APEX 2020 Technical Requirements Document

for

Crossroads and NERSC-9 Systems

LA-UR-15-28541
SAND2016-4325 O

APEX 2020:
Technical Requirements

1   Introduction

Los Alamos National Security, LLC (LANS), in furtherance of its participation in the Alliance for Computing at Extreme Scale (ACES), a collaboration between Los Alamos National Laboratory and Sandia National Laboratories; and in coordination with the Regents of the University of California (UC), which operates the National Energy Research Scientific Computing (NERSC) Center residing within the Lawrence Berkeley National Laboratory (LBNL), is releasing a joint Request for Proposal (RFP) for two next generation systems, Crossroads and NERSC-9 under the Alliance for application Performance at EXtreme scale (APEX), to be delivered in the 2020 time frame.

The successful Offeror will be responsible for delivering and installing the Crossroads and NERSC-9 systems at their respective locations. While it is our preference to award both the Crossroads and NERSC-9 subcontracts to a single Offeror, awards may be made to separate Offerors. Awards will be made by LANS on behalf of ACES and by UC on behalf of NERSC. In total there will be four subcontracts, one Non-Recurring Engineering subcontract for each of ACES and NERSC (see Section 4 herein) and the "build" (system) subcontracts (one issued by LANS for Crossroads and one issued by UC for NERSC-9). The technical requirements in this document describe joint requirements wherever possible. The Offeror shall respond with a single proposal that contains distinct sections showing how and where their proposed Crossroads and NERSC-9 systems differ. Alternative solutions for hardware, software, and/or architecture may also be included in the Offeror's proposal.

An Offeror's Technical Proposal shall include narrative and graphics as appropriate, describing its proposed solutions to technical aspects of the project as seen in numbered sections of this Technical Requirements Document. An Offeror shall incorporate its proposed solutions directly into each section of the Technical Requirements Document to the greatest practical extent. The Technical Requirements Document is provided in MS Word format to facilitate this proposal requirement. The evaluation committee will make no presumption of technical capability when evaluating Offeror responses. Offerors must address each section in a materially responsive manner. The response shall clearly describe the role of any subcontractor(s) and the technology or technologies, both hardware and software, and value added that the subcontractor(s) provide, where appropriate.

If an Offeror chooses to submit alternative solutions to the APEX RFP, complete, separate and distinct proposal packages (to include all applicable/required proposal documents) must be submitted for each

alternative. Failure to comply with these proposal submission instructions may cause an Offeror's proposal(s) to be downgraded.

The scope of work and technical specifications for any subcontracts resulting from this RFP will be negotiated based on the requirements and the options in this document and the successful Offeror's responses.

Crossroads and NERSC-9 each have maximum funding limits over their system lives, to include all design and development, site preparation, maintenance, support and analysts. Total ownership costs will be considered in system selection. The Offeror must respond with a configuration and pricing for both systems.

Application performance and workflow efficiency are essential to these procurements. Success will be defined as meeting APEX 2020 mission needs while at the same time serving as a pre-exascale system that enables our applications to begin to evolve using yet to be defined next generation programming models. The advanced technology aspects of the APEX systems will be pursued both by fielding first of a kind technologies on the path to exascale as part of system build and by selecting and participating in strategic NRE projects with the Offeror and applicable technology providers. A compelling set of NRE projects will be crucial for the success of these platforms, by enabling the deployment of first of a kind technologies in such a way as to maximize their utility. The NRE areas of collaboration should provide substantial value to the Crossroads and NERSC-9 systems with the goals of:

- Increasing application performance.
- Increasing workflow efficiency.
- Increasing the resilience, and reliability of the system.

The details of the NRE are more completely described in section 4.

To support the goals of application performance and workflow efficiency an accompanying whitepaper, "APEX Workflows," is provided that describes how application teams use HPC resources today to advance scientific goals. The whitepaper is designed to provide a framework for reasoning about the optimal solution to these challenges. (The Crossroads/NERSC-9 workflows document can be found on the APEX website.)

1.1   Crossroads

The Department of Energy (DOE) National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) Program requires a computing system be deployed in 2020 to support the Stockpile Stewardship Program. In the 2020 timeframe, Trinity, the first ASC Advanced Technology System (ATS-1), will be nearing the end of its useful lifetime. Crossroads, the proposed ATS-3 system, provides a replacement, tri-lab computing resource

for existing simulation codes and provides a larger resource for ever-increasing computing requirements to support the weapons program. The Crossroads system, to be sited at Los Alamos, NM, is projected to provide a large portion of the ATS resources for the NNSA ASC tri-lab simulation community: Los Alamos National Laboratory (LANL), Sandia National Laboratories (SNL), and Lawrence Livermore National Laboratory (LLNL), during the 2021-2025 timeframe.

In order to fulfill its mission, the NNSA Stockpile Stewardship Program requires higher performance computational resources than are currently available within the Nuclear Security Enterprise (NSE). These capabilities are required for supporting stockpile stewardship certification and assessments to ensure that the nation's nuclear stockpile is safe, reliable, and secure.

The ASC Program is faced with significant challenges by the ongoing technology revolution. It must continue to meet the mission needs of the current applications but also adapt to radical change in technology in order to continue running the most demanding applications in the future. The ASC Program recognizes that the simulation environment of the future will be transformed with new computing architectures and new programming models that will take advantage of the new architectures. Within this context, ASC recognizes that ASC applications must begin the transition to the new simulation environment or they may become obsolete as a result of not leveraging technology driven by market trends. With this challenge of technology change, it is a major programmatic driver to provide an architecture that keeps ASC moving forward and allows applications to fully explore and exploit upcoming technologies, in addition to meeting NNSA Defense Programs' mission needs. It is possible that major modifications to the ASC simulation tools will be required in order to take full advantage of the new technology. However, codes running on NNSA Advanced Technology Systems (Trinity and Sierra) in the 2019 timeframe are expected to run on Crossroads. In some cases new applications also may need to be developed. Crossroads is expected to help technology development for the ASC Program to meet the requirements of future systems with greater computational performance or capability. Crossroads will serve as a technology path for future ASC systems in the next decade.

