

DesignForward R&D Statement of Work – DRAFT v.1

March 4, 2013



U.S. DEPARTMENT OF
ENERGY

Office of
Science

This work is performed under the auspices of the U.S. Department of Energy by the Regents of the University of California / Lawrence Berkeley National Laboratory under Contract DE-AC02-05CH11231.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor the Regents of the University of California 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 the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or the Regents of the University of California, and shall not be used for advertising or product endorsement purposes.

Contents

1	INTRODUCTION	5
2	ORGANIZATIONAL OVERVIEW	6
2.1	The Department of Energy Office of Science	6
2.1.1	Advanced Scientific Computing Research Program.....	6
2.2	National Nuclear Security Administration.....	7
2.2.1	Advanced Simulation and Computing Program	7
3	MISSION DRIVERS.....	7
3.1	Office of Science Drivers	7
3.2	National Nuclear Security Administration Drivers.....	8
4	EXTREME-SCALE TECHNOLOGY CHALLENGES	8
4.1	Power Consumption and Energy Efficiency	8
4.2	Concurrency	9
4.3	Fault Tolerance and Resiliency.....	9
4.4	Programmability/Productivity	10
5	APPLICATIONS	10
6	ROLE OF CO-DESIGN	11
6.1	Overview.....	11
6.2	ASCR Co-Design Centers	11
6.3	ASC Co-Design Project	12
6.4	Proxy Apps	12
7	REQUIREMENTS.....	13
7.1	Description of Requirement Categories	13
7.2	Requirements for Research and Development Investment Areas.....	13
7.3	Common Mandatory Requirements.....	13
7.3.1	Solution Description (MR)	13
7.3.2	Research and Development Plan (MR).....	13
7.3.3	Productization Strategy (MR)	14
7.3.4	Staffing/Partnering Plan (MR)	14
7.3.5	Project Management Methodology (MR)	14
7.3.6	Intellectual Property Plan (MR)	15
8	EVALUATION CRITERIA	15

8.1	Evaluation Team	15
8.2	Evaluation Factors and Basis for Selection	15
8.3	Performance Features	15
8.4	Feasibility of Successful Performance	16
8.5	Supplier Attributes	16
8.5.1	Capability	16
8.6	Price of Proposed Research and Development	16
8.7	Alternate Proposals	17

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 the 2008–2010 timeframe. 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.

In July 2011, Argonne National Laboratory, on behalf of the seven aforementioned DOE labs, released a request for information (RFI) with the purpose of providing DOE SC and NNSA with information for planning the DOE Exascale Computing Initiative (ECI). The RFI responses highlighted numerous challenges on the path to exascale and presented many innovative ideas to address those challenges. Funding for DOE ECI has not yet been secured, but DOE has compelling real-world challenges that will not be met by existing vendor roadmaps. Informed by responses to the exascale RFI, DOE SC and NNSA have identified two areas of strategic research and development (R&D) investment that will provide benefit to future extreme-scale applications:

- System Integration
- Interconnect technology

These R&D activities will initially be pursued through a program called DesignForward. The objective of the DesignForward program is to initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme-scale computing. It is recognized that the broader computing market will drive innovation in a direction that may not meet DOE's mission needs. Many DOE applications place extreme requirements on computations, data movement, and reliability. DesignForward 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 consortium expects to establish an ongoing program to continue innovation in these and additional technology areas. Contracts awarded through this RFP process may be eligible for

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

additional funding to add work scope to accelerate further the critical technology R&D if Congress approves funding for this purpose.

The consortium is soliciting innovative R&D proposals in the areas of system integration and interconnect technology that will maximize energy and concurrency efficiency while increasing the performance, productivity, and reliability of key DOE extreme-scale applications. Due to the focus on extreme-scale applications, overall time to solution is also an important consideration. The goal is to begin addressing 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.

Proposals submitted in response to this solicitation must 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 Co-Design Centers to ensure solutions are aligned with DOE needs. While DOE's extreme-scale computer requirements are a driving factor, these projects must 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, initiated with FastForward, and continued with DesignForward, will aid the development of technology that reduces economic and manufacturing barriers to constructing exaflop-sustained systems, but also further DOE'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 also the broader IT industry. This will result in an increase in DOE's ability to leverage commercial developments. It is not DOE's intent 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 Department of Energy 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 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. Additionally, NNSA federal agents provide safe and secure transportation of nuclear weapons and components and special nuclear materials along with other missions supporting the national security.

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, high-performance simulation and computing capabilities are developed to analyze and predict the performance, safety, and reliability of nuclear weapons and to certify their functionality. 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 significantly advance simulation-based scientific discovery, which includes the objective to "provide computing resources at the petascale and beyond, network infrastructure, and tools to enable computational science and scientific collaboration." All the other research programs within DOE 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 workshops² 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 simulations 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

² <http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/>

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 2010 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, it is recognized that 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 pushes the need for computational capability into the exascale level.

