

New Frontiers TECHNICAL REQUIREMENTS



March 2024



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

Oak Ridge Leadership Computing Facility

**NEW FRONTIERS
REQUEST FOR PROPOSALS**

September 2024

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-Battelle LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

ABBREVIATIONS	ii
ABSTRACT	1
1. INTRODUCTION	1
1.1 PROGRAM OVERVIEW	1
1.1.1 OFFICE OF SCIENCE	1
1.1.2 DOE MISSION NEED	1
1.2 RESEARCH AND DEVELOPMENT THEMES	2
1.3 RFP INFORMATION	2
1.3.1 Who May Respond	3
2. RESEARCH THEMES	5
2.1 HARDWARE TECHNOLOGY	5
2.1.1 Processor and Memory Advancements	5
2.1.2 Disaggregated/Heterogeneous Computing	5
2.1.3 Next Generation Storage Architectures	6
2.1.4 HPC Network Advancements	6
2.2 SOFTWARE TECHNOLOGY	7
2.2.1 Software Complementing Novel Hardware	7
2.2.2 Compilers, software stacks, programming models	7
2.3 CROSS-CUTTING TECHNOLOGY	8
2.3.1 Co-design of algorithms, software, and hardware	8
2.3.2 Energy Efficiency	8
2.3.3 Reliability & Resiliency	8
2.3.4 Thermal Management	8
2.3.5 Quantum Computing	8
2.3.6 Integrated Research Infrastructure	9
3. REQUIREMENTS	10
3.1 DESCRIPTION OF REQUIREMENT CATEGORIES	10
3.2 REQUIREMENTS FOR R&D INVESTMENT AREAS	10
3.3 COMMON MANDATORY REQUIREMENTS	10
3.3.1 Solution Description (MR)	10
3.3.2 Research and Development Plan (MR)	10
3.3.3 Staffing/Partnering Plan (MR)	11
3.3.4 Project Management Methodology (MR)	11
3.3.5 Intellectual Property Plan (MR)	11
3.3.6 Productization Strategy (TR)	11
4. EVALUATION	12
4.1 EVALUATION TEAM	12
4.2 EVALUATION FACTORS AND BASIS FOR SELECTION	12
4.3 PERFORMANCE FEATURES	12

ABBREVIATIONS

AI	Artificial Intelligence
ASCR	DOE Advanced Scientific Computing & Research
DOE	Department of Energy
DRAM	Dynamic Random Access Memory - used in DDR and HBM
FLOP	Floating Point Operation - a measure of compute performance
HBM	High Bandwidth Memory
HPC	High Performance Computing
ML	Machine Learning - a type of AI
Mod/Sim	Modeling & Simulation - traditional scientific/technical computing
OLCF	Oak Ridge Leadership Computing Facility
ORNL	Oak Ridge National Laboratory
RFI	Request For Information
RFP	Request For Proposals
SC	DOE Office of Science
SME	Subject Matter Expert

1. INTRODUCTION

This document contains the request for proposals for the New Frontiers industry research program to enable the next generation of exascale supercomputers beyond 2028. This RFP is being led by the Oak Ridge National Laboratory (ORNL) and the Oak Ridge Leadership Computing Facility (OLCF). This research is required to fulfill the mission needs of the Advanced Scientific Computing Research (ASCR) Program within the Department of Energy (DOE) Office of Science (SC). ORNL is managed by UT-Battelle, LLC, referred to in this document as THE COMPANY.

1.1 PROGRAM OVERVIEW

1.1.1 OFFICE OF SCIENCE

The DOE Office of Science (SC) is the lead Federal agency supporting fundamental scientific research for energy and the nation's largest supporter of basic research in the physical sciences. The SC portfolio has two principal thrusts: direct support of scientific research and direct support of the development, construction, and operation of unique, open-access scientific user facilities. These activities have wide-reaching impact. SC supports research in all 50 States and the District of Columbia, at DOE laboratories, and more than 300 universities and institutions of higher learning nationwide. The SC user facilities provide the Nation's researchers with state-of-the-art capabilities that are unmatched worldwide.