To directly support the ASC Roadmap, which states that "work in this timeframe will establish a strong technological foundation to build toward exascale computing environments, which predictive capability may demand," it is critical for the ASC Program to both explore the rapidly changing technology of future systems and to provide systems with higher performance and more memory capacity for predictive capability. Therefore, a design goal of Crossroads is to achieve a balance between usability of current NNSA ASC simulation codes and adaptation to new computing technologies.

1.2   NERSC-9

The DOE Office of Science (SC) requires a high performance production computing system in the 2020 timeframe to provide a significant upgrade to the current computational and data capabilities that support the basic and applied research programs that help accomplish the mission of DOE SC.

The system also needs to provide a firm foundation for future exascale systems in 2023 and beyond; a need that is called out in the DOE's Strategic Plan 2014-2018, that calls out for "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."

NERSC Center supports nearly 6000 users and about 600 different application codes from a broad range of science disciplines covering all six program offices in SC. The scientific goals are well summarized in the 2012-2014 series of requirements reviews commissioned by the Advanced Scientific Computing Research (ASCR) office that brought together application scientists, computer scientists, applied mathematicians, DOE program managers and NERSC personnel. The 2012-2014 requirements reviews indicated that compute-intensive research and research that attempts scientific discovery through the analysis of experimental and observational data both have a clear need for major increases in computational capability and capacity in the 2017 timeframe and beyond. In addition, several science areas also have a burgeoning need for HPC resources that satisfy an increased compute workload and provide strong support for data-centric workflows and real-time observational science. More details about the DOE SC application requirements are in the reviews located at: http://www.nersc.gov/science/hpc-requirements-reviews/.

NERSC has already begun transitioning the SC user base to energy efficient architectures, with the procurement of the NERSC-8 "Cori" system. In the 2020 time frame, NERSC also expects a need to address early exascale hardware and software technologies, including the areas of processor technology, memory hierarchies, networking technology, and programming models.

The NERSC-9 system is expected to run for 4-6 years and will be housed in the Wang Hall (Building 59) at LBNL that currently houses the "Cori" system and other resources that NERSC supports. The system must integrate into the NERSC environment and provide high bandwidth access to existing data stored by continuing research projects. For more information about NERSC and the current systems, environment, and support provided for our users, see http://www.nersc.gov.

1.3   Schedule

The following is the tentative schedule for the Crossroads and NERSC-9 systems.

Table 1 Crossroads/NERSC-9 Schedule

	Crossroads and NERSC-9
RFP Released	Q3CY16
Subcontracts (NRE/Build) Awarded	Q4CY16
On-site System Delivery Begins	Q2CY20
On-site System Delivery Complete	Q3CY20
Acceptance Complete	Q1CY21

2   Mandatory Requirements

An Offeror shall address all Mandatory Requirements in a materially responsive manner and its proposal shall demonstrate how it meets or exceeds each one. A proposal will be deemed non-responsive/unacceptable, will be rejected, and will not be considered further if each and every one of the following Mandatory Requirements is not met.

2.1.1   The Offeror shall provide a detailed full system architectural description of both the Crossroads and NERSC-9 systems, including diagrams and text describing the following details as they pertain to the Offeror's system architecture(s):

- Component architecture – details of all processor(s), memory technologies, storage technologies, network interconnect(s) and any other applicable components.
- Node architecture(s) – details of how components are combined into the node architecture(s). Details shall include bandwidth and latency specifications (or projections) between components.
- Board and/or blade architecture(s) – details of how the node architecture(s) is integrated at the board and/or blade level. Details should include all inter-node and inter-board/blade communication paths and any additional board/blade level components.
- Rack and/or cabinet architecture(s) – details of how board and/or blades are organized and integrated into racks and/or cabinets. Details should include all inter rack/cabinet communication paths and any additional rack/cabinet level components.
- Platform storage – details of how storage is integrated with the system, including a platform storage architectural diagram.

- ▪ System architecture – details of how rack or cabinets are combined to produce system architecture, including the high-speed interconnects and network topologies (if multiple) and platform storage.
- ▪ Proposed floor plan – including details of the physical footprint of the system and all of the supporting components.

2.1.2   The Offeror shall provide a detailed description of the proposed software eco-system, including a high-level software architectural diagram including the provenance of the software component, for example open source or proprietary and support mechanism for each (for the lifetime of the system including updates).

2.1.3   The Offeror shall describe how the system does or does not fit into the Offeror's long-term product roadmap and a potential follow-on system acquisition in the 2025 and beyond timeframe.

3   Target Design Requirements

This section contains detailed system design targets and performance features. It is desirable that the Offeror's design meets or exceeds all the features and performance metrics outlined in this section. If a Target Design Requirement cannot be met, it is desirable that the Offeror provide a development and deployment plan, including a schedule, to satisfy the requirement.

The evaluation committee will make no presumption of technical capability when evaluating Offeror responses to Target Design Requirements. Offerors that do not address the Target Design Requirements in a materially responsive manner will be downgraded. The Offeror may also propose any hardware and/or software architectural features that will provide improvements for any aspect of the system.

3.1   Scalability

The scale of the system necessary to meet the needs of the application requirements of the APEX laboratories adds significant challenges. The Offeror shall propose a system that enables application performance up to the full scale of the system. Additionally, the system proposed should provide functionality that assists users in obtaining performance at up to full scale. Scalability features, both hardware and software, that benefit both current and future programming models are essential.

3.1.1   The system shall support running jobs up to and including the full scale of the system.

3.1.2   The system shall support launching an application at full system scale in less than 30 seconds. The Offeror shall describe factors (such as executable size) that could potentially affect application launch time.