4 EXTREME-SCALE TECHNOLOGY CHALLENGES

The HPC community has done extensive analysis³ 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, an overview of the most significant of these challenges is provided.

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

³ 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 \$5–10M annually. To achieve an exascale system using current technology, the annual power cost to operate the system would be above \$2.5B per year with a power load of over a gigawatt (more than many power plants currently produce). To keep the operating costs of such a system in some kind of feasible range, a target of 20 Megawatts has been established.

Achieving the power target for exascale systems is a significant research challenge. Even with optimistic expectations of current R&D activities, *there is at least a factor-of-five gap between what we must have and what current research can provide*. To get the additional factor of five improvements in power efficiency over projections, a number of technical areas in hardware design need to be explored. 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.

Allowing applications to efficiently exploit this level of concurrency is a challenge for which there currently are no good solutions. While many of these challenges are associated with memory and processors, and are the target of the previous Fast Forward solicitation, there are many remaining challenges in the space of overall system architecture, system software, and interconnect technology.

Further complicating this is 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.

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.

4.4 Programmability/Productivity

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

5 APPLICATIONS

The applications that will eventually run on future exascale systems will be a product of a process that starts with the DOE Mission Drivers outlined in Section 3 above, and continues on through identifying key science challenges, developing mathematical models, developing algorithms, and finally producing application software.

While the key challenges in many fields have been effectively articulated in the workshop reports referenced in Section, the development of models, algorithms and software is dependent on how hardware and system software designers respond to the challenges of building an exascale system and vice versa. The aim is to resolve many possible trade-offs in the space of applications and architectures by using the co-design methodology described in the next section.

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 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 are establishing multiple co-design centers that will be used as a vehicle to collaborate with vendors on R&D. The existing and planned co-design centers are at varying stages of deployment at the time of this RFP, and DOE is in the process of developing a strategy for issues such as cross-center collaboration, intellectual property protection, and overall governance. It is expected that much of this will be in-place and documented at the time DesignForward awards are made.

6.2 ASCR Co-Design Centers

In mid-2011, ASCR granted the first three co-design awards, and additional ASCR centers may be established in the future. Each of these co-design centers is a distributed collaboration between multiple national laboratories and university partners. Each center has focused on a specific application that is an important driver for exascale and is using development of that application as a way to explore issues of mathematics, algorithms, computer science, systems software, and of course, hardware in the co-design process. For an overview of the ASCR co-design centers see <http://science.energy.gov/ascr/research/scidac/co-design/> and for specific centers see below.

Center for Exascale Simulation of Combustion in Turbulence	ExaCT
Exascale Co-Design Center for Materials in Extreme Environments	ExMatEx
Center for Exascale Simulation of Advanced Reactors	CESAR

6.3 ASC Co-Design Project

The NNSA labs and ASC program are defining 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. The NSApp project will focus on these established applications as the drivers, and participate in co-design largely through proxy applications. Additional information is available at <https://asc.llnl.gov/codesign/> and <http://proxyapps.lanl.gov/>.

6.4 Proxy Apps

DOE will use proxy applications as the means to interact with our vendor partner(s) during the co-design process. These applications will be used both by the vendors to understand the effects of hardware tradeoffs, and also by integrated code team members and DOE researchers wishing 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.

Proxy apps can be grouped into three categories in increasing sophistication and fidelity to the actual applications (or packages) used in integrated design codes:

- **Kernels:** A combination of one or more fundamental primitives (the micro-benchmarks) integrated into a single executable most likely executing on a single type of device (for example, graphics processing unit (GPU) or multi-core CPU).
- **Skeleton apps:** Reproduced data flow of a simplified physics application with little or no attempt to investigate numerical performance. They are primarily useful in investigating network performance characteristics at large scale.
- **Mini apps:** These apps contain some of the dominant numerical kernels (or subsets thereof) contained in an actual application and produce simplifications of physical phenomena.

A more detailed description of the aforementioned proxy apps categories is available at http://proxyapps.lanl.gov/proxyapps_20130106.pdf.