Within SC, the mission of the Advanced Scientific Computing Research (ASCR) program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the DOE. A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures, which will require numerous modifications to today's tools and techniques to deliver on the promise of science in the exascale era.

1.1.2 DOE MISSION NEED

High-performance computing (HPC) and data-driven modeling and simulation are used extensively in advancing DOE missions in science. To maintain leadership and to address the future challenges in science, energy, health, and growing security threats, the United States must continue to push strategic advancements in HPC—bringing about a grand convergence of modeling and simulation, data analysis, workflows, deep learning, artificial intelligence (AI), quantum computing, and other emerging capabilities—across integrated infrastructures in computational ecosystems. New approaches to predictive analysis for scientific discovery and solutions to complex data-driven engineering problems will arise from this convergence.

The DOE exascale systems (i.e., Frontier, Aurora, and El Capitan) are designed to address this emerging convergence. They can run simulations that require the entire platform and take days to weeks to complete. The AI-driven approaches to these systems will be used to perform uncertainty quantification and to discover complex, non-linear relationships in the output of extensive multi-physics simulations and large science experiments. The new capabilities of these systems will revolutionize scientific areas, such as energy production, materials design, chemistry, precision health care, advanced manufacturing, stockpile stewardship, and national security.

For the past decade, the six science programs in the DOE Office of Science have formulated strategic plans

for the disciplines that they steward. These plans rely on HPC in ever-increasing proportion, and, in recent years, the explicit call for HPC at exascale performance levels has been a common and defining theme. Examples include discovery and characterization of next-generation materials; systematic understanding and improvement of chemical processes; analysis of the extremely large datasets resulting from the next generation of particle-physics experiments; and extraction of knowledge from systems-biology studies of the microbiome. Advances in applied energy technologies also are dependent on next-generation simulations, notably whole-device modeling in plasma-based fusion systems. The Exascale Computing Project (ECP) has developed a portfolio of applications and technologies at exascale that will use the DOE exascale systems, while benefiting next-generation systems.

1.2 RESEARCH AND DEVELOPMENT THEMES

In August of 2022, the Oak Ridge National Laboratory, in collaboration with the other ASCR laboratories, released a request for information (RFI) with the purpose of providing DOE SC with information for planning of future DOE HPC Programs. The RFI responses highlighted numerous innovative ideas to address future HPC challenges. Informed by responses to the RFI, the DOE facilities drafted several, broad themes for strategic research and development (R&D) investment that will provide benefit to future extreme-scale applications:

- Hardware Technologies
 - Processor and Memory Advancements
 - Disaggregated and Heterogeneous Computing
 - Next Generation Storage Architectures
 - HPC Network Advancements
- Software Technologies
 - Compilers, Software Stack, and Programming Environments
- Cross Cutting
 - Co-Design of Algorithms, Software, and Hardware
 - Energy Efficiency
 - Reliability and Resiliency
 - Thermal Management
 - Quantum Computing
 - Integrated Research Infrastructure

1.3 RFP INFORMATION

These R&D activities will initially be pursued by THE COMPANY through a program called New Frontiers. The objective of the New Frontiers program is to initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for the next generation of exascale computing. It is recognized that the broader computing market will drive innovation in directions that may not meet DOE's

mission needs. Many DOE applications place extreme requirements on computations, data movement, and reliability. New Frontiers seeks to fund innovative new and/or accelerated R&D of technologies targeted for productization in the 5–10-year timeframe. The period of performance for any subcontract resulting from this request for proposal (RFP) will be two years.

THE COMPANY is soliciting innovative R&D proposals in the areas of hardware technologies, software technologies, and cross-cutting technologies. Due to the focus on extreme-scale application workflows, overall time to solution is also an important consideration. Software technologies will focus on open-source and sustainable software technologies for extreme scale HPC system and the development of techniques necessary to support emerging workloads of integrated facilities across the DOE landscape. The goal is to begin to address long-lead time items that will impact extreme-scale DOE systems later this decade. Technology roadmaps, as they exist today, threaten to have a hugely disruptive and costly impact on development of DOE applications and ultimately a negative impact on the productivity of DOE scientists.

