U.S. Department of Energy
Office of Science
and
National Nuclear Security Administration

Preliminary Conceptual Design for an Exascale Computing Initiative

November 2014

Contents

Executive Summary

Leadership in high-performance computing (HPC) and large-scale data analysis will advance national competitiveness in a wide array of strategic sectors, including basic science, national security, energy technology, and economic prosperity. The U.S. semiconductor and HPC industries can develop the necessary technologies for an exascale computing capability early in the next decade. However, without an integrated approach to the development of hardware, software, and applications, these new resources will go untapped by the science, industry, and DOE mission critical research communities that need them—HPC technology development will fail to meet the needs of the HPC research communities and the researcher will lose the opportunity to develop codes to take advantage of the predicted capabilities. Therefore, the Department of Energy's Exascale Computing Initiative (ECI) will focus and integrate efforts across industry, academia and government to address the technical challenges of exascale computing.

The ECI's goal is to deploy, by 2023, capable exascale computing systems. This is defined as a hundred-fold increase in sustained performance over today's computing capabilities, enabling applications to address next-generation science, engineering, and data problems to advance Department of Energy (DOE) Office of Science and National Nuclear Security Administration (NNSA) missions. The plan includes three distinct components: Exascale Research, Development and Deployment (ExaRD) (see Appendix 1 for detailed technical descriptions); Exascale Application Development (ExaAD) to take full advantage of the emerging exascale hardware and software technologies from ExaRD; and Exascale Platform Deployment (ExaPD) to prepare for and acquire two or more exascale computers.

This document presents a preliminary strategy and contains limited descriptions of project management details. Further details, such as change control, quality assurance, and alternatives analyses, will be provided in subsequent, detailed plans, once project baselines are established.

1 Introduction

The Exascale Computing Initiative will, by 2023, enable deployment of highly productive exascale computing platforms, which are defined as having a hundred-fold increase in sustained performance over today's computing capabilities, for DOE and NNSA mission-critical applications. Exascale computing will extend HPC to a new scale; it will lead to discoveries, make possible the exploration of ensembles of simulations of complex physical phenomena, and provide markedly increased fidelity and quality in simulation results – with quantified uncertainties. Exascale computing will allow scientists to interpret experimental data in conjunction with large-scale modeling and simulation of the experiment to gain increased understanding of physical phenomena. It also will help decision makers act on the knowledge gained from these simulations.

Reaching exascale, however, will be difficult. Incremental improvement over today's technologies would lead to computers with unacceptable reliability and power requirement of hundreds of MW for an exascale system; power constraints necessitate technological innovations to bring that level down, to around 20 MW. The traditional means of increasing processor performance through increased transistor density and faster internal clocks is reaching its end. Performance gain must thus be achieved through increased concurrency and heterogeneity. At the same time, to access sufficient memory, an application will need to negotiate complex memory hierarchies while increasing data movement. These changes in processor and memory technologies necessitate new solutions in system software, programming environments, and applications algorithms and libraries to deal with the anticipated concurrency, heterogeneity, and the memory wall, which is the growing disparity between processors speed and the time it takes to access memory. Finally, the sheer number of newly developed components and their operating parameters in an exascale computer will lead to reliability concerns: soft errors and other faults are expected to increase dramatically. Research and development in computing systems resilience is required to address these issues.

2 Project Background

Since the beginning of the digital era, the U.S. federal government has made pivotal investments in the computer industry at critical times when progress was stagnating. We are once again at a critical turning point in HPC technology where innovations in hardware and software architectures are necessary to drive future advances in computing performance. While the computing industry will continue to advance technologies, the marketplace will drive them in directions orthogonal to HPC interests. Past experience has demonstrated that partnerships between the government and industry have led to incorporating beneficial technologies into product lines in ways that adherence to market forces would have precluded. At this critical juncture the Government needs to directly influence future HPC technology that will result in the design and development of highly-energy-efficient, scalable exascale systems and extreme scale applications. The government must actively engage industry in HPC technology development, as market forces will not support national needs.

Over the past several years, the DOE has become aware that future-generation systems will require significant design changes. The designs industry has proposed to enable necessary energy efficiency through massive parallelism to a degree not experienced previously. Continued growth in processing performance requires breakthroughs that address the Von Neumann memory bottleneck, reducing power consumption, and solving problems of computing at these scales. As such, the DOE's approach to HPC technology challenges is aimed at a broad spectrum of capabilities over the next few years. Meeting these challenges requires a significant investment by the federal government involving strong leadership from DOE headquarters and close coordination among government, national laboratories, industry, and academia.

