

FastForward 2 R&D Draft Statement of Work

April 3, 2014



U.S. DEPARTMENT OF
ENERGY

Office of
Science

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Contents

1	INTRODUCTION	6
2	ORGANIZATIONAL OVERVIEW	7
2.1	The Department of Energy Office of Science	7
2.1.1	Advanced Scientific Computing Research Program	7
2.2	National Nuclear Security Administration	8
2.2.1	Advanced Simulation and Computing Program	8
3	MISSION DRIVERS	8
3.1	Office of Science Drivers	8
3.2	National Nuclear Security Administration Drivers	9
4	EXTREME-SCALE TECHNOLOGY CHALLENGES	9
4.1	Power Consumption and Energy Efficiency	9
4.2	Concurrency	10
4.3	Fault Tolerance and Resiliency	11
4.4	Memory Technology	11
4.5	Programmability	12
5	APPLICATION CHARACTERISTICS	12
6	ROLE OF CO-DESIGN	14
6.1	Overview	14
6.2	ASCR Co-Design Centers	14
6.3	ASC Co-Design Project	15
6.4	Proxy Apps	15
7	REQUIREMENTS	16
7.1	Description of Requirement Categories	16
7.2	Requirements for Research and Development Investment Areas	16
7.3	Common Mandatory Requirements	16
7.3.1	Solution Description (MR)	17
7.3.2	Research and Development Plan (MR)	17
7.3.3	Technology Demonstration (MR)	17
7.3.4	Productization Strategy (MR)	17
7.3.5	Staffing/Partnering Plan (MR)	18

7.3.6	Project Management Methodology (MR)	18
7.3.7	Intellectual Property Plan (MR)	18
7.3.8	Coordination with Current Research (MR)	18
8	EVALUATION CRITERIA	18
8.1	Evaluation Team	18
8.2	Evaluation Factors and Basis for Selection	18
8.3	Performance Features	19
8.4	Feasibility of Successful Performance	19
8.5	Supplier Attributes	20
8.5.1	Capability	20
8.6	Price of Proposed Research and Development	20
ATTACHMENT 1: NODE ARCHITECTURE RESEARCH AND DEVELOPMENT REQUIREMENTS		21
A1-1	Key Challenges for Node Architecture Technologies	21
A1-1.1	Component Integration	21
A1-1.2	Energy Utilization	21
A1-1.3	Resilience and Reliability	21
A1-1.4	On-Chip and Off-Chip Data Movement	21
A1-1.5	Concurrency	22
A1-1.6	Programmability and Usability	22
A1-2	Areas of Interest	22
A1-2.1	Component Integration	22
A1-2.2	Energy Utilization	22
A1-2.3	Resilience and Reliability	23
A1-2.4	On-Chip and Off-Chip Data Movement	23
A1-2.5	Concurrency	23
A1-2.6	Programmability and Usability	24
A1-3	Performance Metrics (MR)	24
A1-4	Mandatory Requirements	25
A1-4.1	Overall Node Design (MR)	25
A1-5.1	Component Integration (TR-1)	26
A1-5.2	NIC integration (TR-2)	26
A1-5.3	Energy Utilization (TR-1)	26
A1-5.4	Resilience and Reliability (TR-1)	26

A1-5.4	On-Chip Data Movement (TR-2)	26
A1-5.5	Processing Near Memory (TR-2).....	26
A1-5.6	Programmability and Usability: Hardware (TR-1).....	27
A1-5.7	Programmability and Usability: Software (TR-1)	27
A1-5.8	System Integration Strategy (TR-1)	27

ATTACHMENT 2: MEMORY TECHNOLOGY RESEARCH AND DEVELOPMENT REQUIREMENTS . 28

A2-1	Key Challenges for Memory Technology	28
A2-1.1	Energy Consumption	28
A2-1.2	Memory Bandwidth and Latency	28
A2-1.3	Memory Capacity.....	28
A2-1.4	Reliability	29
A2-1.5	Error Detection, Correction, and Reporting	29
A2-1.6	Processing in Memory	29
A2-1.7	Integration of NVRAM Technology.....	30
A2-1.8	Ease of Programmability.....	30
A2-1.9	New Abstractions for Memory Interfaces and Operations	30
A2-1.10	Integration of Optical and Electrical Technologies	31
A2-2	Areas of Interest.....	31
A2-3	Performance Metrics (MR)	31
A2-3.1	DRAM Performance Metrics.....	32
A2-4	Multivendor Integration Strategy (MR).....	34
A2-5	Target Requirements	34
A2-5.1	Energy per Bit	34
A2-5.2	Aggregate Delivered DRAM Bandwidth	34
A2-5.3	Memory Capacity per Socket.....	35
A2-5.4	FIT Rate per Node	35
A2-5.5	Error Detection Coverage and Reporting	35
A2-5.6	Advanced Processing in Memory Capabilities.....	35
A2-5.7	NVRAM Performance Metrics	35
A2-5.8	Multivendor Integration Strategy.....	35

1 INTRODUCTION

The Department of Energy (DOE) has a long history of deploying leading-edge computing capability for science and national security. Going forward, DOE's compelling science, energy assurance, and national security needs will require a thousand-fold increase in usable computing power, delivered as quickly and energy-efficiently as possible. Those needs, and the ability of high performance computing (HPC) to address other critical problems of national interest, are described in reports from the ten DOE Scientific Grand Challenges Workshops¹ that were convened in 2008–2010. A common finding across these efforts is that scientific simulation and data analysis requirements are exceeding petascale capabilities and rapidly approaching the need for exascale computing. However, workshop participants also found that due to projected technology constraints, current approaches to HPC software and hardware design will not be sufficient to produce the required exascale capabilities.

In April 2011 a Memorandum of Understanding was signed between the DOE Office of Science (SC) and the DOE National Nuclear Security Administration (NNSA), Office of Defense Programs, regarding the coordination of exascale computing activities across the two organizations. This led to the formation of a consortium that includes representation from seven DOE laboratories: Argonne National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory, and Sandia National Laboratories.

Funding for the DOE Exascale Computing Initiative has not yet been secured, but DOE has compelling real-world challenges that will not be met by existing vendor roadmaps. In response to these challenges, DOE SC and NNSA initiated an R&D program called FastForward that established partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme-scale computing. FastForward funded five companies (two of which have merged into one) starting in July 2012. With the initial two-year FastForward program coming to an end, DOE SC and NNSA are planning a follow-up program called FastForward 2. This new program will focus on two areas: Node Architecture and Memory Technology. The timeframe for the productization of the resulting Node Architecture and Memory Technology projects in 2020-2023. Node Architecture proposals for near-term product development that does not meet exascale needs are not in scope.

The Node Architecture focus area broadens the previous FastForward focus on Processors to include the entire architecture of a compute node. Both the node hardware and any necessary enabling software are in scope. A Node Architecture research proposal can also include several focus areas. For example, if novel runtime techniques or programming models are needed to make a new node architecture usable, research into these technologies could be included in a proposal. (However a software-only proposal would not be in scope.)

The Memory Technology focus area includes technologies that could be used in multiple vendors' systems. Memory technologies that are an integral part of a proprietary node design should be proposed in the Node Architecture focus area. Processor-in-memory (PIM) research

¹ <http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/>

may be proposed in the Memory Technologies focus area if the resulting technologies could be used in multiple vendors' node designs.

Vendors currently funded under FastForward may propose follow-on research under FastForward 2, and DOE also welcomes new research areas and new vendors for this program.

FastForward 2 seeks to fund innovative new or accelerated R&D of technologies targeted for productization in 5–8 years. The period of performance for any subcontract resulting from this request for proposal (RFP) will be approximately 27 months and end on November 1, 2016.

The consortium is soliciting innovative R&D proposals in Node Architecture and advanced Memory Technology that will maximize energy and computational efficiency while increasing the performance, productivity, and reliability of key DOE extreme-scale applications. The proposed technology roadmaps could have disruptive and costly impacts on the development of DOE applications and the productivity of DOE scientists. Therefore, proposals submitted in response to this solicitation should address the impact of the proposed R&D on both DOE extreme-scale mission applications as well as the broader HPC community. Offerors are expected to leverage the DOE SC and NNSA Co-Design Centers to ensure solutions are aligned with DOE needs. While DOE's extreme-scale computer requirements are a driving factor, these projects should also exhibit the potential for technology adoption by broader segments of the market outside of DOE supercomputer installations. This public-private partnership between industry and the DOE will aid the development of technology that reduces economic and manufacturing barriers to building systems that deliver exascale performance, and the partnership will also further DOE's goal that the selected technologies should 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 also the broader IT industry. This breadth will result in an increase in the consortium's ability to leverage commercial developments. The consortium does not intend to fund the engineering of near-term capabilities that are already on existing product roadmaps.