ASC and ASCR co-design centers are in the process of developing and publishing their proxy apps. Some that are available today, and some previously developed ASCR or ASC benchmarks, are listed below:

ExaCT	http://exactcodesign.org/proxy-app-software/
TORCH	http://crd.lbl.gov/groups-depts/ftg/projects/previous-projects/torch-testbed/
Mantevo	http://www.mantevo.org
NERSC	http://www.nersc.gov/research-and-development/performance-and-monitoring-

SSP tools/sustained-system-performance-ssp-benchmark/

LULESH <https://computation.llnl.gov/casc/ShockHydro/>

7 REQUIREMENTS

7.1 Description of Requirement Categories

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

- MR are performance features essential to DOE requirements. An Offeror must satisfactorily address *all* MR to have its proposal considered responsive.
- TR, 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. TR add value to a proposal and are prioritized by dash number. TR-1 is more desirable than TR-2.

TR-1s and MR are of equal value. The aggregate of MRs and TR-1s form a baseline solution. TR-2s are goals that boost a baseline solution, taken together as an aggregate of MRs, TR-1s, and TR-2s, into the moderately useful solution.

7.2 Requirements for Research and Development Investment Areas

Detailed requirements for each of the two targeted R&D areas of investment are provided as Attachments to 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.

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, with emphasis on how it will increase the performance of key DOE extreme-scale applications relative to energy usage while maintaining or increasing reliability and maintaining or decreasing runtimes.

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. 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.

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.

7.3.3 Productization Strategy (MR)

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.

7.3.4 Staffing/Partnering Plan (MR)

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

7.3.5 Project Management Methodology (MR)

Project management and regular project status reporting are required. Offeror shall describe project management methodology and provide communication plan indicating method of communication (for example, written report, teleconference, and/or face-to-face meeting) and frequency (for example, weekly, monthly, and/or quarterly).

7.3.6 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. Please consult RFP letter for information on Federal regulations concerning IP.

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, as well as Federal government representatives. The Regents of the University of California, manager and operator of Lawrence Berkeley National Laboratory (LBNL), 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 they believe 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. LBNL 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.

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

LBNL 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 Attachment area of technology.

8.3 Performance Features

The Evaluation Team will validate that an Offeror's proposal satisfies the MR. The Evaluation Team will assess how well an Offeror's proposal addresses the TR. An Offeror is not solely limited to discussion of these features. An Offeror may propose other features or attributes if the

Offeror believes they may be 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 two-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
- 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

8.7 Alternate Proposals

An Offeror may submit an alternate proposal. The Evaluation Team may evaluate alternate proposals for award consistent with the preceding information or as otherwise deemed necessary.

ATTACHMENT A: SYSTEM INTEGRATION RESEARCH AND DEVELOPMENT REQUIREMENTS

The scope for the DesignForward system design and integration R&D request for proposals spans both hardware and software for exascale systems that will be deployed in the next decade. DOE wants to work with computer system integrators to refine a vision for the overall system design including the design and integration of hardware, interconnect, packaging, and system software. Responders must detail their vision for overall system design and supported execution model(s) as well as the specific DesignForward system integration research that they propose to enable accelerating industry roadmaps or to provide capability that existing market forces would not ensure. Responders should also discuss their HPC commercialization/business strategy for their company's exascale-related projects. Furthermore, responders should describe how their DesignForward effort could feed into and benefit DOE's upcoming pre-exascale procurements.

The DOE is not interested in developing and deploying one-of-a-kind exascale systems. The intent is to impact the HPC ecosystem more broadly. DOE anticipates that the vendors funded for system integration will work closely with vendors funded by the DOE FastForward Program, who are developing processor and memory technologies. However, DOE requests that responders to this topic include an alternatives analysis of expected component technology options. Integrators should also address their requirements for an interconnection fabric, including whether they are relying on FastForward or other efforts for needed developments. System integration vendors are not expected to team up with FastForward awardees in their responses to DesignForward. However, respondents are required to describe the execution model their systems will support.

DesignForward seeks to fund enhancements that benefit the scalability, resilience, and performance of DOE mission applications. Collaboration with DOE facilities, researchers and co-design efforts will be key to the success of the DesignForward Program. We anticipate two-year DesignForward funding to provide a bridge to the start of DOE's ECI effort.

A1-1 Key Challenges for System Integration

A1-1.1 Execution Model and Overall System Architecture

Large-scale systems are increasingly using a range of node types for different functions, such as computation, I/O, analytics, and compilation. This diversity creates additional complexity for all aspects of system integration including packaging, interconnection networks, system software and execution models. Overall system architecture should be robust and maximize delivered value for likely exascale workflows.

The fundamental problem with current High Performance Computing (HPC) systems and programming methods is that they are based on sequential models of computation organized in coarse-grained concurrent processes that cannot efficiently utilize parallelism. A new model of computation or an execution model is required and must be developed that enables the

programmer to perceive the system as a unified and naturally parallel computer system; not merely as a collection of microprocessors and an interconnection network. The execution model provides the conceptual scaffolding for deriving system elements in the context of and consistent with each other. Ideally, the execution model implements a decision chain where each layer contributes to the optimum determination of when, where, and how data placement, data movement, and operation of a computation are performed. Current execution models do not emphasize nor manage the specific characteristics critical to system efficiency and scalability. This leads to poor use of system resources and early limitations on system effectiveness, as measured in terms of metrics such as GFLOPS per watt (GFLOPS/W).

