

AUG 2024

# ALCF-4 DESIGN REVIEW TECHNICAL REQUIREMENTS OVERVIEW



**KEVIN HARMS**

ALCF-4 Technical Director  
Argonne National Laboratory  
Leadership Computing Facility

**J. TAYLOR CHILDERS**

ALCF-4 Deputy Technical Director  
Argonne National Laboratory  
Leadership Computing Facility

# AGENDA

Times	Item	Owner
8:30	Executive Session	Review Chair
9:00	Welcome	Mike Papka
9:10	Project Overview	Jini Ramprakash
9:40	Technical Overview and Early Science	Kevin Harms Chris Knight
10:15	Break	
10:30	Technical Requirements	Taylor Childers
11:30	Benchmarks	Chris
12:15	(Working Lunch) Discussion & Questions from the committee	ALCF-4 Team
12:30	(Working Lunch) Executive Session	Review Chair
13:30	Facilities	Jon Cisek
14:15	ALCF-4 Risks Review	Noah / Jini
15:00	Break	
15:15	Executive Committee Q&A with ALCF-4 team	Review Chair
15:45	Executive Writing Session	Review Chair
17:00	Adjourn / Tour of Aurora	Susan Coghlan
18:00	Dinner	



# CHARGE QUESTIONS

1. Is the technical approach appropriate to support the ALCF-4 Mission Need requirements?
2. **Are the RFP technical requirements reasonable, clear, and consistent with the goals and objectives for the ALCF-4 project?**
3. Does the ALCF facility upgrade plan support the system requirements specified in the RFP for the onsite options?
4. Have the major technical risks and appropriate mitigation strategies been correctly identified for this stage of the project?

# PREFACE

- Both OLCF-6 and ALCF-4 build on text and procedure from the previous CORAL processes.
- The ALCF-4 Technical Requirements explicitly started with the text from the OLCF-6 Technical Requirements, which also followed the text of the NERSC-10 Technical Requirements.
  - The point is to provide vendors with uniform and easier to process RFP processes that facilitate responses with less confusion.
- ALCF-4 does have strategic and technical differences from OLCF-6, which will be highlighted here.



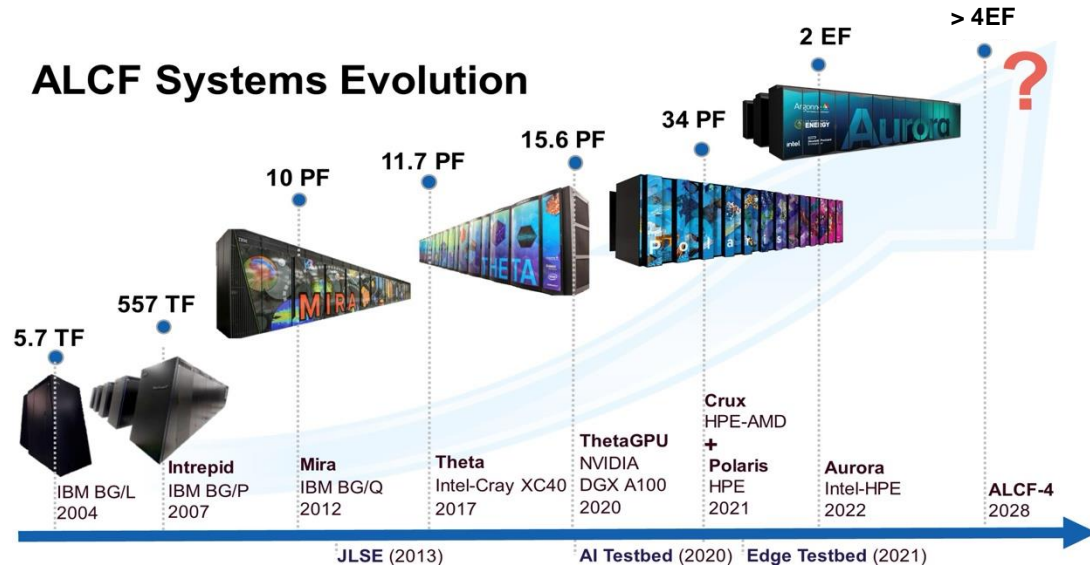
# GOAL FOR TECHNICAL REQUIREMENTS

- **MISSION** Outline the mission-critical drivers shaping the ALCF-4's capabilities and impact.
- **VISION** Define clear objectives for the ALCF-4 system, focusing on innovation and leadership in HPC.
- **SPECIFICATION** Detail the mandatory and optional components required to meet ALCF-4's objectives.
- **CLARIFICATION** Request comprehensive responses from vendors on crucial system design, performance metrics, and future-proofing strategies.



