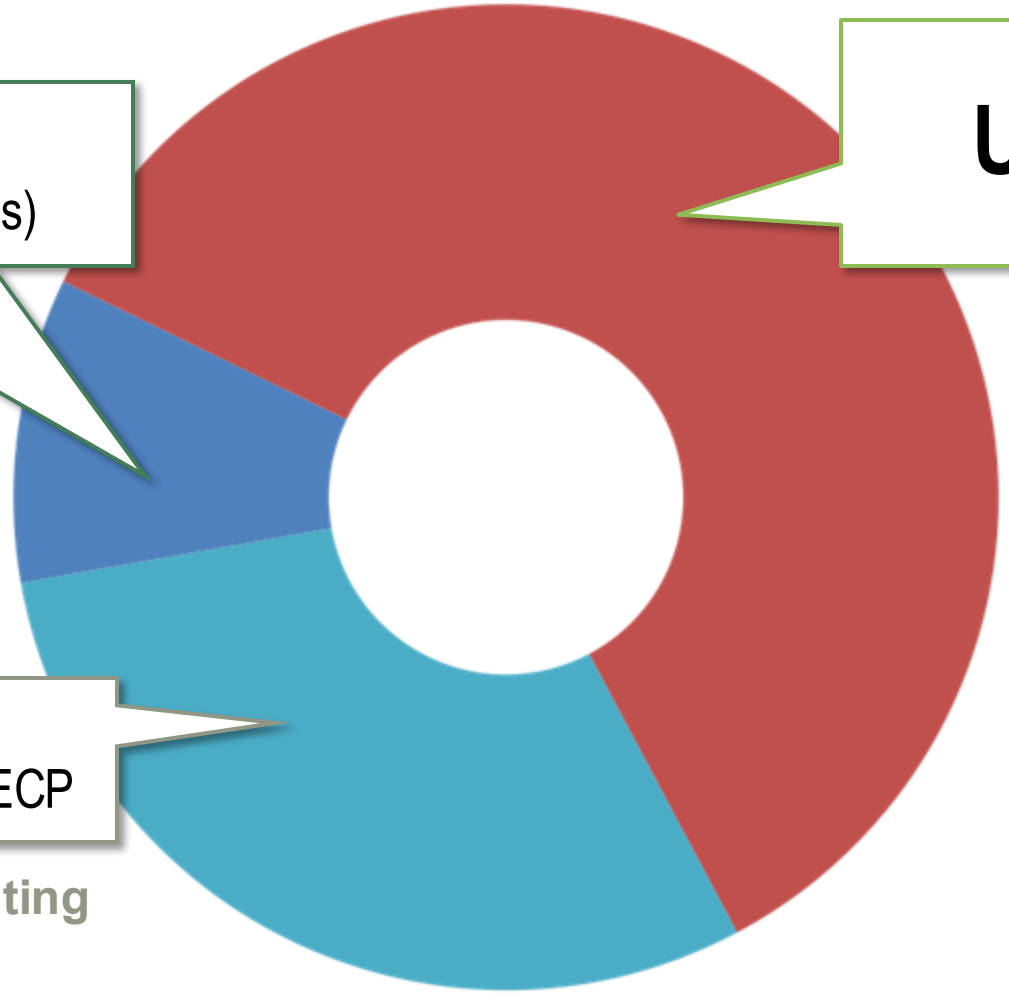

10% Director's Discretionary
(Includes LCF strategic programs)

Up to 60% INCITE

Leadership-class computing

30% ASCR Leadership
Computing Challenge & ECP

DOE/SC capability computing

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 Sciences	• Global climate • Carbon sequestration • Seismology
Biology	• Protein structure • Creation of biofuels • Replicating enzyme functions • Membrane channels • Protein folding
Additional topic areas	• Data and AI • Computer Science • Plasma Physics • Fusion Research • 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

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

Twofold review process

	New proposal assessment	Renewal assessment
1 Peer review: INCITE science panels + (when appropriate) Data/AI panel	• **Scientific and/or technical merit** • Appropriateness of proposal method, milestones given • Team qualifications • Reasonableness of requested resources	• **Scientific and/or technical merit** • **Met milestones** • Change in scope • On track to meet future milestones
2 Technical assessment: 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, senior management	

2025 INCITE Award Summary

- Call for proposals closed June 14, 2024
- 60% of allocable time on the LCF production resources was awarded for CY 2025
- 108 peer reviewers from around the world
- **Submitted:** 146 proposals (110 new, 36 renewal)
- **Awarded:** 81 projects (45 new, 36 renewal)
- 41% of nonrenewal submittals were accepted
- 92% of proposals were meritorious