3.1.3   The Offeror shall describe how applications launch scales with the number of concurrent launch requests (pers second) and scale of each launch request (resources requested, such as the number of scheduleable units etc.), including information such as:

- All system-level and node-level overhead in the process startup including how overhead scales with node count for parallel applications, or how overhead scales with the application count for large numbers of serial applications.

- Any limitations for processes on compute nodes from interfacing with an external work-flow manager, external database or message queue system.

3.1.4   The system shall support thousands of concurrent users and more than 20,000 concurrent batch jobs. The system shall allow a mix of application or user identity wherein at least a subset of nodes can run multiple independent applications from multiple users. The Offeror shall describe details, including limitations of their proposed support for this requirement.

3.1.5   The Offeror shall describe all areas of the system in which node-level resource usage (hardware and software) increases as a job scales up (node, core or thread count).

3.1.6   The system shall utilize an optimized job placement algorithm to reduce job runtime, lower variability, minimize latency, etc. The Offeror shall describe in detail how the algorithm is optimized to the system architecture.

3.1.7   The system shall include an application programming interface to allow applications access to the physical-to-logical mapping information of the job's node allocation – including a mapping between MPI ranks and network topology coordinates, and core, node and rack identifiers.

3.1.8   The system software solution shall provide a low jitter environment for applications and shall provide an estimate of a compute node operating system's noise profile, both while idle and while running a non-trivial MPI application. If core specialization is used, the Offeror shall describe the system software activity that remains on the application cores.

3.1.9   The system shall provide correct numerical results and consistent runtimes (i.e. wall clock time) that do not vary more than 3% from run to run in dedicated mode and 5% in production mode. The Offeror shall describe strategies for minimizing runtime variability.

3.1.10   The system's high speed interconnect shall support a high messaging bandwidth, high injection rate, low latency, high throughput, and independent progress. The Offeror shall describe:

- The system interconnect in detail, including any mechanisms for adapting to heavy loads or inoperable links, as well as a description of how different types of failures will be addressed.
- How the interface will allow all cores in the system to simultaneously communicate synchronously or asynchronously with the high speed interconnect.
- How the interconnect will enable low-latency communication for one- and two-sided paradigms.

3.1.11   The Offeror shall describe how both hardware and software components of the interconnect support effective computation and communication overlap for both point-to-point operations and collective operations (i.e., the ability of the interconnect subsystem to progress outstanding communication requests in the background of the main computation thread).

3.1.12   The Offeror shall report or project the system's node injection/ejection bandwidth.

3.1.13   The Offeror shall report or project the system's bit error rate of the interconnect in terms of time period between errors that interrupt a job running at the full scale of the system.

3.1.14   The Offeror shall describe how the interconnect of the system will provide Quality of Service (QoS) capabilities (e.g., in the form of virtual channels or other sub-system QoS capabilities), including but not limited to:

- An explanation of how these capabilities can be used to prevent core communication traffic from interfering with other classes of communication, such as debugging and performance tools or with I/O traffic.
- An explanation of how these capabilities allow efficient adaptive routing as well as a capability to prevent traffic from different applications interfering with each other (either through QoS capabilities or appropriate job partitioning).
- An explanation of any sub-system QoS capabilities (e.g. platform storage QoS features).

3.1.15 The Offeror shall describe specialized hardware or software features of the system that accelerate workflows or components of workflows such as data analysis or visualization, and describe any limits their scalability on the system. The hardware shall be on the same high speed network as the main compute resources and shall have equal access to other compute resources (e.g. file systems and platform storage). It is desirable that the hardware have the same node level architecture as the main compute resources, but could, for example, have more memory per node.

3.2 System Software and Runtime

The system shall include a well-integrated and supported system software environment. The overall imperative is to provide users with a productive, high-performing, reliable, and scalable system software environment that enables efficient use of the full capability of the system.

3.2.1 The system shall include a full-featured Linux operating system environment on all user visible service partitions (e.g., front-end nodes, service nodes, I/O nodes). The Offeror shall describe the proposed full-featured Linux operating system environment.

3.2.2 The system shall include an optimized compute partition operating system that provides an efficient execution environment for applications running up to full-system scale. The Offeror shall describe any HPC relevant optimizations made to the compute partition operating system.

3.2.3 The Offeror shall describe the security capabilities of the operating systems proposed in technical requirements 3.2.1 and 3.2.2.

3.2.4 The system shall include efficient support for dynamic shared libraries, both at job load time and during runtime. The Offeror shall describe how applications using shared libraries will execute at full system scale with minimal performance overhead compared to statically linked applications.

3.2.5 The system shall include resource management functionality, including job migration, backfill, targeting of specified resources (e.g., platform storage), advance and persistent reservations, job preemption, job accounting, architecture-aware job placement, power management, job dependencies (e.g., workload management), and resilience management. The Offeror may propose multiple solutions for a vendor-supported resource manager and should describe the benefits of each.

3.2.6   The system shall support jobs consisting of multiple individual applications running simultaneously (inter-node or intra-node) and cooperating as part of an overall multi-component application (e.g., a job that couples a simulation application to an analysis application). The Offeror shall describe in detail how this will be supported by the system software infrastructure (e.g., user interfaces, security model, and inter-application communication).

3.2.7   The system shall include a mechanism that will allow users to provide containerized software images without requiring privileged access to the system or allowing a user to escalate privilege. The startup time for launching a parallel application in a containerized software image at full system scale shall not greatly exceed the startup time for launching a parallel application in the vendor-provided image.

3.2.8   The system shall include a mechanism for dynamically configuring external IPv4/IPv6 connectivity to and from compute nodes, enabling special connectivity paths for subsets of nodes on a per-batch-job basis, and allowing fully routable interactions with external services.

3.2.9   The Offeror shall provide access to source code, and necessary build environment, for all software except for firmware, compilers, and third party products. The Offeror shall provide updates of source code, and any necessary build environment, for all software over the life of the subcontract.