2 ORGANIZATIONAL OVERVIEW

2.1 The Department of Energy Office of Science

The 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 at 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 anywhere in the world.

2.1.1 Advanced Scientific Computing Research Program

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 significant modifications to today's tools and techniques to deliver on the promise of exascale science.

2.2 National Nuclear Security Administration

The NNSA is responsible for the management and security of the nation's nuclear weapons, nuclear non-proliferation, and naval reactor programs. It also responds to nuclear and radiological emergencies in the United States and abroad.

2.2.1 Advanced Simulation and Computing Program

Established in 1995, the Advanced Simulation and Computing (ASC) Program supports NNSA Stockpile Stewardship Programs' shift in emphasis from test-based confidence to simulation-based confidence. Under ASC, simulation and computing capabilities are developed to analyze and predict the performance, safety, and reliability of nuclear weapons and to certify their functionality. Modern simulations on powerful computing systems are key to supporting the U.S. national security mission. As the nuclear stockpile moves further from the nuclear test base through either the natural aging of today's stockpile or introduction of component modifications, the realism and accuracy of ASC simulations must further increase through development of improved physics models and methods requiring ever greater computational resources.

3 MISSION DRIVERS

3.1 Office of Science Drivers

DOE's strategic plan calls for promoting America's energy security through reliable, clean, and affordable energy, ensuring America's nuclear security, strengthening U.S. scientific discovery, economic competitiveness, and improving quality of life through innovations in science and technology. In support of these themes is DOE's goal to advance simulation-based scientific discovery significantly. This goal includes the objective to "provide computing resources at the petascale and beyond, network infrastructure, and tools to enable computational science and scientific collaboration." All other research programs within the SC depend on the ASCR to provide the advanced facilities needed as the tools for computational scientists to conduct their studies.

Between 2008 and 2010, program offices within the DOE held a series of ten workshops to identify critical scientific and national security grand challenges and to explore the impact exascale modeling and simulation computing will have on these challenges. The extreme scale workshops documented the need for integrated mission and science applications, systems software and tools, and computing platforms that can solve billions, if not trillions, of equations simultaneously. The platforms and applications must access and process huge amounts of data efficiently and run ensembles of simulations to help assess uncertainties in the results. New simulation capabilities, such as cloud-resolving earth system models and multi-scale materials models, can be effectively developed for and deployed on exascale systems. The petascale machines of today can perform some of these tasks in isolation or in scaled-down combinations (for example, ensembles of smaller simulations). However, the computing goals of many

scientific and engineering domains of national importance cannot be achieved without exascale (or greater) computing capability.

3.2 National Nuclear Security Administration Drivers

Maintaining the reliability, safety, and security of the nation’s nuclear deterrent without nuclear testing relies upon the use of complex computational simulations to assess the stockpile, to investigate basic weapons physics questions that cannot be investigated experimentally, and to provide the kind of information that was once gained from underground experiments. As weapon systems age and are refurbished, the state of systems in the enduring stockpile drifts from the state of weapons that were historically tested. In short, simulation is now used in lieu of testing as the integrating element. The historical reliance upon simulations of specific weapons systems tuned by calibration to historical tests will not be adequate to support the range of options and challenges anticipated by the mid-2020s, by which time the stewardship of the stockpile will need to rely on a science-based predictive capability.

To maintain the deterrent, the U.S. Nuclear Posture Review (NPR) insists that “the full range of Life Extension Program (LEP) approaches will be considered: refurbishment of existing warheads, reuse of nuclear components from different warheads, and replacement of nuclear components.” In addition, as the number of weapons in the stockpile is reduced, the reliability of the remaining weapons becomes more important. By the mid-2020s, the stewardship of the stockpile will need to rely on a science-based predictive capability to support the range of options with sufficient certainty as called for in the NPR. In particular, existing computational facilities and applications will be inadequate to meet the demands for the required technology maturation for weapons surety and life extension by the middle of the next decade. Evaluation of anticipated surety options is raising questions for which there are shortcomings in our existing scientific basis. Correcting those shortcomings will require simulation of more detailed physics to model material behavior at a more atomistic scale and to represent the state of the system. This requirement pushes the need for computational capability into the exascale level.

4 EXTREME-SCALE TECHNOLOGY CHALLENGES

The HPC community has done extensive analysis² of the challenges of delivering exascale-class computing. These challenges also apply more generally to extreme-scale HPC, regardless of whether or not the end result is an exaflop computer. In this section, we provide an overview of the most significant of these challenges.

4.1 Power Consumption and Energy Efficiency

All of the technical reports on exascale systems identify the power consumption of the computers as the single largest challenge going forward. Today, power costs for the largest petaflop systems

² http://science.energy.gov/~media/ascr/ascac/pdf/reports/Exascale_subcommittee_report.pdf;
http://science.energy.gov/~media/ascr/pdf/program-documents/docs/Arch_tech_grand_challenges_report.pdf;
http://science.energy.gov/~media/ascr/pdf/program-documents/docs/Crosscutting_grand_challenges.pdf;
<http://www.cse.nd.edu/Reports/2008/TR-2008-13.pdf>; <http://www.exascale.org/mediawiki/images/2/20/IESP-roadmap.pdf>

are in the range of \$5-10 million annually. To achieve an exascale system using current technology, the annual power cost to operate the system would be around \$250 million per year with a power load of 350 megawatts. To keep the operating costs of such a system in some kind of feasible range, a target of 20 megawatts has been established.

The power consumed by data movement will dominate the power budget of future systems. The power consumed in moving data between memory and processor is of particular concern. Historically a bandwidth/flop ratio of around 1 byte/flop has been considered a reasonable balance. For a current computer operating at 2 petaflop/s, the power required to maintain a 1 byte/flop ratio is about 1.25 MW. Extrapolating the JEDEC roadmap to 2020 and accounting for the expected improvements of DDR-5 technology, the total power consumption of the memory system would jump to 260 MW, well above the posited parameters for an exascale system. Even reducing the byte/flop ratio to 0.2—considered by some experts to be the minimum acceptable value for large-scale modeling and simulation problems—power consumption of the memory subsystem still would exceed 50 MW.

Achieving the power target for exascale systems is a significant research challenge. Even optimistic projections based on current R&D call for power consumption to be three to five times higher than we can tolerate for exascale. To improve power efficiency to the required level, we must explore a number of technical areas in hardware design. These may include: energy efficient hardware building blocks (central processing unit (CPU), memory, interconnect), novel cooling, and packaging, Si-Photonic communication, and power-aware runtime software and algorithms.

4.2 Concurrency

The end of increasing single compute node performance by increasing Instruction Level Parallelism (ILP) and/or higher clock rates has left explicit parallelism as the only mechanism in silicon to increase performance of a system. Scaling up in absolute performance will require scaling up the number of functional units accordingly, projected to be in the billions for exascale systems.

Efficiently exploiting this level of concurrency, particularly in terms of applications programs, is a challenge for which there currently are no good solutions. Memory latency further compounds the concurrency issue. We are already at or beyond our ability to find enough activities to keep hardware busy in classical architectures while long-time events such as memory references occur. While the flattening of clock rates has one positive effect in that such latencies will not increase dramatically by themselves, the explosive growth in concurrency will substantially increase the occurrence of high latency events; and the routing, buffering, and management of these events will introduce even more delay. When applications require synchronization or other interactions between different threads, this latency will exponentially increase the facilities needed to manage independent activities, which in turn forces up the level of concurrent operations that must be derived from an application to hide them.

A further complication arises from the explosive growth in the ratio of energy to transport data versus the energy to compute with it. At the exascale level, this transport energy becomes a front-and-center issue in terms of architecture. Reducing the transport energy will require creative packaging, interconnect, and architecture changes to bring the data needed by a computation energy-wise “closer to” the function units. This closeness translates directly into

reducing the latency of such accesses in creative ways that are significantly better than today's multi-level cache hierarchies.

4.3 Fault Tolerance and Resiliency

Resilience is a measure of the ability of a computing system and its applications to continue working in the presence of system degradations and failures. The resiliency of a computing system depends strongly on the number of components that it contains and the reliability of the individual components. Exascale systems will be composed of huge numbers of components constructed from VLSI devices that will not be as reliable as those in use today. It is projected that the mean time to interrupt (MTTI) for some components of an exascale system will be in the minutes or seconds range. Increasing evidence points to a rise in silent errors (faults that never get detected or get detected long after they generated erroneous results), causing havoc, which will only get more problematic as the number of components rises.