# MISSION DRIVERS

## ALCF Systems Evolution



- High-End Computing Revitalization Act of 2004 established the LCFs to provide access to high-performance computing resources on a competitive, merit-review basis.
- Despite ever growing computational capabilities, INCITE, our premier competitive allocation program, is consistently oversubscribed.
- Next system aims for 5-10x application performance over Aurora



# VISION FOR ALCF-4



Enable  
Scientific  
Discovery



AI for  
Science



Sustainable  
Computing



DOE as an  
Ecosystem



# VISION FOR ALCF-4



Enable  
Scientific  
Discovery

- This is always a priority for DOE and the LCFs
- "significantly increase performance of ModSim workloads over Aurora, e.g., time-to-solution and increased accuracy"
- "significantly increase AI performance over Aurora, e.g., FP16 FLOPS and model training time"
- "support scaling payloads to the full machine or applicable partition"



# VISION FOR ALCF-4



AI for  
Science

- “significantly increase AI performance over Aurora, e.g., FP16 FLOPS and model training time”
- “support novel architectures that provide leadership-class AI training and inference on scientific payloads”
- “The Company has structured these requirements with the objective of allowing Offerors to propose a wide range of solutions. This includes hybrid systems that include specialized AI architectures”
- “The Company expects ALCF-4 will be unlike any previous DOE system with a focus on advancing scientific artificial intelligence”

# VISION FOR ALCF-4



## Sustainable Computing

- “continue to make progress and lead in **dramatically improving energy efficiency** across the ecosystem (e.g., improving FLOP/Watt)”
- “The Company is interested in innovative solutions that offer rolling **future upgrades that reuse system infrastructure** in contrast to the current method of full facility reinstallations.”

# VISION FOR ALCF-4




DOE as an  
Ecosystem

- “The Company expects ALCF- 4 [...] will support **complex workflows** that merge in situ AI and simulation, and **connect ALCF-4 to the broader DOE facility landscape** through a versatile, tightly-integrated edge- service.”
- “Edge Service Nodes (ESN)” as Mandatory Requirement
- “**interoperate with a leadership-scale IRI**, defined as the ability to **interface with and support DOE user facilities** and other LCF infrastructure, e.g., through Edge Service Nodes and system connectivity”

# REQUIREMENT DEFINITIONS

- **Mandatory Requirements (MR):**  
These are non-negotiable features that vendors must meet to be considered. MR forms the foundation of the ALCF-4 system's capabilities.
- **Mandatory Options (MO):**  
Optional features that vendors must propose but are not required to be included in the final system. MOs provide flexibility in system design and future upgrades.
- **Target Requirements (TR):**  
These are highly desirable features that enhance the system's performance and usability. TRs are prioritized by level (TR-1, TR-2, TR-3), with TR-1 being the most critical.
- **Target Options (TO):**  
Optional enhancements that offer additional value. TOs allow for scalability, future upgrades, and customization based on the evolving needs of the ALCF-4 system.



Same classification as OLCF-6 to simplify for vendors.

# PROPOSAL PODIUM

## Proposal's that include:

Bronze:  
All MR + All TR-1



Silver:  
All MR, TR-1,  
and TR-2



Gold:  
All MR,  
TR-1, TR-2, TR-3



"A hexagonal, polished chunk of iron taken from the [Eiffel Tower] is being embedded in each gold, silver and bronze medal"

# EXAMPLE REQUIREMENT

## 6.1.5 External Connectivity [TR-1]

The Offeror shall describe the provision of a high-bandwidth and resilient solution that allows external connectivity to and from the system and the ALCF data center network. Describe how simultaneous transfers are supported and the aggregate throughput. **The Company prefers the aggregate throughput to be capable of at least 4 TB/s.** The Offer shall describe what protocols and interfaces would be supported.

- Most of our requirements do not demand metrics, but ask for descriptions of what the vendor is offering or how they accomplish a task.
- We will express our preferences to help vendors know what we are looking for.
- We will also give examples to help clarify for vendors what we are communicating.

# HIGH-LEVEL REQUIREMENTS [MR]



- Base Architecture Description
  - Details about Chips, Node, Racks, Network, Full Machine, Management systems, etc.
  - **Notable:** Edge Service Nodes are distinct from login nodes, as they are used for long-lived user and/or admin services, data in/egress pipelines, workflow management. Important for our IRI goals.
- Base Software Description
  - All major software components provided and how they support our benchmarks



# HIGH-LEVEL REQUIREMENTS [TR]



- FOMs for Benchmarks
- Total Memory (HBM, RAM, etc)
- Total FLOPS for each architecture and its supported data type
- Reliability, Serviceability, Upgradability, Privacy
- Software sharing, licensing, support
- **Notable: Rolling Upgradability** ALCF seeks partners who can offer rolling upgrades over a 10-15 year timeline, focusing on sustainability by reusing existing infrastructure while integrating the latest compute hardware. This approach reduces waste compared to the traditional model of replacing all infrastructure every 5-6 years.