1.3.1 Who May Respond

Responders may include any vendor of hardware or software that could be deployed in a future advanced computing system. Examples include, but are not limited to, basic hardware component (e.g., memory, processors, interconnects, storage media) vendors, system (e.g., compute, storage) integrators, and software (e.g., compilers, system management, storage management) vendors. Any company that provides hardware or software technologies relevant to delivering multi-exaflop supercomputing may respond.

Proposals submitted in response to this solicitation must address the impact of the proposed R&D on the next generation of exascale machines. While DOE's extreme-scale computer requirements are a driving factor, these projects must also exhibit the potential for technology adoption by other DOE facilities as well as the broader segments of the market outside of DOE supercomputer installations. This public-private partnership between industry and THE COMPANY, initiated with New Frontiers, will aid the development of technology that reduces economic and manufacturing barriers to constructing extreme-scale systems, but also further THE COMPANY's goal that the selected technologies have the potential to impact low-power embedded, cloud/datacenter, and midrange HPC applications. This ensures that DOE's investment furthers a sustainable software/hardware ecosystem supported by applications across not only HPC but the broader IT industry. This will result in an increase in the DOE's ability to leverage commercial developments.

We ask vendors to limit each submission to 50 pages or less.

In line with the BUY AMERICAN ACT, CHIPS ACT [3], and US Supply Chain Act [2]. All respondents must be US-based or US-subsidary companies. Responses from companies that do not meet these requirements will not be reviewed.

1.3.1.1 Inappropriate Proposals

Any proposal with any of the following criteria will not be accepted:

- Requests to fund engineering for items that are already on the Offeror's product roadmaps,
- Requests to fund venture capital,
- Requests to provide endorsements to the Offeror to aid in obtaining venture capital,

- Requests to fund fundamental quantum design and/or engineering. Quantum proposals should be narrowly tailored to integrating quantum systems with traditional HPC Mod/Sim systems, and
- Proposals from non-US companies as outlined above.
- Submissions exceeding 50 pages in total length.

2. RESEARCH THEMES

The following sections reflect the research themes identified during DOE’s Advanced Computing Ecosystem Request For Information in 2022. These are meant to be broad umbrellas that will enable the widest possible set of proposals for technology that will enhance future DOE computing capabilities.

2.1 HARDWARE TECHNOLOGY

This section focuses on several hardware-specific guidelines; however, the following list is not comprehensive

2.1.1 Processor and Memory Advancements

- Need for architectures that accelerate application/workload time to solution
 - Enable strong scaling of a variety of workloads – from dense to sparse, highly regular to significant branching, FLOPs rich to FLOPs ambivalent
- Packaging advancements (e.g., 2.1d -> 2.5d -> 3d)
 - Memory, processor, and accelerator packaging improvements for improved performance, particularly strong scaling
 - Logic / Memory / NVM stacking and integration
- Understand and exploit fabrication process node improvements
- Memory technologies to dramatically reduce the memory wall
 - Logic 3D SRAM or MRAM, Logic near/in DRAM
 - Memory-centric computing, process near memory, process in memory
 - Memory gather, scatter, and other optimized sparse memory operations
- Novel architectures
 - AI/ML accelerators (training & inference)
 - From static systolic arrays to coarse grain reconfigurable architectures
 - Generalized data flow architectures for HPC modeling & simulation

2.1.2 Disaggregated/Heterogeneous Computing

- Enable more efficient use of resources through composable infrastructure
 - Standards to allow integration within/between packages and vendors (e.g., UCle, CXL) — Ability to compose resources at both a fine and coarse level
 - HPC Networks, Network on a Chip, & Silicon Photonics technologies that support high performance, low-latency, low-power compositions of resources
- Architecture and Packaging advancements that support a tight integration of a variety compute and memory technologies involving different vendor’s IP

- Chiplets
- Multipackage & full reticle integration
- Modularity, upgradability, serviceability
 - Balancing the benefits of tight integration with that of field replaceable unit sizes