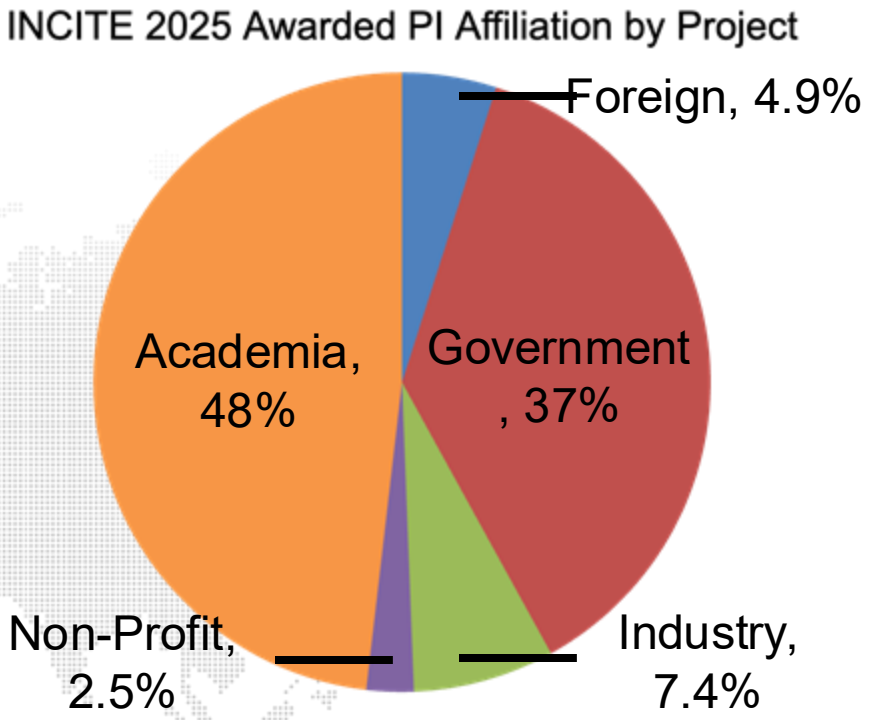

	Requested	Oversubscribed
Frontier	86M	~2.4x
Aurora	48M	~2.4x
Polaris	12M	~6.2x

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

2025 Award Statistics

Number of projects*	55	34	20
Average Project	695,900	665,800	130,000
Median Project	750,000	550,000	200,000
Total Awards (node-hrs in CY2025)	38,275,000	22,640,000	2,600,000

All reported in node-hours native to each resource

- Total of INCITE projects : 81
  - Frontier, Aurora, Polaris : 4
  - Frontier, Aurora : 12
  - Frontier, Polaris: 3
  - Aurora, Polaris : 5

2025 INCITE Science Panels

- 103 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

Estimating Time on LCF Resources

- Best Practice: Estimates based on running on resource of interest
  - Polaris and Frontier are currently available
  - Request a DD at the appropriate center(s)
  - Perform scaling and timings on a representative problem
- Alternative: Using a similar resource
  - Perform timings on a resource as similar as possible
    - Frontier - Use an AMD GPU machine (ideally MI250X), otherwise a comparable GPU machine
    - Polaris - Use an NVIDIA GPU machine (ideally A100), otherwise a comparable GPU machine
    - Aurora - Use a comparable GPU machine or Aurora data if you have it (see next slide)
  - Perform timings on the closest node architecture and scaling on an appropriate sized GPU resource
  - Estimate performance on the targeted machine(s) using publicly available information and include the appropriate details in the proposal

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 o… a parallel file system -> possibly 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.

Getting Started: Know Your Audience

- Remember, INCITE is very broad in scope
  - Reviewers are computational-science-savvy senior scientists and 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

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

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 proposals must also very clearly
  - Describe the approach to solving the problem, its challenging aspects, preliminary results
  - Tie the resources requested to 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 the team may need to build some software)

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 shape 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 Frontier or Aurora)
- 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 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,	STRONG SCALING,
Increase problem size as resources are increased	Increase resources (nodes) while doing the same computation