A1-1.2 Energy Utilization

Energy and power are key design constraints for exascale systems. Power management and energy cost reduction techniques to minimize or constrain power used while maintaining reliable, predictable behavior are needed.

A1-1.3 Resilience and Reliability

The integrated RAS (reliability, accessibility and serviceability) system is critical to the stability and usability of today's supercomputers. Development of exascale RAS systems having the ability to identify, to contain, and to overcome faults quickly is of paramount importance. Incorporating fault tolerance into all layers of the hardware/software stack, including the application programming environment and tools, is needed to ensure that applications can continue to progress even with system faults.

A1-1.4 Data Movement through the System

The scalability and performance of DOE applications is tightly linked to the latency and bandwidth of data movement through the system. Higher performance, higher efficiency (lower energy use) data movement technologies are needed for system-wide data movement between nodes and between cabinets.

A1-1.5 Density

Packaging density largely determines the needed floor space for an exascale system. Methods are needed for improving system density, while managing heat dissipation, are needed to keep the size of the complete exascale system between 100 and 300 cabinets.

A1-1.6 System Software

System software includes the OS, file system, runtime, system monitoring, resource management, job management, accounting and user management, and security. These software components must be improved to cope with the scale of an exascale system, to ensure its overall reliability, and to incorporate power management into their actions. Mechanisms must be developed to support integrated, seamless maintenance of system software and monitoring within the overall computing center.

A1-1.7 Programming Environment

Programming the languages supported, and the required software including compilers, debugging tools, runtime, libraries, and performance tuning. Programming environment advances are needed to improve system integration, programmability, code portability, and usability. Programming environments and tools should facilitate the porting of current DOE code base and ensure that the development of new codes is as straightforward and efficiently as possible.

A1-2 Areas of Interest

The following are examples of objectives and technologies that could be considered in system integration R&D proposals that address DOE's extreme-scale computing needs. Some of the items below may only apply to certain architectures, and some may be mutually exclusive. *Proposals are not limited to these areas, and alternative topics are encouraged.*

A1-2.1 Overall System Architecture

- Advances that simplify changing or upgrading specific node aspects (e.g., processors, memory, coprocessors) especially in the face of faults that may degrade or kill nodes
- Mechanisms to increase flexibility in resource utilization such as ways to share memory capacity across nodes
- Mechanisms to mitigate the tension between production system use, which primarily entails large jobs, and software development for the system, which involves non-computational tasks such as compilation and small jobs for testing and debugging
- Advances that facilitate compiling for a mix of heterogeneous nodes
- Advances to support isolation and flexibility in resource association
- Techniques to support efficient scheduling of diverse resource types
- Coping with heterogeneity in an optimal fashion
- Scalable, adaptive, and unobtrusive monitoring, with real-time analysis of platform state
- Real-time autonomic platform management under production workloads, gracefully handling events without requiring immediate human intervention

A1-2.2 Energy Utilization

- Advances that improve the power efficiency of the system
- Advances in measurement, runtime control and application control of power utilization

- Advances that support system-wide and site-wide power management
- Techniques to reduce cooling energy requirements

A1-2.3 Resilience and Reliability

- Advances that improve the resiliency or reliability of the system, for example, improved fault detection, containment, correction, and response time across the software stack (not just hardware).
- Advances that enable dynamic adaptation to a constantly changing system
- Advances that demonstrate hardware/software resilience synergies to improve overall time to solution
- Techniques to improve fault detection accuracy and root cause analysis (e.g., fewer undetected errors) or to reduce their cost
- Framework for representing hardware and system software dependencies, for interpretation of failure modes and autonomic reasoning about remaining recovery paths

A1-2.4 Data Movement through the System

- Advances that allow extremely low-latency multi-hop messages
- Improvements to the performance and energy efficiency of messaging, remote memory access, and collective operations
- Advances that demonstrate the optical/copper tradeoffs to improve the data movement across the system
- Reliable low-energy, long-distance data movement.
- Efficient data movement for computation and also across levels of storage hierarchy
- Mechanisms to avoid contention, and to provide QoS guarantees

A1-2.5 Density

- Advances in packaging and cooling that improve the density of the system.
- Advances that address potential scaling and concurrency limits that restrict the ultimate size of the system.
- Lightweight packaging and cooling techniques