3.3   Software Tools and Programming Environment

The primary programming models used in production applications in this time frame are the Message Passing Interface (MPI), for inter-node communication, and OpenMP, for fine-grained on-node parallelism. While MPI+OpenMP will be the majority of the workload, the APEX laboratories expect some new applications to exercise emerging asynchronous programming models. System support that would accelerate these programming models/runtimes and benefit MPI+OpenMP is desirable.

3.3.1   The system shall include an implementation of the MPI version 3.1 (or most current) standard specification. The Offeror shall provide a detailed description of the MPI implementation (including specification version) and support for features such as accelerated collectives, and shall describe any limitations relative to the MPI standard.

3.3.2   The Offeror shall describe at what parallel granularity the system can be utilized by MPI-only applications.

3.3.3   The system shall include optimized implementations of collective operations utilizing both inter-node and intra-node features where appropriate, including MPI_Barrier, MPI_Allreduce, MPI_Reduce, MPI_Allgather, and MPI_Gather.

3.3.4   The Offeror shall describe the network transport layer of the system including support for OpenUCX, Portals, libfabric, libverbs, and any other transport layer including any optimizations of their implementation that will benefit application performance or workflow efficiency.

3.3.5   The system shall include a complete implementation of the OpenMP version 4.1 (or most current) standard including, if applicable, accelerator directives, as well as a supporting programming environment. The Offeror shall provide a detailed feature description of the OpenMP implementation(s) and describe any expected deviations from the OpenMP standard.

3.3.6   The Offeror shall provide a description of how OpenMP 3.1 applications will be compiled and executed on the system.

3.3.7   The Offeror shall provide a description of any proposed hardware or software features that enable OpenMP performance optimizations.

3.3.8   The Offeror shall list any PGAS languages and/or libraries that are supported (e.g. UPC, SHMEM, CAF, Global Arrays) and describe any hardware and/or programming environment software that optimizes any of the listed PGAS languages supported on the system. The system shall include a mechanism to compile, run, and debug UPC applications. The Offeror shall describe interoperability with MPI+OpenMP.

3.3.9   The Offeror shall describe and list support for any emerging programming models such as asynchronous task/data models (e.g., Legion, STAPL, HPX, or OCR) and describe any system hardware and/or programming environment software it will provide that optimizes any of the supported models. The Offeror shall describe interoperability with MPI+OpenMP.

3.3.10   The Offeror shall describe the proposed hardware and software environment support for:

- Fast thread synchronization of subsets of execution threads.
- Atomic add, fetch-and-add, multiply, bitwise operations, and compare-and-swap operations over integer, single-precision, and double-precision operands.
- Atomic compare-and-swap operations over 16-byte wide operands that comprise two double precision values or two memory pointer operands.
- Fast context switching or task-switching.
- Fast task spawning for unique and identical task with data dependencies.

- Support for active messages.

3.3.11 The Offeror shall describe in detail all programming APIs, languages, compliers and compiler extensions, etc. other than MPI and OpenMP (e.g. OpenACC, CUDA, OpenCL, etc.) that will be supported by the system. It is desirable that instances of all programming models provided be interoperable and efficient when used within a single process or single job running on the same compute node.

3.3.12 The system shall include support for the languages C, C++ (including complete C++11/14/17), Fortran 77, Fortran 90, and Fortran 2008 programming languages. Providing multiple compilation environments is highly desirable. The Offeror shall describe any limitations that can be expected in meeting full C++17 support based on current expectations.

3.3.13 The system shall include a Python implementation that will run on the compute partition with optimized MPI4Py, NumPy, and SciPy libraries.

3.3.14 The system shall include a programming toolchain(s) that enables runtime coexistence of threading in C, C++, and Fortran, from within applications and any supporting libraries using the same toolchain. The Offeror shall describe the interaction between OpenMP and native parallelism expressed in language standards.

3.3.15 The system shall include C++ compiler(s) that can successfully build the Boost C++ library, http://www.boost.org. The Offeror shall support the most recent stable version of Boost.

3.3.16 The system shall include optimized versions of libm, libgsl, BLAS levels 1, 2 and 3, LAPACK, ScaLAPACK, HDF5, NetCDF, and FFTW. It is desirable for these to efficiently interoperate with applications that utilize OpenMP. The Offeror shall describe all other optimized libraries that will be supported, including a description of the interoperability of these libraries with the programming environments proposed.

3.3.17 The system shall include a mechanism that enables control of task and memory placement within a node for efficient performance. The Offeror shall provide a detailed description of controls provided and any limitations that may exist.

3.3.18 The system shall include a comprehensive software development environment with configuration and source code management tools. On heterogeneous systems, a mechanism (e.g., an upgraded autoconf) shall be provided to create configure scripts to build cross-compiled applications on login nodes.

3.3.19 The system shall include an interactive parallel debugger with an X11-based graphical user interface. The debugger shall provide a single point of control that can debug applications in all supported languages using all granularities of parallelism (e.g. MPI+X) and programming environments provided and scale up to 25% of the system.

3.3.20 The system shall include a suite of tools for detailed performance analysis and profiling of user applications. At least one tool shall support all granularities of parallelism in mixed MPI+OpenMP programs and any additional programming models supported on the system. The tool suite must provide the ability to support multi-node integrated profiling of on-node parallelism and communication performance analysis. The Offeror shall describe all proposed tools and the scalability limitations of each. The Offeror shall describe tools for measuring I/O behavior of user applications.

3.3.21 The system shall include event-tracing tools. Event tracing of interest includes: message-passing event tracing, I/O event tracing, floating point exception tracing, and message-passing profiling. The event-tracing tool API shall provide functions to activate and deactivate event monitoring during execution from within a process.

3.3.22 The system shall include single- and multi-node stack-tracing tools. The tool set shall include a source-level stack trace back, including an API that allows a running process or thread to query its current stack trace.