2.1.3 Next Generation Storage Architectures

- Coupling high-performance storage and networking advancements (e.g., DPUs, Storage accelerators)
- Developing Advanced Computational storage technologies
 - Data/format agnostic offloads to improve efficiency of data movement to/from storage
 - Offloading of analytics enabling capabilities such as indexing

2.1.4 HPC Network Advancements

- Technologies that improve latency and bandwidth at lower power (e.g., Silicon Photonics)
 - On package integration
 - Cache / memory coherent interfaces / protocols
 - ISA and microarchitecture advancements to better support network / processor coordination
- Improvements in network messaging rates
 - Improve strong scaling when surface of communication to volume of compute increases
- Reduced latency and pipeline stalls for heterogeneous device communications
 - These may include CPU to GPU, GPU to GPU, CPU to Accelerator
- Improvements in network telemetry
 - Providing applications and system software with actionable information (e.g., load balancing, process placement, scheduling)
- Congestion management and reduction of tail latency
- Novel uses for SmartNICs/DPUs

2.2 SOFTWARE TECHNOLOGY

This section focuses on software-specific guidelines; however, the following list is not comprehensive.

2.2.1 Software Complementing Novel Hardware

- Software relating to novel hardware capabilities proposed in Sec. 2.1.
- Software relating to other novel hardware capabilities (not proposed).

Software proposed under this section should either be required for the effective creation of computational science and engineering applications and supporting computational infrastructure using the target hardware, or should provide significantly enhanced capabilities.

2.2.2 Compilers, software stacks, programming models

- Feature complete open-source and sustainable modern Fortran implementation with support for modern accelerators
- Automated performance analysis and optimization tools:
 - Tools that automate performance analysis of programs running on accelerators e.g., by automating the collection of performance data from accelerators and host system hardware counters, by automating the analysis of the data determining how the programs used the available hardware resources.
 - Tools to provide optimization suggestions to programs based on multi-faceted models incorporating better performance analysis, source code analysis, and hardware usage. Optimization suggestions would enable the original code to make better use of the available compute devices.

Any such tools should not be limited to specific hardware, but instead should be as general as possible, based, for example, on parametric hardware performance models or other approaches.

Performance analysis and optimization may encompass the analysis and optimization of power consumption. Tools may include compiler-based tools to generate code optimized for power consumption.

- Libraries, primitives, and associated compiler technology for execution on disaggregated and heterogeneous components at the node and component level or at the system level.

2.3 CROSS-CUTTING TECHNOLOGY

This section focuses on Cross-Cutting related themes; however, the following list is not comprehensive.

2.3.1 Co-design of algorithms, software, and hardware

Combined hardware and software techniques for optimizing challenging problems (Sparse Matrix, FFT)

2.3.2 Energy Efficiency

Enable technologies for application driven or application tailored system power management to maximize energy efficiency of HPC applications. Pursue hardware and software co-design to achieve software defined power management that aggressively addresses system level dynamic power and idle power consumption driven by applications needs or workload characteristics. Provide means to accurately account and monitor energy efficiency of HPC applications.

- Enable component level and server level hardware capabilities that provide deep power management towards addressing both dynamic and idle power consumption of the HPC system
- Middleware, runtime, and APIs aiming for software defined energy efficiency driven by either applications themselves and/or intelligent controllers that can maximize energy efficiency at the level of each application run (i.e., job) that involves multiple nodes and components
- Sensors and hardware counters for low overhead monitoring and accounting of both component and server level power and energy consumption accompanied with application performance metrics and on die resource utilization to associate with energy consumption (i.e., floating point arithmetic units (FP16, FP32, FP64), memory capacity & memory usage, interconnect I/O bandwidth, storage I/O bandwidth))
- Software ecosystem that provides fine-grained accounting and monitoring of application energy efficiency targeting user driven application runtime monitoring and system administrative use cases

2.3.3 Reliability & Resiliency

- Improve the overall reliability of extreme scale systems

2.3.4 Thermal Management

- Software based thermal regulation techniques working in conjunction with cooling infrastructure to enable the use of 40C cooling technologies

