

AUG 2024

ALCF-4 DESIGN REVIEW TECHNICAL APPROACH



KEVIN HARMS
ALCF-4 Technical Director

CHRIS KNIGHT
Catalyst Team Lead

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?


TECHNICAL APPROACH

- ALCF building its approach on these core foundation
 - E6 Lab Request For Information
 - Vendor Meetings
 - Analysis of Industry Trends
 - ALCF User Survey
 - Alternatives Analysis
 - Early Science

REQUEST FOR INFORMATION

- E6 Labs released RFI for next generation systems June 2022
- Covered broad range of technologies
 - See chart at right
- 49 responses
- Broad range from startup companies with fewer than 50 employees to major multinational corporations
- Follow-up meetings held with Vendors who requested
 - Following companies were among those identified as being interested in the pending RFP: Advanced Micro Devices, ARM, Atos, Cornelis Networks, DDN Storage, Dell Technologies, Esperanto Technologies, Google Cloud, Hewlett Packard Enterprise, IBM, Intel, Microsoft, NextSilicon, NVIDIA, Penguin Computing, Qualcomm, Rockport Networks, SambaNova Systems, Samsung, SiPearl, Supermicro, Tachyum, Untether AI, and VAST Data

RFI Responses by Product Category




- Interest from CSPs
- 2028: Evolutionary rather than Revolutionary

VENDOR MEETINGS

- Meetings with vendors to review ALCF technical specifications and goals
 - Receive feedback and discuss potential technologies
- Results
 - No major issues identified
 - Still too early to get solid information on compute technology in 2028
 - But we have trends
- Nvidia – August 1, 2024
- AMD – August 12, 2024
- HPE – August 14, 2024
- MS – September 12, 2024
- AWS – October 16, 2024
- Penguin – October 24, 2024
- Dell – TBD

INDUSTRY TRENDS (COMPUTE)

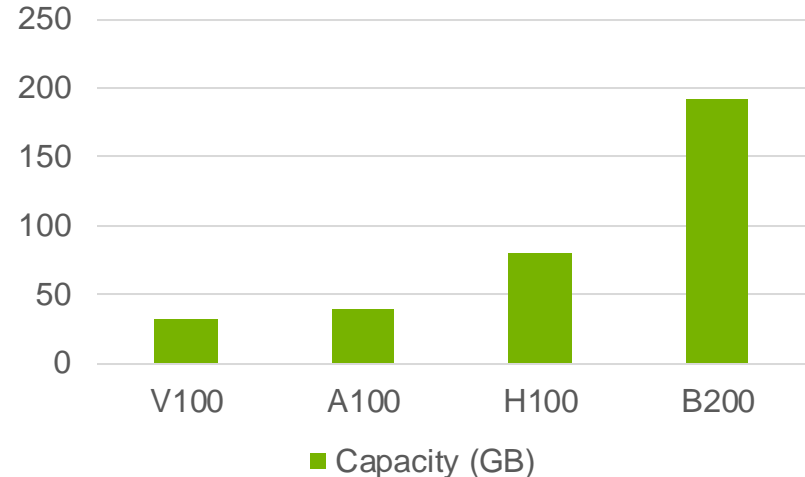
- Compute Capability
 - AI and power constraints pushing accelerators toward reduced precisions below 64-bit floating point
- AI
 - acceleration is excellent
- Mod/Sim
 - Modify codes to take advantage of reduced precision
 - Emulation of higher precision with lower precisions



INDUSTRY TRENDS (MEMORY)

- Absolute memory capacity attached to accelerators continues to increase
- Capability for memory capacity is becoming similar to main memory sizes
- Look for solutions to optimize usage of accelerator memory and reduce dedicated CPU memory
 - CPU shares accelerator memory to save cost of CPU memory


Nvidia GPU Memory



INDUSTRY TRENDS (POWER)

- Moore's Law and Dennard scaling have run out
 - Pushing more power through a single socket to grow performances
 - Complex power and cooling delivery systems required
- Close relationship with Integrators to understand future requirements for cooling and power systems


Compute Power (W)




USER SURVEY

- ALCF issues a survey users every calendar year which looks at ALCF performance but also collects information about user codes
 - Informs needs for future systems

WHICH OF THE FOLLOWING MACHINE LEARNING AND DEEP LEARNING TOOLS DO YOU USE ON YOUR LAPTOP, CLUSTER-BASED SYSTEMS, AND/OR ON ALCF SYSTEMS?



WHICH ACCELERATOR PROGRAMMING MODEL IS YOUR CODE CURRENTLY USING OR PLANNING TO USE?



ALTERNATIVES ANALYSIS



Argonne National Laboratory is a
U.S. Department of Energy Laboratory
managed by UChicago Argonne, LLC.



Argonne Leadership
Computing Facility

ALTERNATIVES

- **Continued operation of Aurora**
 - Continue operating Aurora for an additional five years, 2030-2034
- **Upgrade Aurora**
 - Upgrade the compute blades of Aurora using HPE blades from 2025
 - Potentially management server and storage upgrades
 - Rack infrastructure and interconnect remain the same
- **Off-premises Cloud Computing**
 - Contract with cloud vendor to use compute cycles from the cloud
- **On-premises Cloud Computing**
 - Deploy a dedicated system within ALCF but uses cloud infrastructure
 - System is rented rather than owned
- **New System**

CRITERIA

- **Provide a significant increase in leadership computational and data science capabilities over the Aurora baseline**
- **Interoperate with leadership scale IRI**
- **Improve energy efficiency across the ecosystem**
- **Operational by 2030 before Aurora's planned end-of-life**
- **Operate within the utility and operations budget**
 - 40MW power and cooling
- **Provide a productive programming environment to users**
- **Strategic value to DOE**
 - ALCF should provide a leadership computing facility