A1.2.6 System Software

- Advances that significantly improve resilience support
- Advances that allow efficient dynamic load balancing
- Advances that provide coarse and fine grain power management across the system

- Support for job and resource management at the scale of a million nodes and a billion cores
- Techniques to manage diverse node types and nodes with heterogeneous resources
- Topology, resilience, and energy aware scheduling of resources

A1.2.7 Programming Environment

- Description of the supported programming model(s)
- Advances that significantly improve application resilience
- Advances that help scale applications to a billion-way concurrency
- Development of programming tools that work with million node systems
- Hardware or software techniques to simplify writing and debugging application software and promote portability to different node types and system architectures

A1-3 Performance Metrics

Offeror shall estimate or quantify the impact of the proposed technology over industry roadmaps and trends. This information shall be provided for all of the metrics listed below. If Offeror determines that a particular metric is not applicable to the technology being proposed, then Offeror shall explain why they believe the metric is not relevant and shall replace that metric with an alternate *meaningful* metric.

Quantities specified should reflect solutions that are productized in the 2020 timeframe. These metrics are independent, but a solution that can deliver advances in more than one metric is more desirable than one that addresses 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.

The list of metrics includes:

- Cabinet power requirements
- Scaling limits of the overall system design
- Error detection, correction, and coverage of hard and soft error types
- Expected resilience of the system (MTTI, mean time to system interrupt)
- Computational capacity per cabinet
- Data motion overhead (and avoided data motion)
- Bandwidths and latencies for inter- and intra-node data movement for all layers of the memory hierarchy
- Optimization of global communications performance

- Avoidance/minimization of contention for communication inter- and intra-node
- Total cost of ownership
- Total delivered value for exascale workflows

A1-4 Target Requirements

The requirements below apply to supercomputers that will be deployed by DOE Office of Science and NNSA at the end of this decade to meet the 2 programs' mission needs. As previously stated, Offerors need not address all problem areas, and thus the Offeror need not respond to a given TR below if the proposed capability does not address that problem area. However, the Offerors must describe the proposed execution model and overall system architecture. 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 achieve these goals fully. Offeror should describe in detail how metrics will be evaluated, including the measurement method that will be used (for example, simulation or prototype) and any assumptions that will be made.

A1-4.1 Execution Model and Overall System Architecture (TR-1)

Offeror should describe the execution model(s) supported. An execution model has the following characteristics:

An execution model has the following characteristics:

- Provides the governing principles for system design, operation, management, and application implementation;
- Impacts across all layers of the system design and operation;
- Provides a conceptual framework for the co-design of all system layers;
- Supports the notion of the operation “decision chain,” i.e.,
 - When, where, and why each operation is performed and
 - How every component layer contributes to this decision process
- Permits reasoning and design decisions in addressing critical efficiency factors.

Offeror should describe their overall system architecture, including both hardware and software components. The hardware description should include a description of node types and interconnection networks between them. Offeror should include an alternatives analysis of expected component technology options in processors, memory, and interconnect technology that identifies gaps in current technology options and features that are needed from component developers for successful system integration. Proposed R&D should identify work that will address open questions, lead to identification of the best alternative, and mitigate risk.

A1-4.1 Energy Utilization (TR-1)

A system that achieves high performance on a broad range of DOE applications while minimizing energy use. Solutions should target 20 MW at system scale while maintaining or improving system reliability.

A1-4.2 Resilience and Reliability (TR-1)

Mean Time to System Interrupt (TR-1). Processor designs should make advances that lead to a mean time to a system interrupt requiring user or administrator action of one day or longer in a 2020 exascale system.

Fault Recovery Overhead (TR-1). The overhead to handle automatic fault recovery should not reduce application performance by more than half.

A1-4.3a Data Movement within a cabinet (TR-2)

The performances of applications depend upon many factors, such as message injection rates and contention. This system integration attachment only presents a few data movement targets and these tend to be minimal requirements. Offerors who are interested in the aggressive interconnect targets should consult attachment 2.

The point-to-point bandwidth and latency between any two nodes in the same cabinet should be greater than 50 GB/s with latency of 500ns or less.

A1-4.3b Data Movement across the machine (TR-2)

The point-to-point bandwidth and latency between the farthest nodes in the topology should be greater than 20 GB/s with latency less than 3 us.

Bi-section bandwidth should be between $\frac{1}{4}$ to $\frac{1}{2}$ total system node injection bandwidth.

A1-4.4 Density (TR-2)

To keep the system size and facility infrastructure at a manageable level, an exascale system should fit in a floor space of 15,000 sq. ft. (between 100 and 300 cabinets).

A1-4.5 System Software (TR-1)

Solutions will need a system software stack. Offeror should describe the components in their system software and how these will be improved to handle the scale, resilience, and power management challenges of an integrated exascale system.