# HIGH-LEVEL REQUIREMENTS [TO]

- System Quantization: allowing increase or decrease in overall size
- Early Access & Testing Systems
- Mid-life upgrade
- 6- and 7-year life extensions



# STORAGE & IO SYSTEMS

- Vendors propose Storage as an option to add on to their proposed System [MO]
- POSIX API Required
- Storage Architecture Details:
  - **Capacity:** 60x the aggregate system high performance memory; at least 100 billion entries
  - **Performance:** Describe Bulk, Random, and Small I/O Read/Write Performance
  - **Reliability;** Namespace Traversal
  - **Longevity:** Performance stability through end-of-life
- Storage Software Details:
  - Software support
  - Software Upgrades and Updates
  - System management
- Early Systems to support HPC bring up
- External connectivity to compute and company storage

# HIGH SPEED INTERCONNECT & NETWORKING

- Detailed Network description:
  - number, types, topology, interoperability
  - Bandwidths, latency, node-level connections (e.g., PCIe, NVLink)
  - Smart network capabilities (e.g. GPUDirect)
- Management, Reliability (Mean Time Between Failure), Quality of Service
- External Connectivity ( $\geq 4$  TB/s)
- Start-up readiness time for full system jobs
- Protocols Supported
- Provide details about low level communication API




# SYSTEM MANAGEMENT

- Hardware for System Management:
  - Describe all parts and their function
  - Storage to support node images and telemetry for 30 days
  - Out-of-band network for management system
- Software for System Management:
  - Software design and scalability
  - Configuration storage and management
  - Licensing
  - Uniform system management across all hardware
  - Resilience despite state of compute hardware
- Workload Management (e.g. PBS):
  - Scalable, topology aware job launching
  - Job launch performance metrics
- Operating System (RedHat or SLES required)
- Platform Management:
  - Remote, Scalable Management
  - Ensure no Single Points of Failure
  - Reliability, Health Testing
- System Software Deployment: Scalable Boot & Updates

# SECURITY

- Describe Security Model or Features in place to prevent unauthorized access.
- Describe Identity Management (e.g. LDAP)
- Vulnerability Reporting & Remediation
- Executive Order Compliance (Zero Trust Architecture)
- Secure Distributed Workflows: Confidentiality between user boundaries
- Privileged support access post-acceptance by US-persons only
- Encryption, Ability to Secure Secrets Across System





# FACILITIES

- Details of ALCF Data Center (TCS) space, structural, cooling, and utilities available for ALCF-4.
- Requests description of any requirements placed by the vendor's system on the facility infrastructure.
- Appendix A includes additional Facility details which the vendor is asked to review and point out problems.



# SYSTEM MONITORING

- Describe availability, accessibility, and security of the monitoring system proposed.
- Describe Logging Reliability and Redundancy (e.g. logging through network outages)
- Describe interfaces for searching, filtering, exporting, and other data manipulation.
- Describe monitoring frameworks to view system health, set alarms, etc.



# NON-REOCCURRING ENGINEERING (NRE)

- Describe any proposed NRE required for ALCF-4 to address mission needs.
- One of the few directives is to include support for Center of Excellence.
  - This has been an appreciated resource in helping early science usage success.

# PROJECT MANAGEMENT

- Describe management structure
- Agree to generate required documents
- Propose Payment Milestones
- Describe Risk Management
- Commit to Quarterly Reviews, setting up Working Groups, and Technical Collaborations
- Defines Acceptance framework



# USER ENVIRONMENT

- This is a large section that includes all the software users may come in contact with or use on ALCF-4 to get their science done.
- Python & AI Libraries
  - Describe support for standard & AI modules, differentiable programming frameworks (JAX)
  - Describe optimizations for AI & distributed training
- Software Packaging, Management, & Deployment
  - Describe plans for software management (e.g. environment modules)
  - Describe support for containers
- Compilers & Libraries
  - Describe support for C/C++/Fortran, Julia, language inter-operability
  - Timeline for N-1 test environment
  - Support for OpenMP/SYCL, Kokkos, std::par, MPI, Math & IO libraries
- Tools
  - Describe Profiler & Debug Support
  - Describe tools for energy monitoring and bug detection
- Workflows
  - Container Orchestration for Edge Service Nodes
  - Workflow enablement
  - Support for Streaming data, cloud integration, and robust task launching & management on the compute hardware

# SUPPORT & MAINTENANCE

- Require robust support for life of system
  - 2 system analysts, visit quarterly
- Replacement Parts Supply on-site
- Tracking support for software/hardware updates/failures
- Sub-contracted vendor support at similar levels







# SUMMARY

- **CHARGE Q2:** Are the RFP technical requirements reasonable, clear, and consistent with the goals and objectives for the ACLF-4 project?
- The ALCF-4 Technical Requirements closely follow the OLCF-6 and NERSC-10 texts in order to maintain consistency and ease of vendor response.
- The document is distinct from OLCF-6 and NERSC-10 in vision for the machine with a theme that speaks to our mission goals.
  - A machine that enables scientific AI leadership by encouraging innovative heterogeneous solutions that increase training time.
  - Support Integrated Research Infrastructure with the inclusion of:
    - Edge Service Nodes,
    - high bandwidth External Networking,
    - System Monitoring,
    - Task Management, and
    - Workflow support.





**Argonne**  
NATIONAL LABORATORY



**Argonne Leadership  
Computing Facility**