CONTINUE OPERATION OF AURORA

- Increase Capabilities
 - Maintains the same compute as Aurora given nothing is changed
 - Major risk that hardware can continue to operate or that replacements are available
- Support IRI
 - Aurora was designed with support for good external connectivity
 - Software stack is static
- Improve Energy Efficiency
 - No improvements as no changes
- Utility & Operations Budget
 - Using 60MW rather than 40MW
- Programming Environment
 - Static since support for PVC will be EOL

Criteria	Rating
Increase Capabilities	Fails to Meet
Support IRI	Partially Meets
Improve Energy Efficiency	Fails to Meet
Operational 2030	Fully Meets
Utility & Operations Budget	Fails to Meet
Programming Environment	Partially Meets
Strategic Asset	Fully Meets

UPGRADE AURORA

- Increase Capabilities
 - Significant FLOPs can be added by upgrade 10k blades with MI300A or GH200 or GB100 blades
 - Slingshot will severely limit scaling and underutilize compute
 - Major risk for ability to purchase 2025 blades in 2028
- Support IRI
 - New compute software allows some flexibility in newer software stacks
 - Still constrained by original Aurora design
- Utility and Operations Budget
 - Assumes upgrading all 10k blades running at 60MW beyond the 40MW planned
- Severe impact to the overall facility operation as Aurora would need to be offline during the upgrade which maybe be 1-2 years

Criteria	Rating
Increase Capabilities	Partially Meets
Support IRI	Partially Meets
Improve Energy Efficiency	Fully Meets
Operational 2030	Fully Meets
Utility & Operations Budget	Fails to Meet
Programming Environment	Fully Meets
Strategic Asset	Fully Meets

OFF-PREMISES CLOUD

- Increase Capabilities
 - Traditional cloud allocate ephemeral resources, no guarantee for large scale jobs that they are connected with reasonable all-to-all network
 - Not bound to a single compute type
- Support IRI
 - Great support for remote connectivity and various services
 - Local resources are no longer local (APS, ALCF storage system, ALCF support clusters)
- Strategic Asset
 - Not deployed within Argonne
 - Risk that staff leave/move to vendor and knowledge base is lost, potentially forever tied to "cloud" solutions
- The cost of the traditional cloud services is very unclear, using published costs, the solutions would be in the 3-10 billion dollars but vendors claim custom pricing would be available

Criteria	Rating
Increase Capabilities	Partially Meets
Support IRI	Adequately Meets
Improve Energy Efficiency	Fully Meets
Operational 2030	Fully Meets
Utility & Operations Budget	Fully Meets
Programming Environment	Fully Meets
Strategic Asset	Fails to Meet

For **Improve Energy Efficiency** and **Utility & Operations Budget**, these components are "hidden" and so assumed to be Fully Meets but we can not verify them.

ON-PREMISES CLOUD

- Strategic Asset
 - Machine is sited within Argonne but is still a rented resource which Argonne does not own
- Some risk to understand how the payment model works

Criteria	Rating
Increase Capabilities	Fully Meets
Support IRI	Fully Meets
Improve Energy Efficiency	Fully Meets
Operational 2030	Fully Meets
Utility & Operations Budget	Fully Meets
Programming Environment	Fully Meets
Strategic Asset	Adequately Meets

NEW SYSTEM

- Solution built exactly to the Technical Specifications
- Sited at Argonne
- Purchased through lease to own

Criteria	Rating
Increase Capabilities	Fully Meets
Support IRI	Fully Meets
Improve Energy Efficiency	Fully Meets
Operational 2030	Fully Meets
Utility & Operations Budget	Fully Meets
Programming Environment	Fully Meets
Strategic Asset	Fully Meets

ALCF-4 ALTERNATIVES ANALYSIS

	Continue Aurora	Upgrade Aurora	Off-Prem Cloud	On-Prem Cloud	New System
Increase Capabilities	FAILS TO MEET	PARTIALLY MEETS	PARTIALLY MEETS	FULLY MEETS	FULLY MEETS
Support IRI	PARTIALLY MEETS	PARTIALLY MEETS	ADEQUATELY MEETS	FULLY MEETS	FULLY MEETS
Improve Energy Efficiency	FAILS TO MEET	FULLY MEETS	FULLY MEETS	FULLY MEETS	FULLY MEETS
Operational 2030	FULLY MEETS	FULLY MEETS	FULLY MEETS	FULLY MEETS	FULLY MEETS
Utility & Operations Budget	FAILS TO MEET	FAILS TO MEET	FULLY MEETS	FULLY MEETS	FULLY MEETS
Programming Environment	PARTIALLY MEETS	FULLY MEETS	FULLY MEETS	FULLY MEETS	FULLY MEETS
Strategic Asset	FULLY MEETS	FULLY MEETS	FAILS TO MEET	ADEQUATELY MEETS	FULLY MEETS

EARLY SCIENCE PROGRAM (ESP)



Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.



Argonne Leadership
Computing Facility

SUMMARY

- ALCF has taken a measured technical approach to the ALCF-4 project
 - Evaluate the market and available technologies
 - Engage with vendors on key issues relative to achieving performance goals within the project scope
 - Ensure ALCF applications software transition to ALCF-4 with appropriate support for software dependencies
 - Leverage successful ESP program to guide software transition
- Proceed with competitive procurement for ALCF-4 which allows for both on-premises cloud and new system sited at Argonne within the ALCF



Argonne
NATIONAL LABORATORY



**Argonne Leadership
Computing Facility**