A critical component of a federally funded effort in exascale computing is concurrent applications software research and development to optimally exploit emerging computing architectures. These applications include those that support nuclear weapons stockpile stewardship, scientific discovery, energy technology innovation, renewable electrical generation and distribution, nuclear reactor design and longevity, data assimilation and analysis, and climate modeling. The applications efforts must address the full spectrum of computing, including terascale and petascale as well as exascale applications.

3 Justification of Mission Need

Investment in exascale computing supports the DOE Strategic Plan 2014-2018 Strategic Objective 3: "Deliver the scientific discoveries and major scientific tools that transform our understanding of nature and strengthen the connection between advances in fundamental science and technology innovation. DOE will continue to pursue scientific discoveries that lay the technological foundation to extend our understanding of nature and create new technologies that support DOE's energy, environment, and security missions. Areas of concentration include: Advanced scientific computing to analyze, model, simulate, and predict complex phenomena, including the scientific potential that exascale simulation and data will provide in the future."

The development of capable exascale computing responds to the science and national security mission requirements of two DOE organizations: the Office of Science and the National Nuclear Security Administration, which have responsibilities for advancing U.S. science and maintaining the reliability and safety of the U.S. nuclear stockpile, respectively. To address these mission needs, computational applications critical to advance DOE missions in science, energy and national security are targeted for initial use on capable exascale systems. Examples include: discovery and characterization of new functional materials, improved climate models, and stockpile stewardship. Some efforts have been initiated under ExaAD and are briefly described in the following paragraphs.

3.1 Computational Materials Science

The Computational Materials Sciences effort will advance U. S. leadership in the development of computational codes for materials sciences and engineering. This software development effort will start in FY 2015 within DOE's Office of Basic Energy Sciences and will involve teams of theorists, computational experts, and experimentalists with expertise in synthesis, characterization, and processing/fabrication of materials. In the first year, the computational materials sciences teams will focus on basic science necessary to develop research-oriented, open-source, experimentally validated software and the associated databases required to predictively design materials with specific functionality on emerging hardware, with a goal of delivering production-level software for use on exascale systems by 2023.

3.2 Next-Generation Climate Models

The Accelerated Climate Modeling for Energy (ACME) project,[1] being launched in FY 2015, comprises eight national laboratories, the National Center for Atmospheric Research, four academic institutions, and industry, and is sponsored and led by the Earth System Modeling program within DOE's Office of Biological and Environmental Research. ACME will address critical Earth-system science questions, modeling and related capabilities, and can be flexibly applied by the DOE research community to address mission-specific climate change applications, such as those identified in the report, *U.S. Energy Sector Vulnerabilities to Climate Change and Extreme Weather*.[2]

By 2023, the ACME project will develop and use next-generation computational Earth system and climate models at the leading edge of scientific knowledge and computational capabilities. These models will provide tools for executing and analyzing climate and Earth system simulations that address the most critical scientific questions for the nation and DOE. ACME Earth system and climate models will include non-hydrostatic atmospheric modeling (less than 10 km resolution), more sophisticated ice sheet physics, and a new approach for terrestrial modeling that uses plant functions instead of plant types for more physical representation of biology.

3.3 Stockpile Stewardship

The nation's confidence in the nuclear stockpile relies on high fidelity simulations of all of the physical processes occurring within a nuclear weapon and the processes that support the design, production, maintenance, and evaluation of the nuclear arsenal, including support of life extension programs and weapons dismantlement. Integrated design codes (IDCs) model various aspects of nuclear weapons, and can have several million lines of code to accurately reflect the integrated multi-physics occurring in a nuclear weapon. The accuracy of these IDCs underpins confidence in the nuclear deterrent. For the National Nuclear Security Administration to exploit the multi-level parallelism demanded by emerging architectures leading to exascale, significant new code development would be required over the next 7-10 years.

4 Barriers to Achieving Capable Exascale Computing

Achieving capable exascale computing will entail overcoming significant technology challenges through a coordinated pursuit by government, industry, and academia. It also will require developing extreme-scale applications, each with substantially new capabilities.