A1-4.6 Programming Environment (TR-1)

Solutions will need an integrated software ecosystem that supports the development of new applications, the migration of existing applications, application maintenance, application resilience, and application portability, while enabling DOE scientists to achieve high performance with no more effort than is required for today's high-end supercomputers. Offeror should describe in detail how this will be accomplished and the improvements to be made.

ATTACHMENT B: INTERCONNECT RESEARCH AND DEVELOPMENT REQUIREMENTS

The scope for the DesignForward interconnect R&D request for proposals is focused on innovations required to most effectively enable highly scalable and performant applications, integrate large parallel systems, improve system reliability, and reduce energy usage. Responders must describe how the specific interconnect R&D that they propose will accelerate industry roadmaps or develop a capability that existing market forces would not provide. Collaboration with DesignForward system design and integration projects, FastForward node technology R&D projects (subject to NDA constraints), DOE researchers, and co-design efforts will be key to the success of the DesignForward interconnect R&D projects. We anticipate two-year DesignForward funding to provide a bridge to the start of DOE's ECI effort.

In this RFP, the term interconnect refers to the communication network(s) used to integrate individual nodes together into a single parallel system, and includes both hardware (network interface controllers, switches, etc.) and software (network stack) components. The assumption is that at a node level, one or more processors are integrated on a board (or a fraction of a board), with memory, and one or more network interface controllers (NICs) (may be add-in component or integrated into the processor).

B1-1 Key Challenges for Interconnect Technology

The main challenge for interconnect technology has always been to provide sufficiently high network bandwidth and message rate, while maintaining as low a latency as possible. However, more important than the raw interconnect performance in these areas, is how effectively applications are able to utilize the network. Historically, many interconnects have been built to provide these capabilities without regard to the communication mechanisms being used by the applications running on the system, leading to a bad semantic match between the communication primitives required by the applications and those provided by the network. The resulting poor utilization of the interconnect is bad for application performance, energy efficiency or both.

Technology advances have had a positive impact on both latency and message rate; however, the relative bandwidth increases continue to trail the increases in processor performance and memory bandwidth. Additionally, the potential move to lighter-weight processing cores may have a dramatic impact on the message rate for some communication models (i.e. potentially reduced performance for MPI matching). One key challenge is how to provide continued sufficient interconnect performance in light of these technology trends.

Other challenges are also moving to the forefront. For example, with the expected increase in mixed-model programming (e.g. MPI + X) it is important that the network stack provides the primitives necessary to support the efficient, simultaneous progression of multiple communication libraries in the presence of a large number of threads. Power is becoming a critical focus. Most current interconnects show little difference between idle and fully loaded states. Applications striving for low latency adopt polling for the fastest notifications, which keep the processors' power consumption high as well. Improving methods for blocking (i.e.

non-polling) notifications could help reduce power usage while maintaining acceptable performance. Many interconnect fabrics provide multiple routes from one node to another (e.g. due to fabric topology as well as nodes with multiple NICs) which could reduce congestions, but hide that capability to avoid deadlock or delivery ordering issues. Finally, the increasing size of high performance computing systems is beginning to stretch existing network management and resiliency capabilities. New software and hardware architectures will be required to solve these problems.

B1-2 Areas of Interest

The following are examples of objectives and technologies that could be considered in interconnect R&D proposals that address DOE's extreme-scale computing needs. Some of the items below may only apply to certain architectures, and some may be mutually exclusive. *Proposals are not limited to these areas, and alternative topics are encouraged.* Also, suggested subtopics within each area are not intended to be limiting and alternate areas will also be considered.

B1-2.1 Overall Interconnect Architecture

The interconnect will be a critical component in the design of Exascale systems, and involves a complex integration of many elements. DOE seeks a comprehensive solution that covers the key aspects of interconnect architecture, including:

- Host network interfaces
- Switches and routers
- Fabric topology
- Software stacks
- Programming models

B1-2.2 Interconnect Integration with Processor and Memory

The main function of the interconnect is to provide the processors access to remote memory, as such, the mechanisms used to integrate the network with the rest of the node architecture (processor and memory) can have a dramatic impact on how and what type of features can be effectively supported by the communication runtime. There are many available options for integrating the network. Logically, the network can be connected as an I/O device, or as a peer on the processor or memory network. Additionally, the NIC can be either coherent or non-coherent with the processor cache hierarchy. Physically, the network can be integrated on-die, on-package or on-board. Each integration choice comes with a set of trade-offs.

The DOE is interested in novel ways to integrate the network to provide the highest levels of scalability and reliability. The following are examples of capabilities and technologies that could be considered in interconnect R&D proposals.