Exascale systems will continually experience failures, necessitating significant advances in the methods and tools for dealing with them. Achieving acceptable levels of resiliency in exascale systems will require improvement in hardware and software reliability, better understanding of the root cause of errors, better reliability, availability, and serviceability (RAS) collection and analysis, fault resilient algorithms and applications to assist the application developer, and local recovery and migration. The goal of research in this area is to improve the application MTTI by greater than 100 times, so that applications can run for many hours. Additional goals are to improve by a factor of 10 times the hardware reliability and improve by a factor of 10 times the local recovery from errors.

4.4 Memory Technology

Innovation in memory architecture is needed to address power, capacity, bandwidth, and latency challenges facing extreme-scale systems. The power consumption of current technology memory systems is predicted to be unsustainable for exascale deployment. Without new approaches, meeting power goals will require a drastic reduction in bytes per CPU core, because memory uses a large proportion of system power. Additionally, trends show a decrease in both memory capacity and bandwidth relative to system scale. The rate of memory density improvement has gone from a 4-times improvement every three years to a 2-times improvement every three years (a 30-percent annual rate of improvement). Consequently, the cost of memory technology is not improving as rapidly as the cost of floating-point capability. Thus, without new approaches, the memory capacity of an exascale machine will be severely constrained; it is anticipated that systems in the 2020 timeframe will suffer a 10-times loss in memory capacity relative to compute power. Research in advanced memory technologies, including high-capacity, low-power stacked memory or hybrid DRAM/NVRAM configurations could supply the capacity required while simultaneously balancing the power requirements.

Likewise, reduced memory bandwidth and increased latency will compound the memory capacity challenge. Neither bandwidth nor latency has improved at rates comparable to Moore's Law for processing units. On current petascale systems, memory access at all levels is the limiting factor in most applications, so the situation for extreme-scale systems will be critical. Innovative approaches are needed to provide critical improvements in latency and bandwidth, but techniques for improving the efficiency of data movement can also help. Options include better data analysis to anticipate needed data before it is requested (thus hiding latency),

determining when data can be efficiently recomputed instead of stored (reducing demands for bandwidth), closer integration with on- and off-chip networks, and improved data layouts (to maximize the use of data when it is moved between levels.)

4.5 Programmability

Programmability is the crosscutting property that reflects the ease by which application programs may be constructed. Programmability affects developer productivity and ultimately leads to the productivity of an HPC system as a tool to enable scientific research and discovery.

Programmability itself involves three stages of application development: (1) program algorithm capture and representation, (2) program correctness debugging, and (3) program performance optimization. All levels of the system, including the programming environment, the system software, and the system hardware architecture, affect programmability. The challenges to achieving programmability are myriad, related both to the representation of the user application algorithm and to underlying resource usage.

- **Parallelism**—sufficient parallelism must be exposed to maintain exascale operation and hide latencies. It is anticipated that 10-billion-way operation concurrency will be required.
- **Distributed Resource Allocation and Locality Management**—to make such systems programmable, the tension must be balanced between spreading the work among enough execution resources for parallel execution and co-locating tasks and data to minimize latency.
- **Latency Hiding**—intrinsic methods for overlapping communication with computation must be incorporated to avoid blocking of tasks and low utilization of computing resources.
- **Hardware Idiosyncrasies**—properties peculiar to specific computing resources such as memory hierarchies, instruction sets, and accelerators must be managed in a way that circumvents their negative impact while exploiting their potential opportunities without demanding explicit user control.
- **Portability**—application programs must be portable across machine types, machine scales, and machine generations. Performance sensitivity to small code perturbations should be minimized.
- **Synchronization Bottlenecks**—barriers and other over-constraining control methods must be replaced by lightweight synchronization overlapping phases of computation.

Novel execution models and architectures may increase programmability, thereby enhancing the productivity of DOE scientists.

5 APPLICATION CHARACTERISTICS

Multi-physics simulation is encountered in all missions supported by the DOE. "Multi-physics" numerical simulation is not simply simulation of complex phenomena on complex geometries. In its most simple form, multi-physics modeling involves two or more physical processes or

phenomena that are coupled and that often require disparate methods of solution. For example, turbulent fluid simulations must be coupled to structural dynamics simulations, shock hydrodynamics simulations must be coupled to solid dynamics or radiation transport simulations, and atomic-level defects in electronic devices must be coupled to large-scale circuit simulations.

Computational modeling with multiple physics packages working together faces many challenging issues at the extreme scale. Among these are problems in which coupled physical processes have inherently different spatial and/or temporal attributes, leading to possibly conflicting discretizations of space and/or time, as well as problems where the solution spaces for the coupled physical processes are inherently distinct with some packages working in a real space while other parts of the solution require a higher dimensional solution space. As an example, for coupled radiation-hydrodynamics, the physical processes in the simulation impose inherently distinct demands on the computer architecture. Hydrodynamics is characterized by moderate floating-point computations with regular, structured communication. Monte Carlo particle transport is characterized by intense fixed-point computations with random communication. As a result, multi-physics simulations typically require well-balanced computer architectures in terms of processor speed, memory size, memory bandwidth, and interconnect bandwidth, at a minimum.

Typical simulations are composed of multiple physics packages, which advance a shared set of data throughout the problem simulation time. While the details vary among packages, all implementations require that multiple physics packages run concurrently. The algorithms developed to model these physics processes have disparate characteristics when implemented on parallel computer architectures. The data for the simulation is distributed across a mesh representing the phenomena modeled. For each element of this mesh, the algorithmic demands have been characterized in terms of memory requirements, communication patterns, and computational intensity described in the table below. These packages often have competing computation and communication requirements. Generally, the strategy is to compromise among the various competing needs of these packages, but an overall driving principle for major applications is to attain the maximum degree of accuracy in the minimum amount of time.

One key challenge of the algorithms used in multi-physics applications is a balance of the memory access characteristics where both the patterns and the size requirements differ considerably and may fluctuate dramatically during the course of a calculation. Such variations impact both the communication patterns and the scaling characteristics of the codes. This is summarized in the following table:

Package	Memory per Mesh Element (KB)	Communication and Memory Access Patterns
A	0.2	Predictable with a modest amount of spatial and temporal locality
B	50–80	Predictable, but difficult to optimize, low spatial but high temporal locality
C	0.5–100	Unpredictable memory access, low spatial and low temporal locality
D	0.5	Predictable, with medium to high spatial and temporal locality

Multi-physics codes must also be able to run on capacity-class computer architectures as well as exascale computers. Portability and high-level abstractions in the programming model will be critical. The complexity of the physics interaction in multi-physics codes tends to demand that the implementation have a single, shared code based on all computer architectures (that is, rewriting for boutique vendor hardware can quickly become a maintenance challenge). To date, mechanisms for expressing data hierarchies and optimization accessible by a given hardware realization have been closer to machine-level programming than high-level abstractions. As architectural complexities increase, research into appropriate abstractions in the programming model is needed. Additionally, improvements in the computational environment, such as compilers and tools, are needed. This need will become increasingly critical on exascale computer architectures. Addressing the issues of restrictions due to power constraints and heterogeneous node architectures are additional challenges.

6 ROLE OF CO-DESIGN

6.1 Overview

The R&D funded through this RFP is expected to be the product of a co-design process. Co-design refers to a system-level design process where scientific problem requirements influence architecture design and technology and architectural characteristics inform the formulation and design of algorithms and software. To ensure that future architectures are well-suited for DOE target applications and that DOE scientific problems can take advantage of the emerging computer architectures, major R&D centers of computational science are formally engaged in the hardware, software, numerical methods, algorithms, and applications co-design process.

Co-design methodology requires the combined expertise of vendors, hardware architects, system software developers, domain scientists, computer scientists, and applied mathematicians working together to make informed decisions about the design of hardware, software, and underlying algorithms. The future is rich with trade-offs, and give and take will be needed from both the hardware and software developers. Understanding and influencing these trade-offs is a principal co-design requirement.

ASCR and ASC have established multiple application co-design centers that serve as R&D collaboration vehicles with all aspects of the extreme-scale development ecosystem, especially vendors.