---

[1] http://www.climatemodeling.science.energy.gov/news/accelerated-climate-modeling-energy-acme-project-plan-available

[2] http://energy.gov/downloads/us-energy-sector-vulnerabilities-climate-change-and-extreme-weather

The key challenges[3] that must be addressed to achieve exascale are:

- *Parallelism Challenge:* Design systems to provide efficient exploitation of the extreme levels of parallelism that will be necessary at exascale, both in terms of developing and executing applications programs. This includes providing application developers with programming models and computer/user interfaces that promote ease of use, or more specifically, that facilitate the development of optimally performing, energy-efficient applications while insulating the user from the complexities of unprecedented parallelism.
- *Resilience Challenge:* Achieve system-level resilience to both permanent and transient faults and failures so that applications can "work through" these problems to successful, accurate, reliable execution and completion.
- *Energy Challenge:* Achieve energy efficiencies so that, when run at the targeted computational rates, the entire system will operate within affordable power budgets. Achieving the required energy efficiency will pose challenges that cut across all levels of hardware and software.
- *Memory and Storage Challenge*: Develop memory and storage architectures that provide high capacity for accessing/storing information that support DOE applications at anticipated computational rates, and that operate within the targeted exascale power envelope.

The exascale challenges listed above force consideration of new execution models (i.e., the detailed descriptions of the workflows of operations for carrying out HPC-based simulations and exploiting the results) that enable programmers to perceive the system as a unified and naturally parallel computer system, not as a collection of microprocessors and an interconnection network. Current petascale execution models do not recognize and manage the features/attributes specific to a particular system and will not scale to the exascale level of concurrency, leading to inefficient use of each system's resources and premature saturation of system efficiency. The result of new execution models will be new energy and concurrency efficient, secure computer systems.

The Advanced Scientific Computing Advisory Committee, in order to further resolve the above four exascale challenges and identify possible technical approaches, conducted an in-depth study in 2012, which identified the Top Ten Exascale Research Challenges[4] that must be addressed to make a productive, economically viable exascale system.

The realization of an exascale system will involve complex tradeoffs between hardware (processors, memory, energy efficiency, reliability, interconnectivity), software (programming models, scalability, data management, productivity) and algorithms. To achieve this, a total systems approach will be used, in conjunction with co-design of hardware and software, to plan and execute the ECI – from

[3] http://www.cse.nd.edu/Reports/2008/TR-2008-13.pdf

[4] http://science.energy.gov/~/media/ascr/ascac/pdf/meetings/20140210/Top10reportFEB14.pdf

research to hardware acquisition. In addition to the exascale challenges described above, the ECI responds to emerging security challenges in:

- *Cybersecurity: How we will develop the next generation of computer architectures and software to ensure that security is designed and integrated into systems from the beginning, across all levels?*
- *Global Supply Chain: How we will ensure that systems and their components we procure recognize and respond to issues created by the fact that the supply chain for HPC hardware and software extends beyond U.S. borders?*

5 Technical Approach

The ECI plan includes significant integrated investments in the following three areas:

1. *Exascale Research, Development and Deployment (ExaRD):*
    - *Exascale Co-Design Centers and Beyond Exascale,* exploratory research to co-design hardware and software architecture, for a set of DOE mission-relevant applications. This area may be expanded to include R&D efforts in other federal agencies;
    - *Software Technology Research and Development*, aimed at specific hardware and software technologies;
    - *Vendor Research and Development*, aimed at developing exascale node and system architectures;
2. *Exascale Application Development (ExaAD): Readiness to Use Capable Exascale Systems*, initiating the development of a suite of exascale applications software packages that will be operational in 2023 to ensure maximal scientific and engineering impact of the exascale systems; and
3. Exascale Platform Deployment (ExaPD)*: Coordinated Acquisition Strategy* for exascale platforms, including long-lead site preparations and system platforms.

The following sections describe these areas in more detail. (For a comprehensive description of the proposed ExaRD research thrusts and technical approach, see Appendix 1).

ExaRD: Exascale Co-Design Centers and Beyond Exascale: Given the well-understood challenges of achieving exascale computing, application code developers must recognize trends and opportunities of emerging architectures. At the same time, platform providers must gain deeper understanding of the intended uses of the computers they are developing. Consequently, DOE has committed early investments in exploratory research to co-design of hardware and architecture, software stacks, uncertainty quantification, and numerical methods and algorithms for a set of DOE mission-relevant applications, to determine technical tradeoffs in the design of exascale hardware, system software, and application codes. Lessons learned from these co-design efforts will be incorporated into the emerging exascale systems and will inform the development of exascale applications software.