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

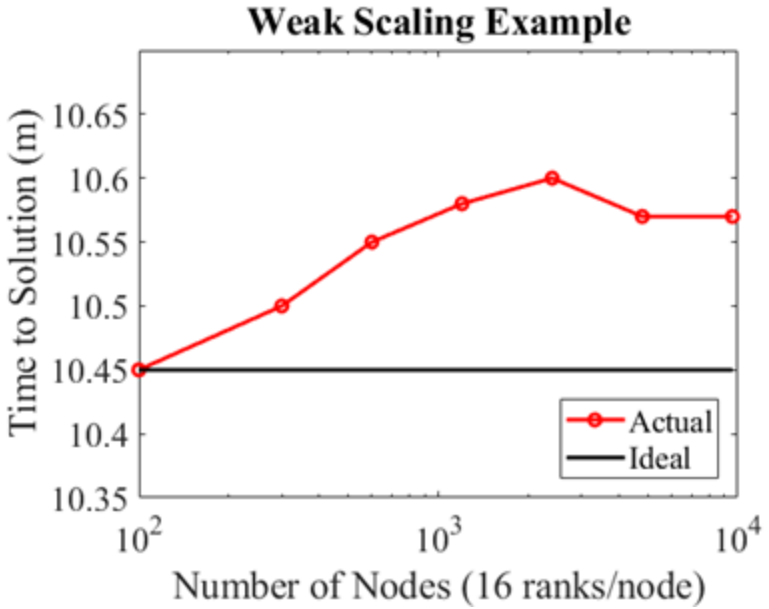

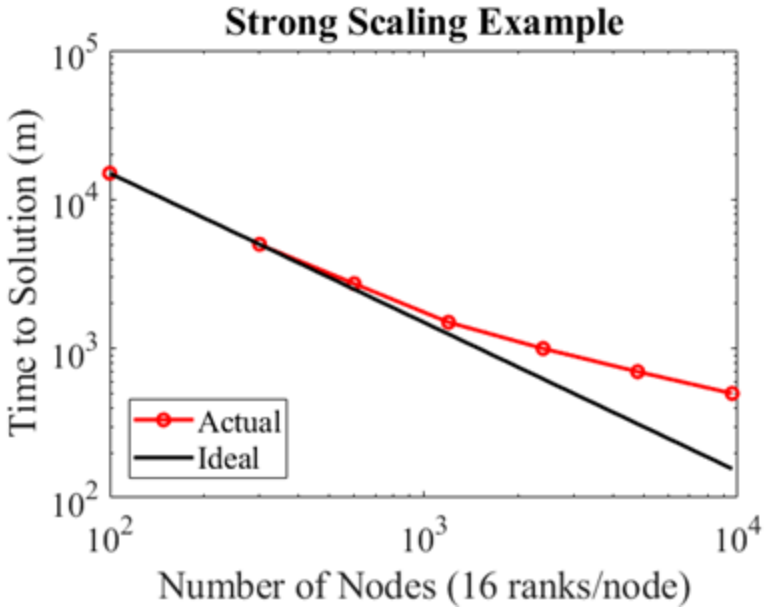

Note: the axes could be samples per second or throughput versus number of nodes. For Data and AI applications, show the scaling (e.g., for convergence) dependence on factors such as hyperparameters and discuss scientific impact.

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)

https://www.doeleadershipcomputing.org/incite-faqs

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

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)

https://www.doeleadershipcomputing.org/incite-faqs

Data Usage (All Projects)

- List your total online and offline (archive) storage for the project
- For each milestone list the total online and offline storage
  - Include the estimated duration the data is need
- For large data projects include a discussion of the data requirements
  - storage amount
  - duration
  - brief description of the type of data
- If you have any special or unusual requirements please be sure to describe them in your proposal
- Make sure you read and are familiar with the center(s) data policies

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.