6.2 ASCR Co-Design Centers

In 2011, ASCR made awards to three application co-design centers. Each center focuses on a specific application that is an important driver for exascale³ and is a distributed collaboration between multiple national laboratories and university partners. Development of that application facilitates exploration of issues of mathematics, algorithms, computer science, systems software, and of course, hardware in the co-design process. For a detailed description of the ASCR co-design centers see <http://science.energy.gov/ascr/research/scidac/co-design/>.

³ <http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/>

6.3 ASC Co-Design Project

The NNSA labs and ASC program have defined a coordinated co-design strategy that leverages the work of the ASCR co-design centers while focusing on the unique needs of the ASC program. ASC is a mission-driven program with applications currently in use that are of importance to run at exascale in support of stockpile stewardship, namely the Engineering and Physics Integrated Codes (EPICs). To meet the key needs of the EPICs, ASC has established the National Security Applications (NSApp) Co-Design project. NSApp focuses on these established applications as the drivers and participates in co-design largely through proxy applications. Additional information is available at <https://asc.llnl.gov/codesign/>.

6.4 Proxy Apps

The DOE co-design centers make extensive use of proxy applications to represent the application workflow and requirements to the exascale ecosystem. These applications codes are used to understand the effects of hardware tradeoffs, and to explore and develop new technologies, runtime systems, languages, programming models, algorithms, tools, file systems, and visualization techniques. Whenever possible, proxy apps are openly available—with occasional need to protect the original source under export-control rules or proprietary access rules in some cases where vendor modifications are supplied back to the co-design center.

In general, a small application code that represents some aspect of the computational workflow of a full application is a proxy app. Proxy apps can be grouped into three categories in increasing sophistication and fidelity to the parent applications and integrated codes:

- **Kernels:** these are small code fragments (algorithms) that are used extensively by the parent application and are deemed essential to perform optimally,
- **Skeleton apps:** these apps reproduce the data flow of a simplified application and make little or no attempt to investigate numerical performance. They are primarily useful in investigating memory management, network performance characteristics, I/O, ...,
- **Mini- or compact apps:** these apps contain the dominant numerical kernels contained in the parent application and represent the computational workflow in as compact a form as possible.

It is important to emphasize that these proxy apps will not be static but will evolve significantly during the co-design process. The co-design centers anticipate the requirement for domain application code-developers to spend significant time with the vendors as well as vendor developers and architects to spend significant time with the co-design centers.

ASC and ASCR co-design centers are developing and publishing their proxy apps. Some that are available today are:

ExMatEx	http://www.exmatex.org
ExaCT	http://www.exactcodesign.org
CESAR	http://cesar.mcs.anl.gov
TORCH	https://ftg.lbl.gov/projects/torch
Mantevo	http://software.sandia.gov/mantevo

NERSC SSP	http://www.nersc.gov/research-and-development/performance-and-monitoring-tools/sustained-system-performance-ssp-benchmark
LULESH	https://computation.llnl.gov/casc/ShockHydro
SNAP	https://github.com/losalamos/snap

These proxy apps and new ones will be reconfigured to express the workflow of the extreme-scale applications on extreme-scale nodes, including the memory sub-system. Both layout and movement for both data and task are of interest. With emerging constraints on memory per compute unit and on the energy of moving data, current algorithms and strategies for data layout and data movement may prove ineffective. The lessons learned from testing the current algorithms and strategies on proposed extreme-scale node architectures will be used by the Co-design Centers to re-express the proxy apps to assess extreme-scale node performance.

7 REQUIREMENTS

7.1 Description of Requirement Categories

Requirements are either mandatory (designated MR) or target (designated TR-1 or TR-2), and are defined as follows:

- MRs are performance features essential to DOE requirements. An Offeror must satisfactorily address *all* MRs 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 DOE but will not result in a nonresponsive determination if omitted from a proposal. TRs add value to a proposal and are prioritized by dash number, as described below.

TR-1s and MRs are of equal value. The aggregate of MRs and TR-1s form a baseline solution. TR-2s are goals that combine with MRs and TR-1s to produce a more useful solution. Therefore, the ideal proposal would meet or exceed all MRs, TR-1s, and TR-2s.

7.2 Requirements for Research and Development Investment Areas

Detailed requirements for each of the targeted R&D areas of investment are provided as Attachments to this document. Offerors wishing to pursue both Node Architecture and Memory Technologies research efforts shall submit a separate proposal for each topic. Offerors wishing to conduct research on Node Architecture that includes more than one focus area shall submit a single proposal that specifies the dependencies between the focus areas. The proposal budget shall clearly specify the budget for each focus area.

7.3 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.

7.3.1 Solution Description (MR)

Offeror shall describe the proposed R&D and how it addresses the target requirements in Attachment 1 or Attachment 2, with emphasis on how it will increase the performance of key DOE extreme-scale applications as represented by the Proxy Applications described in Section 6.4.

Offerors shall discuss the innovative nature of the proposed R&D. Work that funds a company's current roadmap is not desired. Technology acceleration is acceptable if there is a clear DOE 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.

7.3.2 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. Milestone statements must provide sufficient specificity that DOE reviewers can legitimately determine that a report fails to document meeting the milestone requirements. Pricing shall be assigned to each milestone and deliverable. A schedule for periodic technical review by the DOE laboratories shall also be provided.

The R&D funded through this RFP is expected to be the product of a co-design process. More specifically, Offerors are expected to engage in co-design activities with DOE's ASC and ASCR Exascale Co-design Centers. The R&D plan shall include a discussion of how Offeror plans to collaborate with DOE researchers on co-design, with a detailed description of planned co-design efforts if known.

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.

7.3.3 Technology Demonstration (MR)

Offeror shall provide a plan for demonstrating the utility and feasibility of the technologies developed under this program. The most compelling demonstration would be a working prototype. Such demonstrations may be supplemented with a simulation or analysis that assesses the impact of a proposed development. If funding provided through this RFP is insufficient to produce a prototype, Offeror shall propose emulations, simulations, or other demonstrations that assesses the impact of the proposed development.

7.3.4 Productization Strategy (MR)

Offeror shall describe how the proposed technology will be commercialized, productized, or otherwise made available to customers. Offeror shall include identification of target customer base/market(s) for the technology. Offeror 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. Offeror shall indicate the projected timeline for productization.

7.3.5 Staffing/Partnering Plan (MR)

Offeror shall describe staffing categories and levels for the proposed R&D activities. Any collaboration with other industry partners and/or universities shall be identified, and the names of any key personnel from these partners/subcontractors shall be provided together with a description of their contributions to the overall effort.

7.3.6 Project Management Methodology (MR)

Project management and regular project status reporting are required. Offeror shall describe project management methodology and provide a communication plan that indicates methods of communication (for example, written report, teleconference, and/or face-to-face meeting) and frequency (for example, weekly, monthly, and/or quarterly). Offeror shall include quarterly reviews in the project plan. These reviews shall be scheduled semi-annually at a DOE facility and at a time to be determined by SC and NNSA. Offeror shall propose a format for the other half of the required quarterly reviews.

7.3.7 Intellectual Property Plan (MR)

Proposals shall include a plan for how each item of existing intellectual property (IP) and any new IP developed under this statement of work will be handled, including requested IP ownership and licensing. Please consult RFP letter for information on Federal regulations concerning IP.

7.3.8 Coordination with Current Research (MR)

Offerors who are currently funded under a FastForward or DesignForward subcontract shall describe how their proposed research continues or complements existing research under those programs.

8 EVALUATION CRITERIA

8.1 Evaluation Team

The Evaluation Team includes representation from seven DOE laboratories: Argonne National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory, and Sandia National Laboratories. Lawrence Livermore National Security (LLNS), as the entity awarding subcontracts as a result of this RFP, will act as the source selection official.

8.2 Evaluation Factors and Basis for Selection

Evaluation factors are mandatory requirements, performance features, supplier attributes, and price that the Evaluation Team will use to evaluate proposals. The Evaluation Team has identified the mandatory requirements, performance features, and supplier attributes listed above and in each Attachment that should be discussed in the proposal. Offerors may identify and discuss other performance features and supplier attributes that may be of value to the Evaluation Team. If the Evaluation Team agrees, consideration may be given to them in the evaluation

process. The Evaluation Team's assessment of each proposal's evaluation factors will form the basis for selection. LLNS intends to select the responsive and responsible Offerors whose proposals contain the combination of price, performance features, and supplier attributes offering the best overall value to DOE. The Evaluation Team will determine the best overall value by comparing differences in performance features and supplier attributes offered with 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 and supplier attributes in enhancing the likelihood of successful performance or otherwise best achieving the DOE's objectives for extreme scale computing.