3.3.23 The system shall include tools to assist the programmer in introducing limited levels of parallelism and data structure refactoring to codes using any proposed programming models and languages. Tool(s) shall additionally be provided to assist application developers in the design and placement of the data structures with the goal of optimizing data movement/placement for the classes of memory proposed in the system.

3.3.24 The system shall include software licenses to enable the following number of simultaneous users on the system:

	Crossroads	NERSC-9
Compiler	20	100
Debugger	20	20

3.4   Platform Storage

Platform storage is certain to be one of the advanced technology areas included in any system delivered in this timeframe. The APEX laboratories anticipate these emerging technologies will enable new usage models. With this in mind, an accompanying whitepaper, "APEX Workflows," is provided that describes how application teams use HPC resources today to advance

scientific goals. The whitepaper is designed to provide a framework for reasoning about the optimal solution to these challenges. The whitepaper is intended to help an Offeror develop a platform storage architecture response that accelerates the science workflows while minimizing the total number of platform storage tiers. The Crossroads/NERSC-9 workflows document can be found on the APEX website.

3.4.1   The system shall include platform storage capable of retaining all application input, output, and working data for 12 weeks (84 days), estimated at a minimum of 36% of baseline system memory per day.

3.4.2   The system shall include platform storage with an appropriate durability or a maintenance plan such that the platform storage is capable of absorbing approximately four times the systems baseline memory per day for the life of the system.

3.4.3   The Offeror shall describe how the system provides sufficient bandwidth to support a JMTTI/Delta-Ckpt ratio of greater than 200 (where Delta-Ckpt is less than 7.2 minutes).

3.4.4   The Offeror shall describe the projected characteristics of all platform storage devices for the system, including but not limited to:

- Usable capacity, access latencies, platform storage interfaces (e.g. NVMe, PCIe), expected lifetime (warranty period, MTTF, total writes, etc.), and media and device error rates
- Relevant software/firmware features
- Compression technologies used by the platform storage devices, the resources used to implement the compression/decompression algorithms, the expected compression rates, and all compression/decompression-related performance impacts

3.4.5   The Offeror shall describe all available interfaces to platform storage for the system, including but not limited to:

- POSIX
- APIs
- Exceptions to POSIX compliance.
- Time to consistency and any potential delays for reliable data consumption.
- Any special requirements for users to achieve performance and/or consistent data.

3.4.6   The Offeror shall describe the reliability characteristics of platform storage, including but not limited to:

- Any single point of failure for all proposed platform storage tiers (note any component failure that will lead to temporary or permanent loss of data availability).
- Mean time to data loss for each platform storage tier provided.
- Enumerate platform storage tiers that are designed to be less reliable or do not use data protection techniques (e.g., replication, erasure coding).
- The magnitudes and duration of performance and reliability degradation brought about by a single or multiple component failures for each reliable platform storage tier.
- Vendor supplied mechanisms to ensure data integrity for each platform storage tier (e.g., data scrubbing processes, background checksum verification, etc.).
- Enumerate any platform storage failures that potentially impact scheduled or currently executing jobs that impact the platform storage or system performance and/or availability.
- Login or interactive nodes access to platform storage when the compute nodes are unavailable.

3.4.7   The Offeror shall describe system features for platform storage tier management designed to accelerate workflows, including but not limited to:

- Mechanisms for migrating data between platform storage tiers, including manual, scheduled, and/or automatic data migration to include rebalancing, draining, or rewriting data across devices within a tier.
- How platform storage will be instantiated with each job if it needs to be, and how platform storage may be persisted across jobs.
- The capabilities provided to define per-user policies and automate data movement between different tiers of platform storage or external storage resources (e.g., archives).
- The ability to serialize namespaces no longer in use (e.g., snapshots).
- The ability to restore namespaces needed for a scheduled job that is not currently available.
- The ability to integrate with or act as a site-wide scheduling resource.
- A mechanism to incrementally add capacity and bandwidth to a particular tier of platform storage without requiring a tier-wide outage.
- Capabilities to manage or interface platform storage with external storage resources or archives (e.g., fast storage layers or HPSS).

3.4.8   The Offeror shall describe software features that allow users to optimize I/O for the workflows of the system, including but not limited to:

- Batch data movement capabilities, especially when data resides on multiple tiers of platform storage.

- Methods for users to create and manage platform storage allocations.
- Any ability to directly write to or read from a tier not directly (logically) adjacent to the compute resources.
- Locality-aware job/data scheduling.
- I/O utilization for reservations.
- Features to prevent data duplication on more than one platform storage tier.
- Methods for users to exploit any enhanced performance of relaxed consistency.
- Methods for enabling user-defined metadata with the platform storage solution.

3.4.9   The Offeror shall describe the method for walking the entire platform storage metadata, and describe any special capabilities that would mitigate user performance issues for daily full-system namespace walks; expect at least 1 billion objects.

3.4.10 The Offeror shall describe any capabilities to comprehensively collect platform storage usage data (in a scalable way), for the system, including but not limited to:

- Per client metrics and frequency of collection, including but not limited to: the number of bytes read or written, number of read or write invocations, client cache statistics, and metadata statistics such as number of opens, closes, creates, and other system calls of relevance to the performance of platform storage.
- Job level metrics, such as the number of sessions each job initiates with each platform storage tier, session duration, total data transmitted (separated as reads and writes) during the session, and the number of total platform storage invocations made during the session.
- Platform storage tier metrics and frequency of collection, such as the number of bytes read, number of bytes written, number of read invocations, number of write invocations, bytes deleted/purged, number of I/O sessions established, and periods of outage/unavailability.
- Job level metrics describing usage of a tiered platform storage hierarchy, such as how long files are resident in each tier, hit rate of file pages in each tier (i.e., whether pages are actually read and how many times data is re-read), fraction of data moved between tiers because of a) explicit programmer control and b) transparent caching, and time interval between accesses to the same file (e.g., how long until an analysis program reads a simulation generated output file).