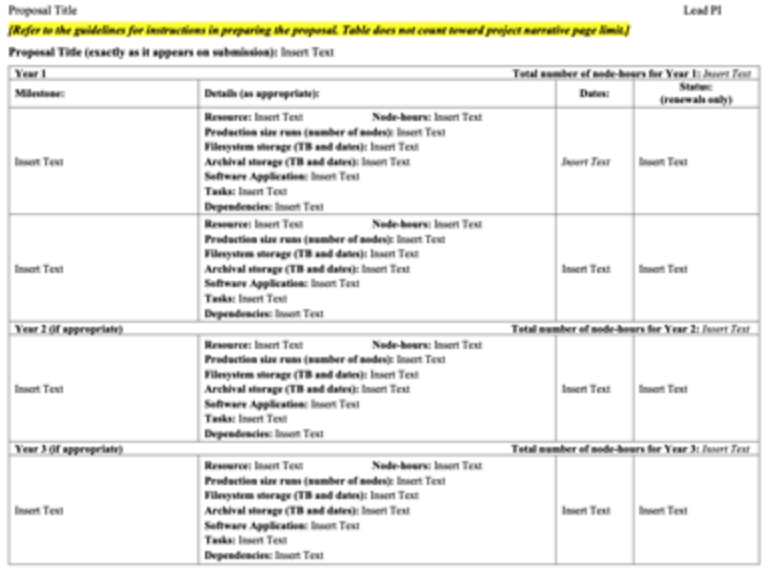

Proposal Title — Lead PI

[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:** *Insert Text*
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:** *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 3 (if appropriate)			**Total number of node-hours for Year 3:** *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