LLNS desires to select two Offerors for each area of technology discussed in the Attachments to this SOW. However, LLNS reserves the right, based on the proposals received in response to the RFP, to select none, one, or more than two for any area of technology.

LLNS reserves its rights to: 1) make selections on the basis of initial proposals and 2) negotiate with any or all Offerors for any reason.

8.3 Performance Features

The Evaluation Team will validate that an Offeror's proposal satisfies the MRs. The Evaluation Team will assess how well an Offeror's proposal addresses the TRs. An Offeror is not solely limited to discussion of these features. An Offeror may propose other features or attributes if the Offeror believes that they are of value. If the Evaluation Team agrees, consideration may be given to them in the evaluation process. In all cases, the Evaluation Team will assess the value of each proposal as submitted.

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

- How well the proposed solution meets the overall programmatic objectives expressed in the SOW
- The degree to which the technical proposal meets or exceeds any TR
- 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
- Realism and completeness of the project work breakdown structure

8.4 Feasibility of Successful Performance

The Evaluation Team will assess the likelihood that the Offeror's proposed research and development efforts can be meaningfully conducted and completed within the anticipated three-year subcontract period of performance. The Evaluation Team will also assess the risks, to both the Offeror and the DOE laboratories, associated with the proposed solution. The Evaluation

Team will evaluate how well the proposed approach aligns with the Offeror's corporate roadmap and the level of corporate commitment to the project.

8.5 Supplier Attributes

The Evaluation Team will assess the following supplier attributes.

8.5.1 Capability

The Evaluation Team will assess the following capability-related factors:

- The Offeror's experience and past performance engaging in similar R&D activities
- The Offeror's demonstrated ability to meet schedule and delivery promises
- The alignment of the proposal with the Offeror's product strategy
- The expertise and skill level of key Offeror personnel (All lead and key personnel should be identified by name and brief CV's for these personnel should be provided.)
- The contribution of the management plan and key personnel to successful and timely completion of the work

8.6 Price of Proposed Research and Development

The Evaluation Team will assess the following price-related factors:

- Reasonableness of the total proposed price in a competitive environment
- Proposed price compared to the perceived value
- Price tradeoffs and options embodied in the Offeror's proposal
- Financial considerations, such as price versus value

ATTACHMENT 1: NODE ARCHITECTURE RESEARCH AND DEVELOPMENT REQUIREMENTS

The focus of this effort is to investigate node architectures for future exascale computer systems. This includes the set of hardware and software features that jointly enable productive use of a compute node within a future exascale system. A compute node is a terminal node in computer system interconnection network. The node hardware is composed of a collection of (possibly heterogeneous) processor and memory components. It has a network attachment, but the multi-node network fabric itself is not within the scope of this effort. The term *processor* typically refers to the set of capabilities within a single microprocessor chip, or a tightly integrated set of capabilities that span several chips (for example, chip stacks, chip carriers, chip sets, and other such approaches). Key challenges include energy usage, performance, data movement, concurrency, reliability, and programmability, all of which are interrelated. Of particular interest is the development of mechanisms for examining trade-offs between these interrelated aspects.

A1-1 Key Challenges for Node Architecture Technologies

A1-1.1 Component Integration

A tightly-coupled node architecture can improve design flexibility, operational efficiency, and robustness, and it can reduce costs. Further, the location of node-based components and the functionalities that they support impacts node energy usage. Methods to reduce energy consumption include coupling components more tightly, locating related functionality on the same component, and ensuring that the capabilities of different components match well.

A1-1.2 Energy Utilization

Energy and power are key design constraints for exascale machines. Techniques to minimize or constrain power used by computations while maintaining predictable behavior are needed. Possible areas include architectural features to improve application efficiency, advanced power gating techniques, near threshold operation, as well as packaging techniques such as 3D integration.

A1-1.3 Resilience and Reliability

Node reliability is a critical concern, especially since future DOE supercomputers will utilize hundreds of thousands of nodes. If FIT (Failures in Time) rates cannot be improved, the MTBI (Mean Time Between Interrupts) will fall to unacceptable levels. Machines with frequently failing components will require continual operator maintenance. Techniques that increase the MTBR (Mean Time Between Repairs) by decreasing how often or the urgency of operators servicing node failures are of interest. The ability to identify, contain, and overcome faults quickly with as little human intervention as possible is of paramount importance.

A1-1.4 On-Chip and Off-Chip Data Movement

Improved methods are needed for on-chip and off-chip data movement. The ability to move data efficiently limits the performance of many HPC applications. The energy required to move one

bit of data within the processor and to memory must be reduced to a few picoJoules. In addition, improved memory interfaces can increase the effective bandwidth delivered to applications. Also, having processing capabilities as close as possible to the storage of data may be desirable.

A1-1.5 Concurrency

Future increases in clock speeds are expected to be limited. As a consequence, processor companies are dramatically increasing concurrency (for example, more cores, greater instruction bundling, and multithreading) as feature sizes decrease. Managing this concurrency and the associated data movement is a considerable challenge. Many technologies could address the associated challenges in exploiting the available concurrency, including improved synchronization mechanisms, flexible atomic operations, and transactional memory. Architectural mechanisms to handle work queue models efficiently could also improve application performance.

A1-1.6 Programmability and Usability

Achieving high performance on next-generation processors will be a challenge. Application developers will need to deal with massive concurrency and may need to manage locality, power, and resilience. A software ecosystem is needed to support the development of new applications and the migration of existing codes. Novel architectures and execution models may increase programmability and enhance the productivity of DOE scientists. Issues include the programmability of proposed architectures both in terms of complexity and the effort that will be required on the part of DOE scientists to achieve high performance.

A1-2 Areas of Interest

The following topics are examples of concerns that a Node Architecture proposal could address. Some of these topics may apply only to certain architectures, and some may be mutually exclusive, so a proposal need not address all of them. Proposals may also address other topics relevant to the design of an exascale compute node. Proposals that address a coherent subset of topics in depth are preferable to those that address all of them superficially.

A1-2.1 Component Integration

- Development of mechanisms to understand the trade-offs between power, resilience and performance, both statically (when the node is designed), and dynamically (at runtime)
- Integration of standard building blocks into a balanced node architecture for HPC

A1-2.2 Energy Utilization

- Advances that improve the power efficiency of processors
- Advances in measurement and application control of power utilization
- Advances that support high-performance, power-efficient processor integration with memory, optics, and networking

- Techniques to reduce cooling energy requirements

A1-2.3 Resilience and Reliability

- Advances that improve the resiliency or reliability of nodes, for example, improved fault detection and correction
- Advances that permit automatic rollback (within a window) after a fault or synchronization error
- Advances that demonstrate hardware/software resilience tradeoffs to improve overall time to solution
- Advances that lower the impact or cost of partial component failures or yield issues without significantly increasing total cost of ownership, such as hot sparing with automatic failover, overprovisioning of resources and ability to operate in a degraded state

A1-2.4 On-Chip and Off-Chip Data Movement

- Advances that allow extremely low-latency response to incoming messages
- Advances to enable very efficient latency hiding techniques
- Improvements to the performance and energy efficiency of messaging, remote memory access, and collective operations
- Advances that allow explicit (software controlled) movement of data in and out of various on-chip memory regions (for example, levels of cache)
- Hardware support for large numbers of short messages to achieve low latency
- Integration of the network interface as a first-class component with the processor and memory system, enabling higher communication efficiencies.
- Other hardware mechanisms for eliminating overhead
- Integration of processing elements near to where the data is stored. (Processing in/near Memory.)

A1-2.5 Concurrency

- Advances that improve the scalability of processor designs as the number of processing units per chip increase
- Advances that address the inherent scaling and concurrency limits in applications
- Advances that improve the efficiency of process or thread creation and their management
- Advances that reduce the synchronization and activation time of large numbers of on-chip threads
- Advances to assist in identification of active performance constraints within the system, such as latency or throughput limited sections, memory and network bottlenecks

A1.2.6 Programmability and Usability

- Advances that significantly improve the performance and energy efficiency of arithmetic patterns common to DOE applications but are not well supported by today's processors, for example, short vector operations such as processing in vector registers
- Advances that allow efficient computation on irregular data structures (for example, compressed sparse matrices and graphs)
- Research to determine the most effective option(s) for cache and memory coherency policies; configurable coherency policies and configurable coherence or NUMA domains may be options; coherency policies might also be a power management tool.
- Research on efficient mapping of multiple levels of application parallelism to node architecture parallelism
- Advances in software and hardware that allow a user or runtime system to measure and to understand node activities and to adjust implementation choices dynamically
- Advances that enable a programmer to understand and reason about optimally programming the node, and exposing the right architectural details for consideration. Development of a target independent programming system.