2.3.5 Quantum Computing

Quantum computing is an emerging paradigm that relies on novel hardware and algorithmic methods for performing computation using the precepts of quantum mechanics. Recent advances in quantum computing devices and systems raise questions about the feasibility and readiness of this technology for solving computational problems. The following research themes are identified to assess the feasibility and accelerate the ability to integrate quantum computing resources with classical computing resources:

- Design and architecture of quantum computing software-hardware interfaces for managing interactions between quantum computing systems and classical computing systems
 - For example, networking, security, and data management between quantum computers and HPC systems
- Developing and profiling applications that integrate quantum computing and HPC
 - For example, demonstrating time, power, and accuracy metrics for hybrid use cases in specific applications domains
- Develop and demonstrate algorithmic and software methods for quantum computing system management
 - For example, methods for multi-user management and resource provisioning and scheduling, i.e., operating systems for quantum computers.

2.3.6 Integrated Research Infrastructure

Work on integrating scientific instruments both within and outside of the DOE Laboratory space is a major strategic initiative that will play an major role in upcoming facility projects. This strategic priority is called the Integrated Research Infrastructure [1]. Vendor contributions to enabling this vision of more productive science by creating an environment in which major scientific tools can be improved through tight coupling with extreme scale super computing is highly desired.

- Enable technologies for scheduling resources across multiple different user facilities
- Develop technologies for on-demand and efficient scientific data streaming services
- Data backplane services for pipelining and caching data across the infrastructure based on access patterns and look-ahead
- Integration of flexible, on-demand, and portable computational applications with large-scale scientific HPC environments
 - Reliable, Flexible, and easy to use software for bridging network connections between lab ESNet connections and HPC RDMA fabrics
- Edge side compression and decompression/data reduction techniques, using hardware acceleration to enable peta-scale data experiments between connected facilities
- Scalable workflow compilation and performant execution across integrated research infrastructures

3. REQUIREMENTS

3.1 DESCRIPTION OF REQUIREMENT CATEGORIES

Requirements are either mandatory (designated MR) or target (designated TR), and are defined as follows:

- MRs are performance features essential to THE COMPANY's requirements. An Offeror must satisfactorily address all MR to have its proposal considered responsive.
- TRs, identified throughout this Statement of Work, are features, components, performance characteristics, or other properties that are important to THE COMPANY but will not result in a non-responsive determination if omitted from a proposal.

3.2 REQUIREMENTS FOR R&D INVESTMENT AREAS

Detailed requirements for each of the targeted R&D areas of investment are discussed in this document. **A single proposal may address multiple areas of investment, that is, an Offeror need not submit a unique proposal for each area of investment on which it chooses to propose.** Each proposal shall address all of the common MRs listed below. All of the MRs in each area of investment shall be included in the proposal.

3.3 COMMON MANDATORY REQUIREMENTS

The following items are mandatory for all proposals. That is, they must be present in any proposal for that proposal to be considered responsive and eligible for further evaluation.

3.3.1 Solution Description (MR)

Offeror shall describe the proposed R&D, with emphasis on how it will provide improvement in the targeted R&D area. Offerors shall discuss the innovative nature of the proposed R&D. Work that funds engineering for a company's current roadmap is not acceptable. Technology acceleration is acceptable if there is a clear benefit, and it is part of a broader strategy. The primary intent is to fund long-lead-time R&D objectives where significant advances can be made during the term of this program.

3.3.2 Research and Development Plan (MR)

Offeror shall provide a plan for conducting the proposed R&D, including timelines, milestones, and proposed deliverables. Deliverables shall be meaningful and measurable. Pricing shall be assigned to each milestone and deliverable. A schedule for periodic technical review by the DOE laboratories shall also be provided.

Some projects may develop a hardware prototype that demonstrates the value of the proposed concept. Others may perform a simulation or analysis that assesses the impact (or feasibility) of a proposed development. If funding provided through this RFP is insufficient to effectively demonstrate a concept or produce a prototype, Offerors shall provide a separate, non-binding budgetary estimate for follow-on work that would be needed to achieve this result. Do NOT include the estimated amount for this activity in the price for the R&D being proposed in response to this RFP. This follow-on work could be proposed in response to a future RFP, if one is issued.