3.4.11 The Offeror shall propose a method for providing access to platform storage from other systems at the facility. In the case of tiered platform storage, at least one tier must satisfy this requirement.

3.4.12 The Offeror shall describe the capability for platform storage tiers to be repaired, serviced, and incrementally patched/upgraded while running different versions of software or firmware without requiring a storage tier-wide outage. The Offeror shall describe the level of performance degradation, if any, anticipated during the repair or service interval.

3.4.13 The Offerer shall specify the minimum number of compute nodes required to read and write the following data sets from/to platform storage:

- A 1 TB data set of 20 GB files in 2 seconds.
- A 5 TB data set of any chosen file size in 10 seconds. Offeror shall report the file size chosen.
- A 1 PB data set of 32 MB files in 1 hour.

3.5   Application Performance

Assuring that real applications perform well on both the Crossroads and NERSC-9 systems is key for their success. Because the full applications are large, often with millions of lines of code, and in some cases are export controlled, a suite of benchmarks have been developed for RFP response evaluation and system acceptance. The benchmark codes are representative of the workloads of the APEX laboratories but often smaller than the full applications.

The performance of the benchmarks will be evaluated as part of both the RFP response and system acceptance. Final benchmark acceptance performance targets will be negotiated after a final system configuration is defined. All performance tests must continue to meet negotiated acceptance criteria throughout the lifetime of the system.

System acceptance for Crossroads shall also include an ASC Simulation Code Suite comprised of at least two (2) but no more than four (4) ASC applications from the three NNSA laboratories, Sandia, Los Alamos and Lawrence Livermore.

The Crossroads/NERSC-9 benchmarks, information regarding the Crossroads acceptance codes, and supplemental materials can be found on the APEX website.

3.5.1 The Offeror shall provide responses to the benchmarks (SNAP, PENNANT, HPCG, MiniPIC, UMT, MILC, MiniDFT, GTC, and Meraculous) provided on the *Crossroads/NERSC-9 benchmarks* link on the APEX website. All modifications or new variants of the benchmarks (including makefiles, build scripts, and environment variables) are to be supplied in the Offeror's response.

- The results of all problem sizes (baseline and optimized) shall be provided in the Offeror's Scalable System Improvement (SSI) spreadsheets. SSI is the calculation used for measuring improvement and is documented on the APEX website, along with the SSI spreadsheets. If predicted or extrapolated results are provided, the methodology used to derive them should be documented.
- The Offeror shall provide licenses for the system for all compilers, libraries, and runtimes used to achieve benchmark performance.

3.5.2 The Offeror shall provide performance results for the system that may be benchmarked, predicted, and/or extrapolated for the baseline MPI+OpenMP (or UPC for Meraculous) variants of the benchmarks. The Offeror may modify the benchmarks to include extra OpenMP pragmas as required, but the benchmark must remain a standard-compliant program that maintains existing output subject to the validation criteria described in the benchmark run rules.

3.5.3 The Offeror shall optionally provide performance results from an Offeror optimized variant of the benchmarks. The Offeror may modify the benchmarks, including the algorithm and/or programming model used to demonstrate high system performance. If algorithmic changes are made, the Offeror shall provide an explanation of why the results may deviate from validation criteria described in the benchmark run rules.

3.5.4 For the Crossroads system only: in addition to the *Crossroads/NERSC-9 benchmarks*, an ASC Simulation Code Suite representing the three NNSA laboratories will be used to judge performance at time of acceptance. The Crossroads system shall achieve a minimum of at least 6 times (6X) improvement over the ASC Trinity system (Knights Landing partition) for each code, measured using SSI. The Offeror shall specify a baseline performance greater than or equal to 6X at time of response. Final acceptance performance targets will be negotiated after a final system configuration is defined. Information regarding ASC Simulation Code Suite run rules and acceptance can be found on the APEX website. Source code will be provided to the Offeror but will require compliance with export control laws and no cost licensing agreements.

3.5.5   The Offeror shall report or project the number of cores necessary to saturate the available node baseline memory bandwidth as measured by the Crossroads/NERSC-9 memory bandwidth benchmark found on the APEX website.

- If the node contains heterogeneous cores, the Offeror shall report the number of cores of each architecture necessary to saturate the available baseline memory bandwidth.
- If multiple tiers of memory are available, the Offeror shall report the above for every functional combination of core architecture and baseline or extended memory tier.

3.5.6   The Offeror shall report or project the sustained dense matrix multiplication performance on each type of processor core (individually and/or in parallel) of the system node architecture(s) as measured by the Crossroads/NERSC-9 multithreaded DGEMM benchmark found on the APEX website.

- The Offeror shall describe the percentage of theoretical double-precision (64-bit) computational peak, which the benchmark GFLOP/s rate achieves for each type of compute core/unit in the response, and describe how this is calculated.

3.5.7   The Offeror shall report, or project, the MPI two-sided message rate of the nodes in the system under the following conditions measured by the communication benchmark specified on the APEX website:

- Using a single MPI rank per node with MPI_THREAD_SINGLE.
- Using two, four, and eight MPI ranks per node with MPI_THREAD_SINGLE.
- Using one, two, four, and eight MPI ranks per node and multiple threads per rank with MPI_THREAD_MULTIPLE.
- The Offeror may additionally choose to report on other configurations.

3.5.8   The Offeror shall report, or project, the MPI one-sided message rate of the nodes in the system for all passive synchronization RMA methods with both pre-allocated and dynamic memory windows under the following conditions measured by the communication benchmark specified on the APEX website using:

- A single MPI rank per node with MPI_THREAD_SINGLE.
- Two, four, and eight MPI ranks per node with MPI_THREAD_SINGLE.
- One, two, four, and eight MPI ranks per node and multiple threads per rank with MPI_THREAD_MULTIPLE.
- The Offeror may additionally choose to report on other configurations.