A1-3 Performance Metrics (MR)

Offeror shall estimate or quantify the impact of the proposed technology over industry roadmaps and trends. This information shall be provided for each applicable metric listed below. If a proposed technology will not substantially improve a metric listed below, the proposal shall state that. If Offeror determines that alternative metrics would better represent the benefits of a proposed technology, then Offeror shall explain why they believe the alternative metrics are more meaningful and estimate the impact of the proposed technology based on those metrics.

Estimated metric values shall reflect solutions that are productized in 5 to 8 years. These metrics are independent, but a solution that can deliver advances in more than one metric is more desirable than one that solves only one metric at the expense of the others. The most meritorious improvements will make substantial gains over industry roadmaps/trends and substantiate a convincing path to achieving the extreme-scale technology characteristics required by DOE.

- Node and socket power requirements
- Processor computational capability per watt
- FIT rate per socket and node
- Error detection, correction, and coverage of hard and soft error types
- Improvements in application MTBI (Mean Time Between Interrupt)
- Continued functionality in the presence of partial node component failures, extending the MTBR (Mean Time Between Repairs.)
- Energy per bit for data transfers

- Computational capacity per node
- Effective bandwidth delivered to application from memory subsystem
- Efficient operation as measured by a weighted sum of time and energy to solution, chosen to approximate the likely balance of capital and operating expenses for a node

A1-4 Mandatory Requirements

The following are mandatory requirements for Node Architecture proposals.

A1-4.1 Overall Node Design (MR)

The Offeror shall provide a description of the hardware architecture of the overall node design that their proposed work is based upon. A software-only solution would not be acceptable. The description shall include:

- a high-level architectural diagram to include all major components and subsystems;
- descriptions of all the major architectural hardware components in the node to include: processing units, memory subsystem, network interface, and relevant software;
- a concise description of the areas of the node architecture that the proposed work is intended to address as well as how it will enable the determination of an optimal overall node architecture.

The proposed node architectural investigation shall address the key challenges specified in Section A1-1 of this appendix. The proposed effort shall include:

- an evaluation of the proposed initial execution model;
- the development of the conceptual node design;
- an analysis of the proposed design that shows the impact of the design on the key challenges; and
- initial metrics for evaluating the success of the design.

A1-5 Target Requirements

The requirements below apply to supercomputers that will be deployed within 5 to 8 years to meet DOE mission needs. As previously stated, Offerors need not address all problem areas, and thus the Offeror need not respond to a TR below if the proposed capability does not address that problem area. In all TR responses that are provided, Offeror should discuss what progress will be made in the next two years and describe what follow-on efforts will be needed to fully achieve these goals. Offeror should describe in detail how the metric will be evaluated, including the measurement method that will be used (for example, simulation or prototype) and any assumptions that will be made.

In the discussion below, a node is defined as the smallest physical unit of hardware that contains a processor chip(s), memory, and at least one network connection to connect to other such units.

A1-5.1 Component Integration (TR-1)

Mechanisms for producing a highly optimized node that has a tight coupling of components and are highly optimized for HPC are desired. Solutions should describe how they would achieve this goal. To keep system sizes manageable, the overall performance of a node should be greater than 10 teraflop per second.

A1-5.2 NIC integration (TR-2)

Solutions should discuss how new functionality enabled by tight integration will contribute to increased communication efficiency. The target interconnection performance for the node is:

- MPI Applications: 500 Million messages per second
- PGAS Applications: 2 Billion messages per second
- Back-to-back (no router in path) message latency of 500ns

A1-5.3 Energy Utilization (TR-1)

An energy and concurrency efficient node that achieves high performance on a broad range of DOE applications (for example, the co-design center applications described previously) is highly desired. Solutions should realize greater than 50 GF/Watt at system scale while maintaining or improving system reliability.

A1-5.4 Resilience and Reliability (TR-1)

Mean Time to Application Failure (TR-1). Processor designs should make advances that lead to a mean time to application failure requiring user or administrator action of six days or greater in an exascale system, as determined by estimates of system component FIT rates and application recovery rates.

Mean Time Between Repair (TR-1). The Mean Time Between Repair (MTBR) for a single node should be greater than 30 years. Repair is required whenever the node functionality drops below the expected minimum level, necessitating operator service or part replacements.

A1-5.4 On-Chip Data Movement (TR-2)

Nodes that meet the capacity and bandwidth demands of extreme scale applications are expected to contain multiple types of memory to meet the DOE's cost and power constraints. Solutions will address how best to balance these memory systems in terms of bandwidth and capacity within a node to optimize for application performance and programmer productivity at minimum cost. Offeror should describe in detail how this will be accomplished.

A1-5.5 Processing Near Memory (TR-2)

Solutions will describe 1) which levels of the memory hierarchy are appropriate targets for processing near (or in) memory; 2) the processing capabilities that they will explore; and 3) an estimate of the potential benefits of the proposed approach. While node architecture includes processing near memory, processing in memory components that independent of the integrated node technology are not in scope.

Solutions should describe the extent to which the proposed technology will augment the capabilities of the memory subsystem while preserving programmability.

A1-5.6 Programmability and Usability: Hardware (TR-1)

Solutions will describe novel features of the hardware that allow applications to use the proposed architecture more efficiently. For example, structures that de-emphasize the importance of which loops are threaded or SIMD vectorization rates are of particular interest, as are structures that enable adaptive runtimes. How the proposed solutions increase performance without increasing programmer effort should be highlighted.

A1-5.7 Programmability and Usability: Software (TR-1)

Solutions will need a software ecosystem that supports the development of new applications, the migration of existing applications, identification of performance bottlenecks, runtime performance introspection, application maintenance, and application portability, while enabling DOE scientists to achieve high performance with no more effort than is required for today's high-end computers. Offeror should describe in detail how the solution improves programmability and usability. In addition, Offeror should describe how the proposed solutions will support abstractions for a particular hardware feature and, thus, an independent programming system and the co-design of the next generation of applications.

A1-5.8 System Integration Strategy (TR-1)

Although this RFP does not address System Integration, research into Node Architectures should include planning for how an exascale system can be built from a node. Therefore, proposals for Node Architecture research should include milestones that call for the Offeror to make contact with one or more potential system integration teams (either within the Offeror's company or externally) and establish the feasibility of building an exascale system from a proposed node design.

ATTACHMENT 2: MEMORY TECHNOLOGY RESEARCH AND DEVELOPMENT REQUIREMENTS

While considerable progress has been made through industry and Fast Forward efforts to reduce power and increase capacity and bandwidth, more research is needed to meet DOE requirements for HPC systems while at the same time developing memory parts of value in the commercial sector. Effort must continue to focus on ways to reduce power consumption, in the DRAM itself, through memory organization and architecture, and through approaches that include new memory devices, such as NVRAMs. As noted earlier, research funded in this focus area must work toward memory technologies that can be used in multiple vendors' systems.

A2-1 Key Challenges for Memory Technology

The following items are some areas of emphasis in memory technology based on the requirements of DOE's application workload. None of these need to be construed as pointing to specific prescribed solutions.

A2-1.1 Energy Consumption

Power consumption is a leading design constraint for future systems. Chief among these concerns is the power consumed by memory technology, which would easily dominate the overall power consumption of future systems if we continue along current technology trends. The target for an exaflop system in the 2020+ timeframe is 20 megawatts for the complete system. If we extrapolate commodity DDR memory technology trends, the memory system alone would eclipse the target power budget and make future HPC systems of all scales less effective. FastForward 2 seeks to develop memory technologies to improve the energy efficiency of memory while improving capacity, bandwidth, and resilience.

A2-1.2 Memory Bandwidth and Latency

Memory bandwidth has always been a major bottleneck for the performance of HPC applications. As core count of processors has increased, the memory bandwidth available to each core has significantly decreased. Higher memory bandwidth enables a wider array of algorithms to utilize available computing performance fully. Approaches to reducing (perceived) latency to the end application, perhaps through adaptive prefetching are encouraged. FastForward 2 will emphasize the development and acceleration of technology to increase memory bandwidth and to reduce latency while keeping cost, reliability, and power consumption under control.