- Remote atomic operation support
- Virtual memory interactions to reduce overhead of memory registration
- Advanced message completion notification
- Communication protocol offload/onload tradeoffs
- Integration with complex memory hierarchies which may include deep cache hierarchies, local scratchpads, on-package memory, DRAM and the possibility of fast memory dedicated to the NIC
- Active messages support and lightweight mechanisms to control computation based on network events
- Processor core enhancements to better support network integration
- Improve fault isolation in the communication stack to improve resiliency

B1-2.3 Multiple Communication Library Progression and Interaction

Although DOE's application suite is largely MPI-everywhere today, it is likely that new programming models will be leveraged as applications consider moving away from MPI everywhere as part of an exascale transition. Although MPI+X is a likely path for many applications, leveraging other programming models, such as PGAS libraries and languages (e.g., UPC, Coarray Fortran or Fortran 2008 Coarrays, Global Arrays, and OpenSHMEM) in MPI applications is also likely. With that in mind, we are interested in proposals to:

- Refine network abstraction layers to allow a more natural integration between communication methodologies, improving the ability of MPI and other languages or communication libraries, such as UPC, GASNet, Global Arrays, Coarray Fortran, and OpenSHMEM, to interoperate
- Improve the interaction between communication libraries and on-node tasking or threading interfaces such as OpenMP or Threading Building Blocks

B1-2.4 Interconnect Fabrics and Management

The interconnection fabric is required to provide system wide data transfer and management capabilities, while operating under possible power constraints and in the presence of system component failures. The fabric design shall be as energy efficient as feasible and have sufficient resilience to remain functional during a variety of failures. The fabric must allow for the interaction between all system components at suitably high data rates and richness in topology to best support the execution of DOE exascale applications. The fabric should allow for the system to achieve the highest possible system availability, and support the highest possible utilization for the diverse exascale workload.

The DOE is interested in the construction of novel interconnection fabrics to support extreme scalability for exascale. The following are examples of possible innovatory capabilities and features that are of interest

- Modeling and simulation capabilities for extreme-scale fabric design
- High radix hierarchical topologies to reduce overall network diameter, and to allow for increased locality

- Mechanisms for fine grained power management of network components to allow for reduced levels of operation, as well as reduced topologies when under power constraints
- Automatic regeneration of new paths in the presence of link failures
- Adaptive checksum capabilities that can be exploited by software when transmission failures occur
- Management software that is able to manage all pieces of the fabric while healthy and in a failing state and scale to millions of endpoints and routers.
- Novel techniques for monitoring, and logging dynamic changes of the fabric
- Congestion control support that is capable of eliminating or greatly reducing the impact of global catastrophic congestion and loss of throughput caused by local network hot spots.
- Trade-offs between enabling multi-pathing to avoid congestion and relaxed ordering due to the use of alternate paths
- Adaptive routing techniques that take advantage of local and global network information to alleviate the onset of network congestion without the side effect of deadlock.

B1-2.5 Protocol Support

Under the assumption that the interconnect will become highly integrated with other components, it will be important to maintain open standards wherever possible to avoid fully proprietary solutions that stifles competitively-driven progress and increases costs. In addition, the system interconnect does not exist in isolation. A fully integrated interconnect should provide capabilities for efficient interoperation with high-volume networks for I/O or communication, such as Ethernet or InfiniBand. Topics of interest include:

- Network protocols to be used for exascale interconnects at all layers of the network stack.
- Open protocol standards or open intermediate or high-level APIs that would provide access to alternative, low-level hardware interconnects.
- Network stack models for software support of the interconnect
- Efficient interoperation with other network technologies through approaches such as extension, encapsulation or routing.
- Opportunities to leverage high-volume/commodity-driven networking components to reduce costs.

B1.2.6 Scalability

It is highly desirable for the interconnect to enable rather than inhibit the scalability of applications and system services that require high-performance communication. Design decisions at the interconnect level can have direct impact on the amount of non-network resources required to enable efficient communication at full scale. For example, flow control and message buffering strategies that are based on usage patterns and that are independent of system size or the number of communicating endpoints have been shown to significantly reduce memory usage and increase network efficiency. Some fundamental operations, such as translating virtual to physical endpoint addresses, may also be infeasible as the number of network endpoints continues to grow.

The DOE is interested in advanced features that enable the scalability of upper-level protocols by minimizing the use of resources needed by applications and system services to efficiently use the interconnect fabric. The following are examples of possible innovative capabilities and features that are of interest:

- Flow control strategies that maximize network efficiency for multiple upper-layer protocol consumers
- Mechanisms that minimize the use of on-node resources dedicated to network data movement services
- Approaches to network addressing that scale independent of the number of endpoints or network clients

B1-3 Performance Metrics

Offeror shall estimate or quantify the impact of the proposed technology over industry roadmaps and trends. This information shall be provided for all metrics listed below which are relevant within the scope of Offeror's proposal. If Offeror determines that a particular metric is not applicable to the technology being proposed, then Offeror shall explain why they believe the metric is not relevant and shall replace that metric with an alternate *meaningful* metric.