There are three ongoing DOE Office of Science Advanced Scientific Computing Research (ASCR) and one NNSA Advanced Simulation and Computing (ASC) exascale co-design centers, which cover materials in extreme environments, simulation of advanced nuclear fusion reactors, turbulent combustion, and stockpile stewardship. ECI proposes to add new co-design centers – in collaboration with other federal agencies – to ensure the exascale needs of those agencies will be addressed in capable exascale systems.

ExaRD: Software Technology Research, Development and Deployment: These activities will support software technology projects in industry, laboratories, and academia. To achieve the full potential of exascale computing, a software stack must be developed that includes new programming models and metrics for evaluating system status with a focus on new and revised application implementations. The scope of the software effort will span the spectrum, from low-level operational software to high-level application development environments. This includes the software infrastructure to support data management for the DOE and NNSA computational science activities at exascale.

To date, Office of Science and NNSA exascale software research and development projects have focused on early-stage efforts. Software research and development efforts must substantially increase to enable the deployment of highly productive, exascale systems in FY 2023.

ExaRD: Vendor Research and Development: ASCR and ASC have initiated the Fast Forward and Design Forward subprograms to form partnerships with key vendors to initiate and accelerate the R&D of node architectures and exascale system designs, and to ensure the commercialization of promising emerging technologies. This public-private partnership between industry, ASCR and ASC supports the development of innovative technologies critical to constructing sustained-exaflop systems, and to reducing economic and manufacturing barriers to their commercial production. As they mature, these initial programs will transform into Path Forward and System Design Phases. The Path Forward Phase will fund computer vendors to develop component technologies needed to build exascale nodes, including the required software. The System Design Phase will fund computer vendors to perform the required engineering, research and development projects that will eventually result in an exascale computer.

ExaAD: Application Readiness to Use Capable Exascale Systems: Exascale systems will enable opportunities for pioneering scientific, engineering and national security progress, and must be closely coordinated with exascale application development. The key issues from the application development and computational science perspectives include: extreme parallelism, reliability and resiliency, scaling to larger systems, and data-intensive science.

The exascale application readiness effort will focus on research and development programs that address critical DOE missions and involve communities that are well versed in HPC use and the current state of petascale application readiness.

ExaPD: Coordinated Acquisition Strategy: Together, DOE and NNSA facilities are responsible for acquiring and transitioning to operation forefront computational capabilities, including a series of pre-exascale systems in fiscal years 2016 through 2022 and the initial capable exascale system in FY 2023-2024. As part of this effort, a dedicated team of computational scientists at the ASCR and NNSA facilities will further identify mission-critical requirements and develop and implement strategies for application readiness, outreach and training to ensure broad use of the new platforms. Acquisition and deployment of production-ready systems will follow Office of Science and NNSA policies and procedures for major scientific/computing facility upgrades.

6 Project Management

The Office of Science and NNSA have a long track record of successfully executing large, technically complex, scientific projects. Following past exemplars, the ECI will be organized as a project and will be executed within a tailored framework that follows DOE Order 413.3B (DO 413.3B), which defines critical decision points, overall project management, and requirements for control of baselined schedule and cost. A single federal official will have overall responsibility for execution of the project, will report to the cognizant DOE Headquarters program offices (Office of Science and NNSA), and will be accountable to an Acquisition Executive, as defined in DO 413.3B. Project execution will be governed by a baselined schedule and cost envelope, in accordance with DO 413.3B, and will follow the defined processes for change control and management of contingency once the performance baseline for ECI is established.

The management agreement between DOE and NNSA that governs execution of the project is codified in the Memorandum of Understanding (MOU) titled "Memorandum of Understanding between the U.S. Department of Energy, Office of Science and the U.S. Department of Energy, National Nuclear Security Administration, Office of Defense Programs for the Coordination of Exascale Activities" that was signed on April 13, 2011. The MOU builds upon long-standing cooperation between the two organizations. (For full MOU, see Appendix 2.)

Because of the breadth and complexity of the development and deployment of an exascale computer, an Integrated Project Team (IPT) will be established through an IPT charter with defined roles and responsibilities. Membership in the IPT will draw on technical expertise from across Office of Science and NNSA laboratories and may evolve based on the needs of the project as it progresses. The IPT will support the federal official, who will lead the IPT through the lifetime of the project.