A2-1.3 Memory Capacity

The rate of improvement for DRAM density has slowed in recent year from quadrupling every three years to doubling every three years. In comparison, logic density and the cost of flops is improving at a much more rapid rate. The consequence is that memory capacity per peak computational performance will decrease compared to past machines. This trend impacts DOE applications significantly because increased problem resolution, which requires larger memory capacity, is at least as important for many scientific applications as improvements in computational performance. Further, technology roadmaps out to the 2020s forecast high-capacity memory with low bandwidth and low-capacity memory with high bandwidth—but not

both. However, the DOE mission need and scientific objectives require improvements both in increased problem sizes (limited by memory capacity) and in performance (limited by memory bandwidth) in the same memory space. A solution that delivers one or the other (but not both), will fail to meet mission objectives. FastForward 2 seeks to accelerate and develop new technology options that can deliver both capabilities (bandwidth and capacity) in the same cost-effective package.

A2-1.4 Reliability

Components that are otherwise reliable in consumer applications that contain only a handful of devices have high aggregate failure rates for scalable HPC systems that typically include millions of components. Even in today's HPC systems and large-scale data centers, memory DIMMS are among the most common sources of hardware failure. A large-scale field study by Google and the University of Toronto has shown that DRAM failure rates are much higher than originally anticipated⁴. For scalable systems, FastForward 2 seeks to develop and accelerate technologies that dramatically reduce DRAM component failure rates over a baseline that is largely set by smaller scale consumer devices.

A2-1.5 Error Detection, Correction, and Reporting

With respect to component failure rates (reliability), modern error detection and correction technology may not be sufficient for the increased rate of transient errors. For scalable HPC systems and large-scale data centers, there is an increased observation of uncorrectable (double-bit or burst) errors. Even more worrisome, silent errors are already apparent in modern HPC systems and an increased incidence of them has been observed. More comprehensive error detection and reporting technology (for example, S.M.A.R.T. technology for system boards) would greatly improve the usability of these resilience features. Furthermore, new techniques for error detection and correction are possible: multidimensional error coding, multimode memory hierarchies with configurable error correction and detection, and integration with system software and programming features. FastForward 2 is seeking technologies to improve and even to scale our ability to detect and to correct transient errors, and to reduce the possibility of silent errors in large-scale systems.

A2-1.6 Processing in Memory

An alternative approach to improving effective memory bandwidth is to embed computing operations within the memory component to reduce pressure on memory bandwidth. At a minimum, this approach includes embedding basic element/word-granularity operations such as atomic memory operations and synchronization primitives in the memory to eliminate round trips of data movement between the processor and the memory. At a medium level of integration, one could embed vector-primitives such as strided gather operations, general gather/scatter, indirect address chaining, smart prefetchers, and checksum operations (for end-to-end error detection) in the memory system to reorganize areas of memory to improve data transfer

⁴ B. Schroeder, E. Pinheiro, W-D. Weber, "DRAM Errors in the Wild: A Large-Scale Field Study," SIGMETRICS/Performance'09, ACM, Seattle WA. 2009.

performance. General-purpose processing-in-memory is the most extreme and general approach to embedding a processing capability into the memory subsystem. In addition to direct connection to the attached node, the ability of a memory package to inject and receive data directly into/from the network is of interest. FastForward 2 seeks novel ideas for embedding processing in memory to improve data transfer efficiency or even to eliminate the need to move data off the memory chip. Solutions shall present a standard interface and be usable by any CPU. Further, proposals may suggest approaches to transfer data between memory subsystems through the interconnection network without node intervention.

A2-1.7 Integration of NVRAM Technology

Solid-state storage technology (FLASH and other forms of NVRAM) has found a way into consumer and HPC systems primarily as disk/file system technology. However, we see many opportunities for improved performance and capability if NVRAM were integrated directly into the memory hierarchy rather than as a disk replacement. For example, node-local NVRAM that is substantially more trustworthy, secure, and reliable than the DRAM memory that holds active application state would offer substantial benefits for checkpointing/resilience technology. Furthermore, NVRAM will need to improve durability in order to be included in an extreme-scale system. On-chip NVRAM could preserve local register or pipeline state to support micro-checkpointing for resilience or instant power-down operation for chips (which are useful in the consumer space, too). NVRAM can be used to hold tables and data items that are seldom written to relieve some pressure from the DRAM portion of the memory system. NVRAM-backed DRAM could enable power-off of areas of memory that are unused or under-utilized. FastForward 2 is seeking novel applications and solutions involving deeper integration of NVRAM technology in the memory hierarchy.

A2-1.8 Ease of Programmability

As novel technologies are added to computer systems, the application may need to manage increased memory hierarchy complexity. NUMA main memory is prevalent, and frequently requires programmer optimization. In addition, new technologies, such as high-speed scratchpad memory, heterogeneous cores, or software-managed caches, create disparate memory spaces with varying performance characteristics and capacities. Support of a broader ecosystem of software for the device would ensure that the features will continue to be supported across systems. Fast Forward 2 is seeking novel hardware and software solutions to simplify the management of deep memory hierarchies.

A2-1.9 New Abstractions for Memory Interfaces and Operations

Separation of high level memory operations (e.g., load or store) from low level aspects of managing vendor-specific devices (e.g., DRAM timing parameters, bank organization) could provide increased flexibility to support multiple memory types, and could allow combining of heterogeneous memory devices. It would provide a transparent upgrade path as denser or lower power devices become available. The separation may enable a simpler memory interface to the node, hiding the possible complexity of in-package memory hierarchy, BIST, reliability, and other memory functions. Fast Forward 2 is seeking memory controller architectures that present a high level abstract interface to the node, and manage memory-part-specific functions within the memory package.

A2-1.10 Integration of Optical and Electrical Technologies

Recent research on silicon photonics has produced promising results that may support the development of bandwidth and latency improvements in both on-die and off-die interconnects. Fast Forward 2 is seeking memory technologies that naturally integrate with photonics capabilities to provide higher bandwidth and reduced latency to memory-resident data.

A2-2 Areas of Interest

Below are some areas of technology development and acceleration that could be considered in memory R&D proposals to address DOE's extreme-scale computing needs. ***Proposals are not limited to these areas, and alternative topics in memory technology for exascale systems are encouraged.***

- Technologies to improve the energy efficiency of memory while improving capacity, bandwidth, and resilience.
- Technology to increase memory bandwidth while not substantially increasing cost, reliability, and power consumption.
- New technology options that can deliver both bandwidth and capacity in the same cost-effective package.
- Technology innovations to reduce latency to application memory requests.
- Technologies that dramatically reduce DRAM and NVRAM component failure rates over a baseline that is largely set by smaller scale consumer devices.
- Technologies to improve and to scale the ability to detect and to correct transient errors, and to prevent the incidence of silent errors in large-scale systems.
- Novel ideas for self-contained, CPU-agnostic embedding of processing in memory to improve data transfer efficiency to local or remote CPUs or even to eliminate the need to move data off the memory chip.
- Novel applications and solutions involving deeper integration of NVRAM technology in a multi-level memory hierarchy.
- Novel hardware and software solutions to simplify the management of deep memory hierarchies, including tools for to enhance programmability as well as to explore trade-offs between hardware, runtime, and programmer-managed levels of memory.
- Approaches to abstract and standardize memory interfaces to support memory interfaces that are independent from the specific memory implementation technology.

A2-3 Performance Metrics (MR)

Offeror shall estimate or quantify the impact of the proposed technology over industry roadmaps and trends. Offeror shall identify which of the metrics listed below apply to its proposal, and respond to each applicable metric and its associated target requirement (if stated). Offeror is

encouraged to provide alternative meaningful metrics and estimates relevant to their proposal. Offeror shall respond fully to at least one category below (e.g., DRAM Performance Metrics) or provide alternatives.

Quantities specified shall reflect solutions that are productized in 5 to 8 years. These metrics are independent, but a solution that can deliver advances in more than one metric is more desirable than one that solves only one metric at the expense of the others. The most meritorious improvements will make substantial gains over industry roadmaps/trends and substantiate a convincing path to achieving the extreme-scale technology characteristics required by DOE.

In addition to the quantities reflecting beginning-of-decade goals, Offeror shall discuss what progress will be made in the subcontract period and describe what follow-on efforts will be needed to achieve these goals fully. Offeror shall describe in detail how the metric will be evaluated, including the measurement method that will be used (for example, prototype or simulation) and any assumptions that will be made.