Quantities specified should reflect solutions that are productized in the 2020 timeframe. Offerors may optionally provide a roadmap over the next decade.

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.

- Message rate per initiator (e.g. MPI rank) and total for the NIC
- Message latency as a function of the size of the fabric (i.e. give minimum latency and explain how it varies as the fabric grows)
- Message bandwidth as a function of message size for a single communication path
- Peak and maximum sustained (peak minus overhead) bandwidth for interconnect links
- Cross-section bandwidth as a function of fabric size
- Architectural and physical limits (e.g. fabric size, port counts, message size)
- Raw bit error rate
- Residual bit error rate. That is, the bit error rate for uncorrectable but detectable transmission errors (including errors from the network end points, i.e. NICs, PCIe, etc.) after any error-correction or recovery mechanisms have been applied (e.g. forward error correction, link-layer retry, end-to-end retry, etc.)
- Worse-case sustained bandwidth maintained at the quoted raw bit error rate for a continuous flow of maximum-sized messages

- Undetected bit error rate (end-to-end)
- Energy usage at standby, at maximum message rate for small messages, and at full bandwidth

If the interconnect has multiple operational modes (e.g. MPI and SHMEM, varying link types) with significantly different characteristics, provide the metrics for each.

Offeror should show how they estimate performance on these metrics.

B1-4 Target Requirements

The DesignForward Interconnect program seeks to fund network Research and Development that benefits the scalability, resilience, and energy efficiency performance of exascale HPC systems for DOE mission applications. The requirements below apply to interconnect R&D that will be deployed in future DOE exascale systems. Offerors need not address all problem areas, and thus the Offeror need not respond to a given TR below if the proposed capability does not address that problem area. In all TR responses that are provided, Offeror should discuss what deliverables will be provided in the next two years and describe what follow-on efforts will be needed to achieve these goals fully. Offeror should describe in detail how metrics will be evaluated, including the measurement method that will be used (for example, design study reports, simulation or emulation results, or prototype software, etc.) 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.

B1-4.1 Overall Interconnect Architecture (TR-1)

Offeror should describe a comprehensive overall interconnect architecture including, host interface, switches and routers, fabric topology, software stacks, protocols and expected programming models. Offeror should include an alternatives analysis of expected technology options and identify any gaps inhibiting successful integration. Proposed R&D may include design studies, simulation efforts, development projects or prototyping activities necessary to further development of a total interconnect solution. Evaluation will be based on the completeness of the proposed approach.

B1-4.2 Processor/Network/Memory Integration (TR-1)

One of the key questions for DesignForward is to understand the design choices that are possible for levels of integration into the interconnection network fabric. NICs can be integrated into the processor socket, into the package, or on the board. Responses in this problem area should propose a design study that quantifies the performance, scalability, reliability, and cost tradeoffs for these different integration options. Processor designers have the further option to integrate a NIC either coherently or non-coherently with the processor cache hierarchy.

The target interconnection network performance for the node is:

- MPI Applications: 500 Million messages per second

- PGAS Applications: 2 Billion messages per second
- Node injection/ejection BW: 1TB/sec, each direction (assuming 5-10TF/s nodes)

B1-4.3 Multiple Communication Library Progression and Interaction (TR-2)

Offeror should describe how enhancements to the networking abstraction layers would provide for improved interoperability between multiple upper layer communication libraries and/or describe how interactions between communication libraries and on-node tasking or threading interfaces will be accomplished.

B1-4.4 Fabrics (TR-1)

The target interconnection network performance is:

- Node injection/ejection BW: 1TB/sec, each direction (assuming 5-10 TF/s nodes)
- Message Latency: 500 nanoseconds, nearest neighbor; 2 microseconds furthest neighbor
- Bi-section Bandwidth: $\frac{1}{4}$ to $\frac{1}{2}$ total system node injection bandwidth
- Residual bit error rate: no more than 1 error per month of operation.
- Undetected error rate: no more than a 50% chance of an undetected error within a 6 year system life.

B1-4.5 Protocol Support (TR-2)

Protocols selected for use in the interconnect should provide:

- Open interface standards or non-discriminatory licensing enabling alternative (commercial) implementations of key components.

B1.4.6 Scalability (TR-1)

The target system scale in terms of interconnection network endpoints is:

- 100,000 to 1,000,000