6.1 Communications

A communications plan will be established to ensure effective internal and external communications during the life of ECI. ECI is committed to promoting communication among team members and with key stakeholders including OMB, OSTP, Congress, other federal agencies, professional societies and HPC users, including industry. Software engineering tools and processes and the establishment of crosscutting councils will be used in ExaRD in part to coordinate efforts. Other communication methods could include regular meetings within each component's subprojects; annual project review meetings with representatives from all of the subprojects; quarterly ECI IPT calls that also involve the chairs of the component IPTs; monthly meetings between DOE, NNSA and national laboratories involved in ECI; quarterly meetings either virtual or face-to-face between the leads of the co-design centers and applications development teams; weekly meetings between the ASCR and ASC ECI project team; briefings as needed for other agencies, legislators and the executive branch; presentations to ASCAC, the Council on Competitiveness and professional societies to engage the broader research community. Other communication technologies such as websites and portals will be used to provide progress updates.

7 Integrated Baseline

7.1 Scope Baseline

The Exascale Computing Initiative comprises the following activities or subprojects. Each subproject may establish its own baseline that will roll up to the overall ECI baseline**.**

WBS 1.0 Exascale Research, Development and Deployment (ExaRD): Perform the research, development, and design efforts that are required to deploy exascale systems by 2023, including the vendor exascale systems and the extreme-scale software environment. This is the largest element of this planning effort. (See Appendix **1** for detailed technical descriptions of the subprojects.)

WBS 2.0 Exascale Application Development (ExaAD): Develop extreme-scale applications with increased physics fidelity and advanced numerical algorithms to scale to and take full advantage of the emerging exascale hardware and software technologies from ExaRD.

WBS 3.0 Exascale Platform Deployment (ExaPD): Prepare for and acquire two or more exascale computers. Included in this component are the planning and execution of the procurements and any necessary site preparation, such as providing electrical power and cooling.

Activities in these interdependent components will be actively coordinated through the project management approach described in Section 6. Performers of these components will include the DOE national laboratories, academia, small businesses, and the computer industries.

Once the ECI budget has been finalized, the key performance parameters will be established. For this stage of project definition, we have identified measureable outcomes, which are listed in the following sections.

7.2 System Performance

In 2024, after the capable exascale machine is accepted and before operations, benchmark codes will be run with the goal of using approximately 20 MW while achieving a hundred-fold increase over the performance of today's Titan and Sequoia systems. Other measurable outcomes, such as mean time to failure or interrupt, also will be considered.

7.3 Application Development

In Section 3, three mission-critical programs were identified. By 2023, these efforts will produce applications that will be used in system acceptance and in early science runs on capable exascale computers. It is anticipated that additional exascsale application development program will be initiated.

7.4 Software Development

The primary deliverable from the ExaRD component is the exascale Bring Up System (XBUS). The XBUS environment will be an integration of most of the software developed under the ECI effort. This will include runtime environment, libraries, programming environments and tools. It will enable the testing of the exascale software on current computer platforms and eventually the testing of the initial exascale platforms. By 2023, the critical components of XBUS will be integrated with the vendor software stack that is deployed with an exascale system.

7.5 Platform Deployment

Platform deployment will be managed as a subproject and will use a tailored DOE Order 413.3B process. The outcome for platform deployment is that this subproject completes on time and within the ECI projected scope and schedule.

7.6 Schedule Baseline

Figure 1 is a notional timeline for the ECI. The timeline shows approximately when various efforts in ExaRD will occur. The bar labeled "Software Technology" encapsulates all ExaRD programs described in Section 7.1, except co-design. Anchoring the timeline is the availability of several research prototypes, labeled P0, P1, and P2, in the timeline. Presently, it is envisioned that P0 will be a node prototype. P0 may have some vendor-specific interfaces that allow DOE researchers to make measurements and study application behaviors. P1 is envisioned to be a petascale-capable cabinet, consisting of P0 or its variant with interconnects, with a minimal vendor software stack that can be used for system software and applications testing. Finally, P2 is an exascale prototype that, when properly scaled, would satisfy outcomes for the final exascale initial-delivery system. This prototype is envisioned to have a nearly complete system software and software environment that the vendor would provide in a final system to allow for advanced testing of DOE-developed software stack, math libraries, and other tools and packages and to enable application readiness. P2 may or may not be a scaled-down version of the eventual exascale system; all prototypes, however, are expected to inform the platform acquisition process in terms of hardware specifications, application readiness, power requirement, site preparation, etc. The B1 system (Build 1) in the following timeline represents the first delivery of cabinets for an exascale computer. Full system delivery and checkout is expected to take up to a year, based on experience with current-generation petascale systems.