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

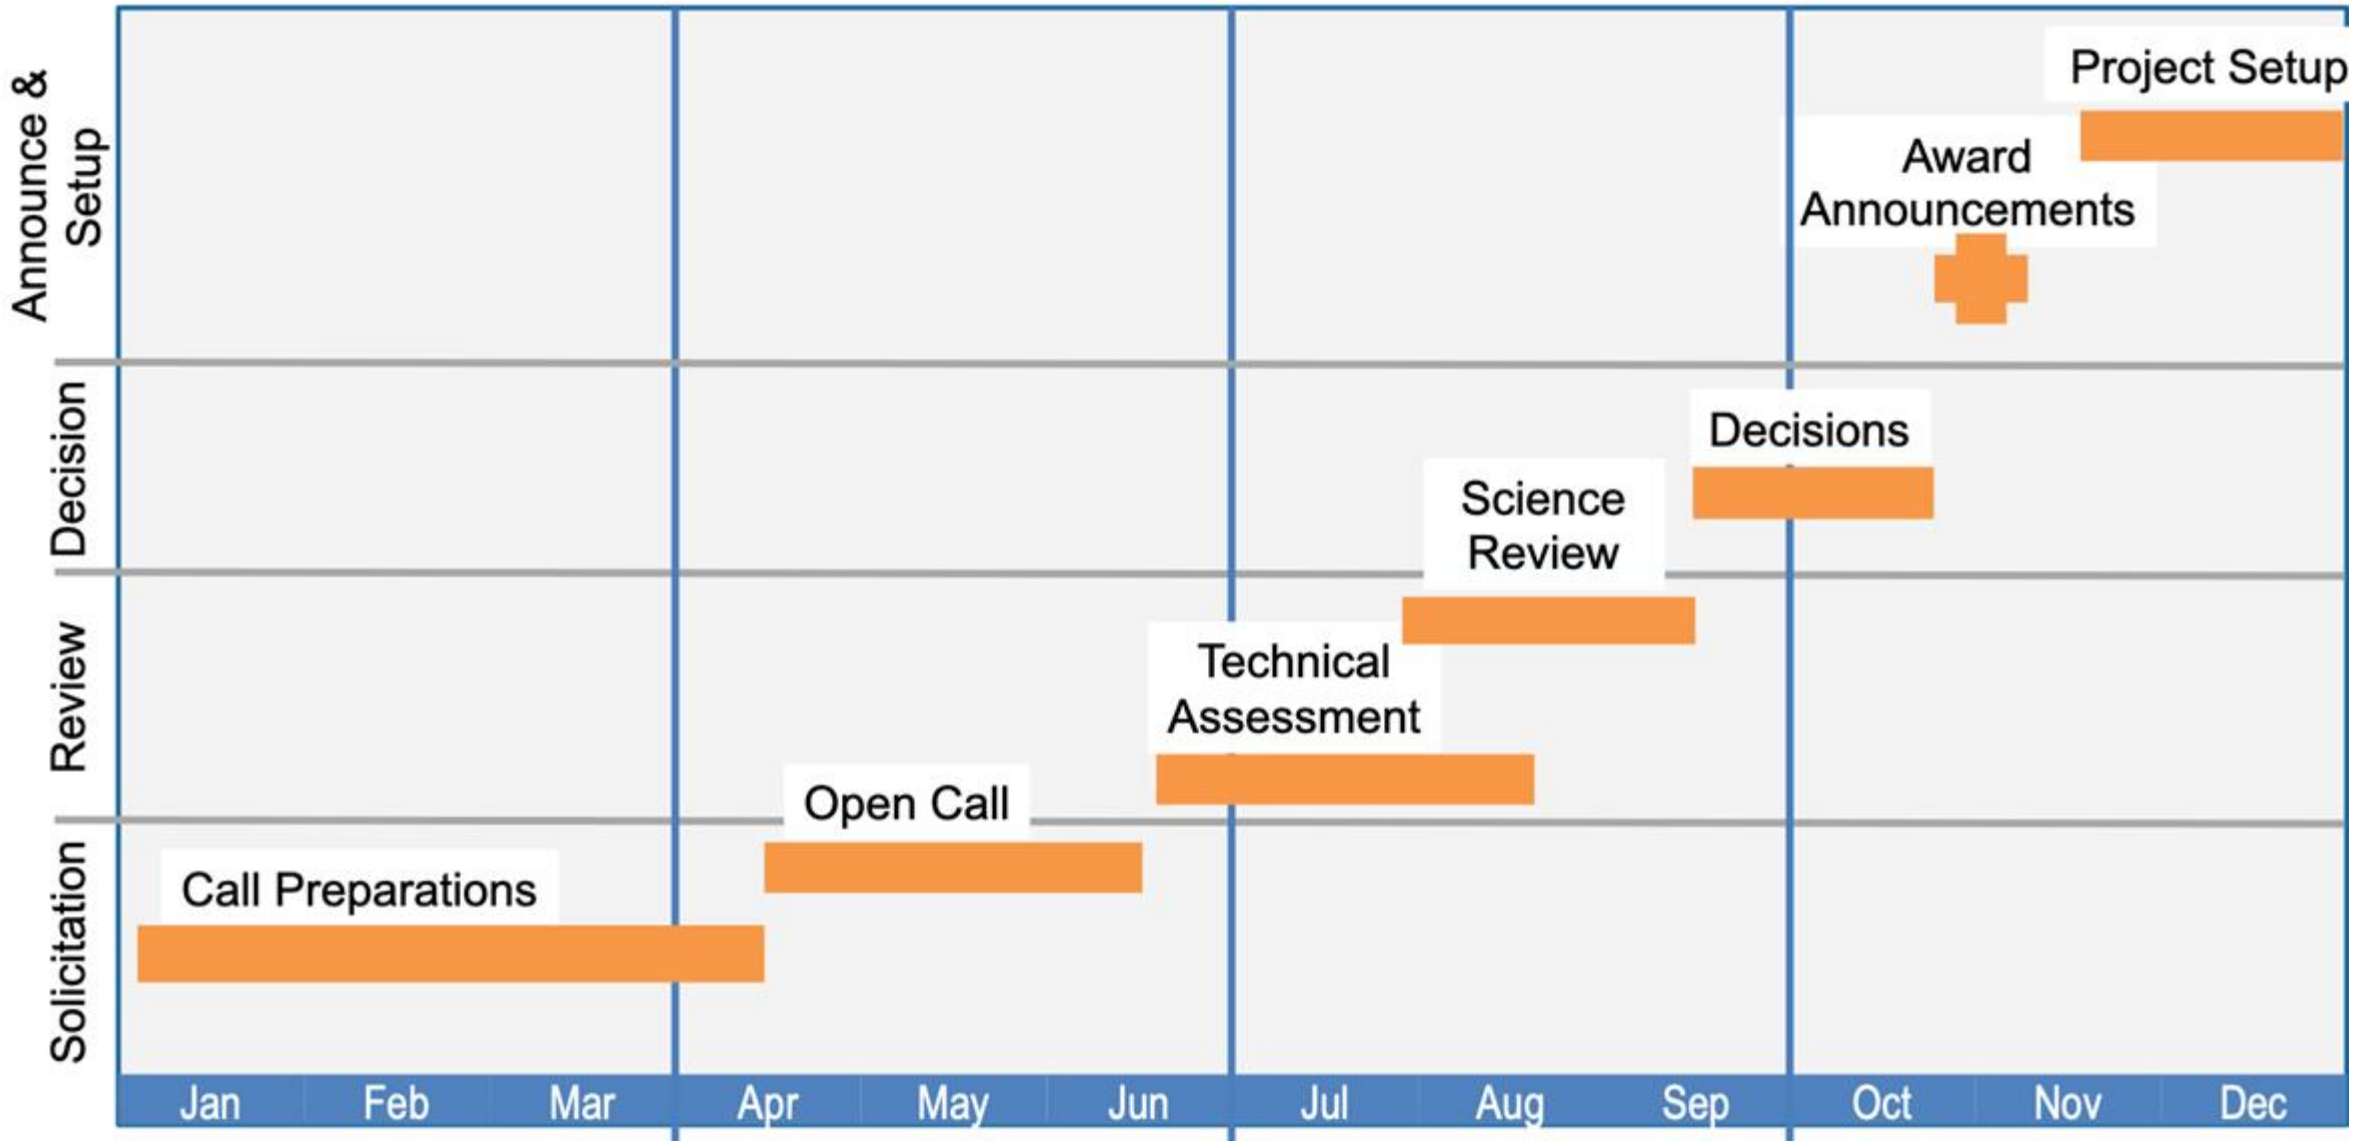

Are You Ready to Apply Now?