We recognize that innovation involves risk. Proposals shall discuss technical and programmatic risk factors and the strategy to manage and to mitigate risk. If the planned R&D is not achieving the expected results, what alternatives will be considered? The amount of risk must be commensurate with the potential impact. Higher risk projects may be acceptable if the impact of the project is also high.

3.3.3 Staffing/Partnering Plan (MR)

Offerors shall describe staffing categories and levels for the proposed R&D activities. Any collaboration with other industry partners and/or universities shall be identified.

3.3.4 Project Management Methodology (MR)

Offerors shall provide a project management plan. The COMPANY expects quarterly milestone status updates on selected work packages.

3.3.5 Intellectual Property Plan (MR)

Proposals shall include a plan for how each intellectual property (IP) item from each portion of the proposed R&D work will be handled, including requested IP ownership and licensing. Traditionally, Government retains all rights unless collaborating vendors contributed 40% of the total to the work project.

3.3.6 Productization Strategy (TR)

Offeror shall describe how the proposed technology will be commercialized, productized, or otherwise made available to customers. Offerors shall include identification of target customer base/market(s) for the technology. Offerors shall describe impact specifically on the HPC market as well as the potential for broad adoption. Solutions that have the potential for broader adoption beyond HPC are highly desired. Offerors shall indicate projected timeline for productization.

4. EVALUATION

4.1 EVALUATION TEAM

The Evaluation Team will include SMEs from Oak Ridge National Laboratory and observers from other DOE computing labs to advise and consult. The Oak Ridge Leadership Computing Facility (OLCF), as the entity awarding subcontracts because of this RFP, will act as the source selection official.

4.2 EVALUATION FACTORS AND BASIS FOR SELECTION

Evaluation factors include mandatory requirements and perceived value of proposed research, addressed below in **performance features**, for next generation extreme scale HPC system. The Evaluation Team has identified the mandatory and target requirements for proposal responses. Offerors may identify and discuss additional research themes they believe may be of value. If the Evaluation Team agrees, consideration may be given to them in the evaluation process.

The Evaluation Team will determine the best overall research proposal by comparing differences in performance features and differences in price, striking the most advantageous balance between expected performance and the overall price. Offerors must, therefore, be persuasive in describing the value of their proposed performance features in enhancing the likelihood of successful performance or otherwise best achieving the DOE's objectives for extreme scale computing.

THE COMPANY reserves its rights to: 1) make selections on the basis of initial proposals; 2) negotiate with any or all Offerors for any reason; and 3) award subcontract(s) based on a single proposal that addresses more than one area of technology.

4.3 PERFORMANCE FEATURES

The Evaluation Team will validate that an Offeror's proposal satisfies the MR's. The Evaluation Team will assess how well an Offeror's proposal addresses the TR's. An Offeror is not solely limited to discussion of these themes. An Offeror may propose other areas of investigation if the Offeror believes they may be of value to future extreme scale HPC systems. If the Evaluation Team agrees, consideration may be given to them in the evaluation process. In all cases, the Evaluation Team will assess the merit of each proposal as submitted.

The Evaluation Team will evaluate the following performance features as proposed:

- The degree of innovation in the proposed R&D activities
- The extent to which the proposed R&D achieves substantial gains over existing industry roadmaps and trends
- The extent to which the proposed R&D will impact HPC and the broader marketplace
- Credibility that the proposed R&D will achieve stated results
- Credibility of the productization plan for the proposed technology

References

- [1] Benjamin L. Brown, William L. Miller, Deborah Bard, Amber Boehnlein, Kjersten Fagnan, Chin Guok, Eric Lançon, Sreeranjani Ramprakash, Mallikarjun Shankar, and Nicholas Schwarz. Integrated research infrastructure architecture blueprint activity (final report 2023). 7 2023.
- [2] Joseph R. Biden Jr. Executive order on america’s supply chains. <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/02/24/executive-order-on-americas-supply-chains/>.
- [3] Joseph R. Biden Jr. Executive order on the implementation of the chips act of 2022. <https://www.whitehouse.gov/briefing-room/presidential-actions/2022/08/25/executive-order-on-the-implementation-of-the-chips-act-of-2022/>.