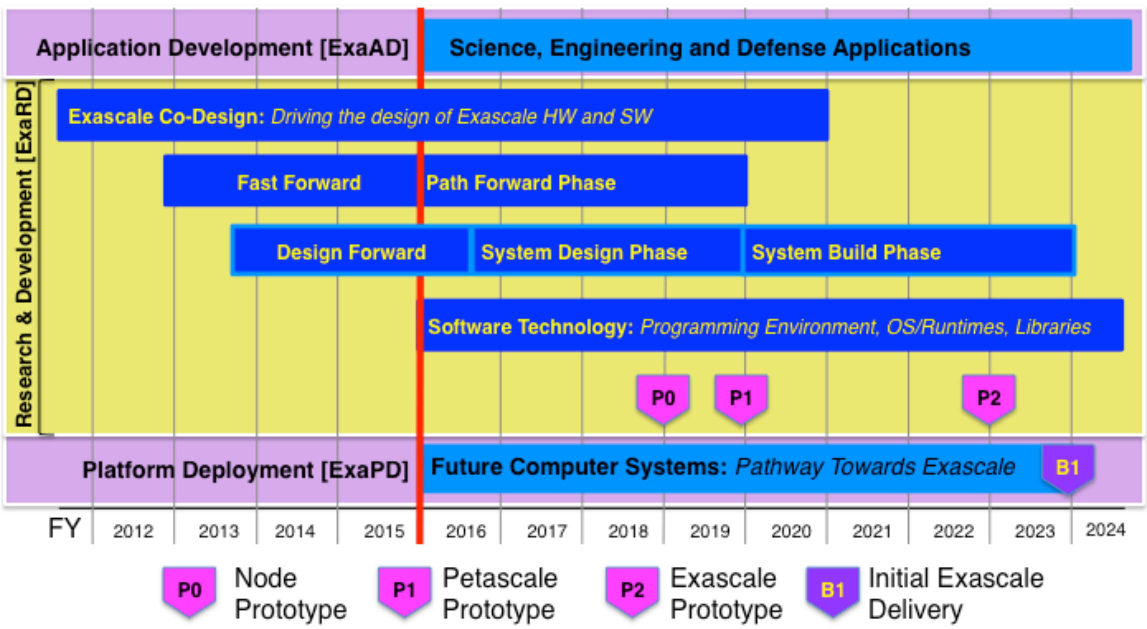

Figure 1 - Schedule Baseline

7.7 Milestones

The preliminary project milestones are identified in Table 1.

FY	Event
2015	Application Development Begins
2015	Release Software Technology FOAs
2015	Release Path Forward RFP
2015	Release System Design RFP
2015	Release Co-Design Center FOAs or process renewals
2015	*Multi-agency ECI Review (annual event)*
2016	**ECI Funding Starts**
2016	Initiate all new contracts
2016	*ECI Project Review (annual event)*
2018	Delivery of P0
2019	Delivery of P1
2019	Release of phase 2 RFP
2019	Facilities release joint RFP for B1
2020	Initiate System Build Phase
2021	Site preparation activities initiated at facilities
2022	Delivery of P2
2023	Start Deployment of B1
2025	**ECI Ends**

Table 1

7.8 Preliminary Work Breakdown Structure

The ECI Project will be managed in accordance with DOE Order 413.3B. There are three major subprojects and each subproject baseline will be established using an integrated analysis of logic-driven, resourced-loaded activities that follow its respective work breakdown structure (WBS).

The WBS for the ECI is consistent with the conventional structure for DOE-managed projects, under which each major WBS element that represents a major deliverable is further subdivided into a hierarchy of lower-level elements. The lowest-level WBS elements describe specific activities that must be performed to achieve the deliverables represented at the higher levels.

Many scheduled activities for the ECI rely on the development and delivery schedules of the vendors, from which the computing equipment is acquired and appropriate vendor milestones will be incorporated into the ECI WBS, as they are identified, to define and track those dependencies.

The preliminary high-level WBS is shown in Figure 2 below and a detailed description of the individual ExaRD tasks and subtasks can be found in Appendix 1.