A2-3.1 DRAM Performance Metrics

Energy per Bit. This metric is defined as the energy needed to completely run memory, counted per bit of data moved, including a short length of interconnect (~2 cm) and the end-points (the complete memory chip, SerDes, wire losses, and memory controller on the CPU side). Offeror shall specify projected energy per bit for proposed DRAM solutions. Offeror shall describe any assumptions used in calculating this metric and how it will be measured. Seven picojoules per bit is considered the baseline value for this metric.

Aggregate Bandwidth per Socket (DRAM or Suitable Replacement for DRAM). This is defined as the data bandwidth delivered to a processor chip comprising the “socket.” A socket is defined as the smallest physical unit of hardware that contains one processor chip, memory, and at least one network connection to connect to other such units. Offeror shall specify both the peak performance as well as what *measured* performance can be expected for different access patterns, and how bandwidth would be measured. One TB/s is considered the baseline value for this metric.

Memory Capacity per Socket. This metric is defined as the *usable* data capacity per socket. Offeror shall specify the projected DRAM capacity and how it relates to other memory metrics such as bandwidth. Four hundred GB is considered the baseline value for this metric.

FIT Rate per Node. This metric is the total soft-error FIT rate for the portion or fraction of a memory system, per node. A node is defined as the smallest physical unit of hardware that contains processor chip(s), memory, and at least one network connection to connect to other such units. The FIT rate is defined as the number of unrecoverable soft errors per billion hours of operation. This FIT rate is not the sum of FIT rates but assumes additional error detection and recovery, for example, possibly with spare components. Offeror shall describe how the FIT rate will be measured, the cost of recovery from transient errors (time/power), and assumptions used in the fault model. A FIT rate of less than 1000 is considered the baseline value for this metric.

Error Detection. Offeror shall describe technologies that will significantly improve error detection, recovery, and reporting. Offeror shall describe in detail tests that would demonstrate how error detection coverage, reporting, and recovery have been improved over the baseline. ECC + bit steering is considered the baseline for this metric.

Processing in Memory. Offeror shall describe the degree to which any proposed processing in memory technology will reduce data movement in target DOE codes. Offeror shall describe the programming model that will make these features productive for software developers. At a minimum, solutions must include support for atomics in memory.

Programmability/Usability. Offeror shall describe how any proposed memory technology feature would be integrated into a productive programming environment. Offeror shall specify projected improvements in productivity of end users and software developers. At a minimum, solutions must make existing programming models easier to use.

A2-3.2 NVRAM Performance Metrics

NVRAM Integration. Offeror shall describe the cell technology and architecture for NVRAM integration, and at what level of the node architecture this NVRAM would be integrated (for example, tightly integrated devices such as NVRAM-backed register files within a CPU versus loosely integrated SSD-like devices for node-level data storage).

Energy per Bit. This metric is largely the same as the DRAM energy per bit. However, the manner for calculating the energy will be highly dependent on where the NVRAM is integrated into the system. Offeror shall specify projected energy per bit for proposed NVRAM solutions. Offeror shall specify projected read and write energy separately. Offeror shall describe all assumptions and specific tests that would be used to assess this energy metric. Offeror shall explain how the energy per bit and performance relates to wear-out rates for storage cells, if applicable to the proposed NVRAM technology.

Aggregate Bandwidth per Socket. This metric is defined as the data bandwidth delivered to the processor chip that comprises the “socket.” A socket is defined as the smallest physical unit of hardware that contains one processor chip, memory, and at least one network connection to connect to other such units. Offeror shall specify both the peak performance for NVRAM as well as the *measured* performance that can be expected for different access patterns, and how bandwidth would be measured.

Capacity per Socket. This metric is defined as the *usable* data capacity per socket. Offeror shall specify the projected NVRAM capacity. Eight hundred GB is considered the baseline for this metric.

FIT Rate per Node. This metric is the total soft-error FIT rate for the portion or fraction of a memory system, per node. A node is defined as the smallest physical unit of hardware that contains processor chip(s), memory, and at least one network connection to connect to other such units. The FIT rate is defined as the number of unrecoverable soft errors per billion hours of operation. This FIT rate is not the sum of FIT rates but assumes additional error detection and recovery, for example, possibly with spare components. Offeror shall describe how the FIT rate would be measured, the cost of recovery from transient errors (time/power), and the assumptions of their fault model. We are particularly interested in how NVRAM technologies can be made substantially less prone to failure so that they can be used as a reliable backing store to recover from errors/faults at the node level.

Durability. Offeror shall describe the durability of any proposed NVRAM technologies. At a minimum, this description should include a range of total number of read or write operations to a NVRAM technology or device under normal operating conditions expected before permanent

failure. Offeror shall describe any specific hardware or software technologies, such as a translation layer, that will influence the durability as seen by the application.

Error Detection. Offeror shall describe technologies that can significantly improve NVRAM error detection, recovery, and reporting. Offeror shall describe in details tests that would demonstrate how error detection coverage, reporting, and recovery have been improved over the baseline.

Programmability/Usability. Offeror shall describe how any proposed NVRAM memory technology feature would be integrated into a productive programming environment. Offeror shall specify projected improvements in productivity of end users and software developers.

A2-4 Multivendor Integration Strategy (MR)

Offeror shall describe how the proposed memory technology could be integrated into multiple vendors' node architectures.

A2-5 Target Requirements

The requirements below apply to supercomputers that will be deployed at the end of this decade to meet DOE mission needs. As previously stated, Offerors need not address all problem areas, and thus the Offeror need not respond to a TR below if the proposed capability does not address that problem area. In all TR responses that are provided, Offeror should discuss what progress will be made in the next two years and describe what follow-on efforts will be needed to fully achieve these goals. For metrics listed below, the Offeror should describe in detail how the metric will be evaluated, including the measurement method that will be used (for example, simulation or prototype) and any assumptions that will be made.

A2-5.1 Energy per Bit

- **Reduced Energy per Bit (TR-1)**
Energy per bit should be 5 picojoules or less end-to-end. End-to-end is defined as including full path from memory to register on processor chip, including the memory component and cost of accessing the memory cell in the memory component.
- **Greatly Reduced Energy per Bit (TR-2)**
Energy per bit should be 2 picojoules end-to-end.

A2-5.2 Aggregate Delivered DRAM Bandwidth

- **Improved Aggregate Delivered DRAM Bandwidth Per Socket (TR-1)**
Aggregate delivered bandwidth per socket for DRAM or equivalent should be 4 TB/s or greater over a distance of 5 cm or more.
- **Greatly Improved Aggregate Delivered DRAM Bandwidth Per Socket (TR-2)**
Aggregate delivered bandwidth per socket for DRAM or equivalent should be 10 TB/s or greater over a distance of 5cm or more.

A2-5.3 Memory Capacity per Socket

- **Increased DRAM Capacity per Socket (TR-1)**
Memory capacity per socket for DRAM or equivalent should be 1.6 TB or greater with preference for “fast” memory per the aggregate bandwidth requirements above.
- **Greatly Increased DRAM Capacity per Socket (TR-2)**
Memory capacity per socket for DRAM or equivalent should be 4 TB or greater with preference for “fast” memory per the aggregate bandwidth requirements above.

A2-5.4 FIT Rate per Node

- **Improved FIT Rate per Node (TR-1)**
FIT rate per node should not exceed 100.
- **Greatly Improved FIT Rate per Node (TR-2)**
FIT rate per node should not exceed 10.

A2-5.5 Error Detection Coverage and Reporting

- **Reduction in Silent Errors (TR-1)**
Solution should propose and estimate ways to greatly reduce possible rates of silent errors.
- **End-to-End Error Detection and Recovery (TR-2)**
Solution should provide complete end-to-end error detection and recovery, including data paths.

A2-5.6 Advanced Processing in Memory Capabilities

- **Vector Operations and/or Gather/Scatter (TR-1)**
Processing in memory solutions should include vector operations and/or gather/scatter.
- **CPU-independent Processor in Memory (TR-2)**
Offeror should implement a CPU-independent processor-in-memory solution that can be attached to any CPU and function as a memory/PIM part.

A2-5.7 NVRAM Performance Metrics

- **Increased NVRAM Capacity per Socket (TR-1)**
Memory capacity per socket for NVRAM or equivalent should be 3.2 TB or greater with preference for greatly improved reliability.

A2-5.8 Multivendor Integration Strategy

- **Description of the memory integration strategy (TR-1)**
The description should include sufficient detail to demonstrate that integrating the proposed memory technology in an exascale computer can be accomplished using hardware and software interfaces that are available to any vendor. (Note: providing a multivendor integration strategy is mandatory, and this TR addresses the quality of that strategy.)