3.5.9   The Offeror shall report, or project, the time to perform the following collective operations for full, half, and quarter machine size in the system and report on core occupancy during the operations measured by the communication benchmark specified on the APEX website for:

- An 8 byte MPI_Allreduce operation.
- An 8 byte per rank MPI_Allgather operation.

3.5.10 The Offeror shall report, or project, the minimum and maximum off-node latency of the system for MPI two-sided messages using the following threading modes measured by the communication benchmark specified on the APEX website:

- MPI_THREAD_SINGLE with a single thread per rank.
- MPI_THREAD_MULTIPLE with two or more threads per rank.

3.5.11 The Offeror shall report, or project, the minimum and maximum off-node latency for MPI one-sided messages of the system for all passive synchronization RMA methods with both pre-allocated and dynamic memory windows using the following threading modes measured by the communication benchmark specified on the APEX website:

- MPI_THREAD_SINGLE with a single thread per rank.
- MPI_THREAD_MULTIPLE with two or more threads per rank.

3.5.12 The Offeror shall provide an efficient implementation of MPI_THREAD_MULTIPLE. Bandwidth, latency, and message throughput measurements using the MPI_THREAD_MULTIPLE thread support level shall have no more than a 10% performance degradation when compared to using the MPI_THREAD_SINGLE support level as measured by the communication benchmark specified on the APEX website.

3.5.13 The Offeror shall report, or project, the maximum I/O bandwidths of the system as measured by the IOR benchmark specified on the APEX website.

3.5.14 The Offeror shall report, or project, the metadata rates of the system as measured by the MDTEST benchmark specified on the APEX website.

3.5.15 The Offeror shall be required at time of acceptance to meet specified targets for acceptance benchmarks, and mission codes for Crossroads, listed on the APEX website.

3.5.16 The Offeror shall describe how the system may be configured to support a high rate and bandwidth of TCP/IP connections to external services both from compute nodes and directly to and from the platform storage, including:

- Compute node external access shall allow all nodes to each initiate 1 connection concurrently within a 1 second window.

- Transfer of data over the external network to and from the compute nodes and platform storage at 100 GB/s per direction of a 1 TB dataset comprised of 20 GB files in 10 seconds.

3.6   Resilience, Reliability, and Availability

The ability to achieve the APEX mission goals hinges on the productivity of system users. System availability is therefore essential and requires system-wide focus to achieve a resilient, reliable, and available system. For each metric specified below, the Offeror must describe how they arrived at their estimates.

3.6.1   Failure of the system management and/or RAS system(s) shall not cause a system or job interrupt. This requirement does not apply to a RAS system feature, which automatically shuts down the system for safety reasons, such as an overheating condition.

3.6.2   The minimum System Mean Time Between Interrupt (SMTBI) shall be greater than 720 hours.

3.6.3   The minimum Job Mean Time To Interrupt (JMTTI) shall be greater than 24 hours. Automatic restarts do not mitigate a job interrupt for this metric.

3.6.4   The ratio of JMTTI/Delta-Ckpt shall be greater than 200. This metric is a measure of the system's ability to make progress over a long period of time and corresponds to an efficiency of approximately 90%. If, for example, the JMTTI requirement is not met, the target JMTTI/Delta-Ckpt ratio ensures this minimum level of efficiency.

3.6.5   An immediate re-launch of an interrupted job shall not require a complete resource reallocation. If a job is interrupted, there shall be a mechanism that allows re-launch of the application using the same allocation of resource (e.g., compute nodes) that it had before the interrupt or an augmented allocation when part of the original allocation experiences a hard failure.

3.6.6   A complete system initialization shall take no more than 30 minutes. The Offeror shall describe the full system initialization sequence and timings.

3.6.7   The system shall achieve 99% scheduled system availability. System availability is defined in the glossary.

3.6.8   The Offeror shall describe the resilience, reliability, and availability mechanisms and capabilities of the system including, but not limited to:

- Any condition or event that can potentially cause a job interrupt.
- Resiliency features to achieve the availability targets.

- Single points of failure (hardware or software), and the potential effect on running applications and system availability.
- How a job maintains its resource allocation and is able to relaunch an application after an interrupt.
- A system-level mechanism to collect failure data for each kind of component.

3.7 Application Transition Support and Early Access to APEX Technologies

The Crossroads and NERSC-9 systems will include numerous pre-exascale technologies. The Offeror shall include in their proposal a plan to effectively utilize these technologies and assist in transitioning the mission workflows to the systems. For the Crossroads system only, the Offeror shall support efforts to transition the Advanced Technology Development Mitigation (ATDM) codes to the systems. ATDM codes are currently being developed by the three NNSA weapons laboratories, Sandia, Los Alamos, and Lawrence Livermore. These codes may require compliance with export control laws and no cost licensing agreements. Information about the ATDM program can be found on the NNSA website.

3.7.1   The Offeror shall provide (thus the Offeror shall propose) a vehicle for supporting the successful demonstration of the application performance requirements and the transition of key applications to the Crossroads and NERSC-9 systems (e.g., a Center of Excellence). Support shall be provided by the Offeror and all of its key advanced technology providers (e.g., processor vendors, integrators, etc). The Offeror shall provide experts in the areas of application porting and performance optimization in the form of staff training, general user training, and deep-dive interactions with a set of application code teams. Support shall include compilers to enable timely bug fixes as well as to enable new functionality. Support shall be provided from the date of subcontract execution through two (2) years after final acceptance of the systems.

3.7.2   The Offeror shall describe which of the proposed APEX hardware and software technologies (physical hardware, emulators, and/or simulators) , will be available for access before system delivery and in what timeframe. The proposed technologies should provide value in advanced preparation for the delivery of the final APEX system(s) for pre-system-delivery application porting and performance assessment activities.