- 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

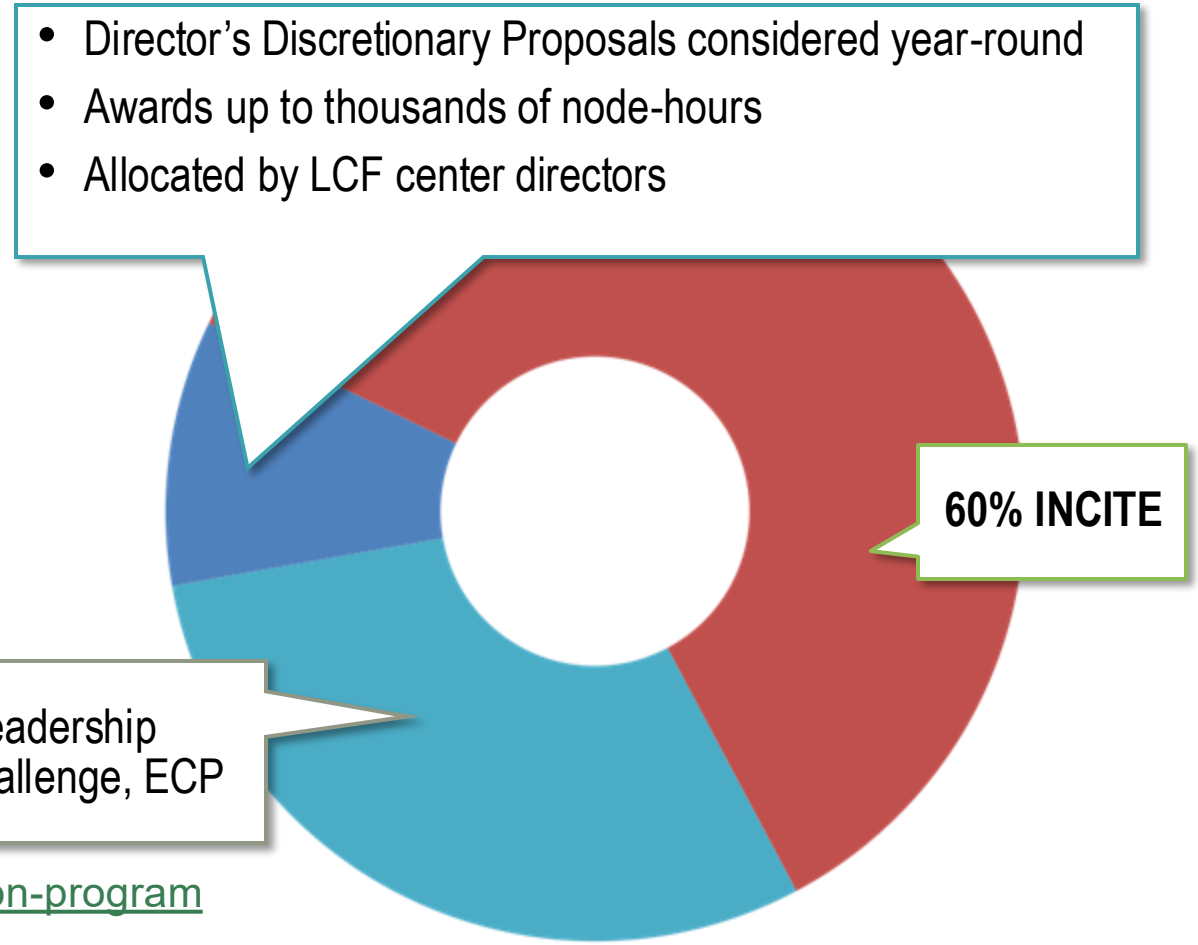

Argonne DD Program:
https://www.alcf.anl.gov/science/directors-discretionary-allocation-program

Oak Ridge DD Program:
https://my.olcf.ornl.gov/project-application-new

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

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