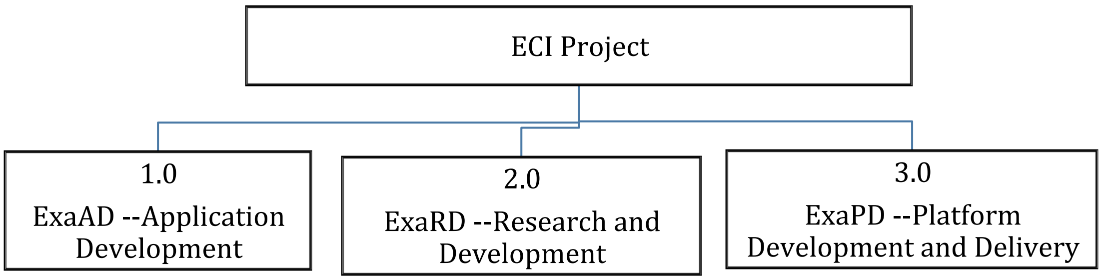

Figure 2 - Preliminary High-level Work Breakdown Structure (WBS)

Risk Management

Below we tabulate some of the high-level risks of ExaRD overall and present possible risk mitigation strategies.

Risk	**Impact (H/M/L)**	**Likelihood (H/M/L)**	**Mitigation**
Insufficient funding within either NNSA or SC	H	M	Prioritize elements within the ECI and, if necessary, de-scope ECI and/or spread the acquisition costs over a longer time period
Full system is unreliable	H	M	Invest in robust, multilayered approaches to manage or resolve faults
Failure to achieve critical integration of DOE-developed software into vendor software	H	M	Early development of XBUS and integration with funded vendors software environments
Software environments do not satisfy DOE application needs	H	M	Determine the workload requirements of critical applications as early as possible
Key algorithms that do not scale may not have timely, suitable alternatives	H	L	Early investments in exploratory algorithms research

Risk	Impact (H/M/L)	Likelihood (H/M/L)	Mitigation
Exascale computer architecture departs significantly from expected designs, after significant R&D in software, tools, and algorithms were invested based on the expected designs	H	M	A clearly defined communication channel with vendors and effective communication so that the ECI participants are well-informed regarding evolving architectures and new directions

Table 2

8 External Stakeholders

8.1 HPC Vendors

For many years DOE has been successful in working with the computer industry to adapt commodity technology into highly usable systems without incurring the high costs of fully custom designs. With the anticipated challenges of the exascale system, vendor engagement throughout the entire life of the ECI is critical to success. The co-design philosophy will guide the ECI community in its close work with vendors throughout the process so that DOE mission requirements are considered when vendors plan and conduct their in-house R&D. The co-design process also is expected to increase the availability of vendor products to the DOE community for testing, prototyping, measurements, and analysis, and for further iterations and refinements by both DOE and industry developers.

8.2 Medium and Small Businesses

Medium and small businesses are expected to play an important role in refining and deploying software to a broader community of HPC users across industry, academia, and government. Partnering with these companies early in the development process will facilitate and hasten the hand-off of successful technologies. DOE will explore a number of strategies to engage this community throughout implementation of the ECI that builds upon the success of efforts such as the Small Business Innovative Research and Technology Transfer programs, incubators and vendor partnership programs, Cooperative R&D Agreements, and consortia.

8.3 Academia

As exascale computing presents difficult challenges in scalability, programmability, reliability, massive parallelism, and data-intensive sciences, many researchers in universities and other academic research organizations are ready to embark on research in applied mathematics and computer science to tackle these topics. Research environments in universities are conducive to stimulating new ideas and often provide longer time horizons for research. Universities also are an essential source for future HPC workers. ECI researchers in the DOE Laboratories will collaborate with university and industry colleagues to ensure promising high-quality, innovative research is examined, adopted, hardened, and applied to DOE mission applications. ASC program has already embarked on this path with its Predictive Science Academic Alliance Program (PSAAP) which can serve as a model for the ECI academic collaborations.

8.4 Other Federal Agencies

DOE exascale efforts are expected to benefit HPC applications across the Federal government. To this end, DOE is soliciting input from and engagement by other Federal Agencies on technical requirements for exascale computing, including supporting joint research and development efforts, and expects to maintain frequent formal and informal communications on progress in ECI with other federal agencies.