Dated 05-23-16

3.8   Target System Configuration

Table 2 Target System Configuration

	Crossroads	**NERSC-9**
Baseline Memory Capacity *Excludes all levels of on-die-CPU cache*	**> 3 PiB**	**> 3 PiB**
Benchmark SSI increase over Edison system	**> 20X**	**> 20X**
Platform Storage	**> 30X Baseline Memory**	**> 30X Baseline Memory**
Wall Plate Power	**< 20 MW**	**< 20 MW**
Peak Power	**< 18 MW**	**< 18 MW**
Nominal Power	**< 15 MW**	**< 15 MW**
Idle Power	**< 10% Wall Plate Power**	**< 10% Wall Plate Power**
Job Mean Time To Interrupt (JMTTI) *Calculated for a single job running in the entire system*	**> 24 Hours**	**> 24 Hours**
System Mean Time To Interrupt (SMTTI)	**>720 Hours**	**> 720 Hours**
Delta-Ckpt	**< 7.2 minutes**	**< 7.2 minutes**
JMTTI/Delta-Ckpt	**> 200**	**> 200**
System Availability	**> 99%**	**> 99%**

3.9   System Operations

System management shall be an integral feature of the overall system and shall provide the ability to effectively manage system resources with high utilization and throughput under a workload with a wide range of concurrencies. The Offeror shall provide system administrators, security

officers, and user-support personnel with productive and efficient system configuration management capabilities and an enhanced diagnostic environment.

3.9.1   The system shall include scalable integrated system management capabilities that provide human interfaces and APIs for system configuration and its ability to be automated, software management, change management, local site integration, and system configuration backup and recovery.

3.9.2   The system shall include a means for tracking and analyzing all software updates, software and hardware failures, and hardware replacements over the lifetime of the system.

3.9.3   The system shall include the ability to perform rolling upgrades and rollbacks on a subset of the system while the balance of the system remains in production operation. The Offeror shall describe the mechanisms, capabilities, and limitations of rolling upgrades and rollbacks. No more than half the system partition shall be required to be down for rolling upgrades and rollbacks.

3.9.4   The system shall include an efficient mechanism for reconfiguring and rebooting compute nodes. The Offeror shall describe in detail the compute node reboot mechanism, differentiating types of boots (warmboot vs. coldboot) required for different node features, as well as how the time required to reboot scales with the number of nodes being rebooted.

3.9.5   The system will include a mechanism whereby all monitoring data and logs captured are available to the system owner, and will support an open monitoring API to facilitate lossless, scalable sampling and data collection for monitored data. Any filtering that may need to occur will be at the option of the system manager. The system will include a sampling and connection framework that allows the system manager to configure independent alternative parallel data streams to be directed off the system to site-configurable consumers.

3.9.6   The system shall include a mechanism to collect and provide metrics and logs which monitor the status, health, and performance of the system, including, but not limited to:

- Environmental measurement capabilities for all systems and peripherals and their sub-systems and supporting infrastructure, including power and energy consumption and control.
- Internal HSN performance counters, including measures of network congestion and network resource consumption.
- All levels of integrated and attached platform storage.

- The system as a whole, including hardware performance counters for metrics for all levels of integrated and attached platform storage.

3.9.7   The Offeror shall describe what tools it shall provide for the collection, analysis, integration, and visualization of metrics and logs produced by the system (e.g., peripherals, integrated and attached platform storage, and environmental data, including power and energy consumption).

3.9.8   The Offeror shall describe the system configuration management and diagnostic capabilities of the system that address the following topics:

- Detailed description of the system management support.
- Any effect or overhead of software management tool components on the CPU or memory available on compute nodes.
- Release plan, with regression testing and validation for all system related software and security updates.
- Support for multiple simultaneous or alternative system software configurations, including estimated time and effort required to install both a major and a minor system software update.
- User activity tracking, such as audit logging and process accounting.
- Unrestricted privileged access to all hardware components delivered with the system.

3.10 Power and Energy

Power, energy, and temperature will be critical factors in how the APEX laboratories manage systems in this time frame and must be an integral part of overall Systems Operations. The solution must be well integrated into other intersecting areas (e.g., facilities, resource management, runtime systems, and applications). The APEX laboratories expect a growing number of use cases in this area that will require a vertically integrated solution.

3.10.1 The Offeror shall describe all power, energy, and temperature measurement capabilities (system, rack/cabinet, board, node, component, and sub-component level) for the system, including control and response times, sampling frequency, accuracy of the data, and timestamps of the data for individual points of measurement and control.

3.10.2 The Offeror shall describe all control capabilities it shall provide to affect power or energy use (system, rack/cabinet, board, node, component, and sub-component level).

3.10.3 The system shall include system-level interfaces that enable measurement and dynamic control of power and energy relevant characteristics of the system, including but not limited to:

- AC measurement capabilities at the system or rack level.
- System-level minimum and maximum power settings (e.g., power caps).
- System-level power ramp up and down rate.
- Scalable collection and retention all measurement data such as:
- point-in-time power data.
- energy usage information.
- minimum and maximum power data.

3.10.4 The system shall include resource manager interfaces that enable measurement and dynamic control of power and energy relevant characteristics of the system, including but not limited to:

- Job and node level minimum and maximum power settings.
- Job and node level power ramp up and down rate.
- Job and node level processor and/or core frequency control.
- System and job level profiling and forecasting.
  - o e.g., prediction of hourly power averages >24 hours in advance with a 1 MW tolerance.

3.10.5 The system shall include application and runtime system interfaces that enable measurement and dynamic control of power and energy relevant characteristics of the system including but not limited to:

- Node level minimum and maximum power settings.
- Node level processor and/or core frequency control.
- Node level application hints, such as:
  - o application entering serial, parallel, computationally intense, I/O intense or communication intense phase.

3.10.6 The system shall include an integrated API for all levels of measurement and control of power relevant characteristics of the system. It is preferable that the provided API complies with the High Performance Computing Power Application Programming Interface Specification (http://powerapi.sandia.gov).

3.10.7 The Offeror shall project (and report) the Wall Plate, Peak, Nominal, and Idle Power of the system.

3.10.8 The Offeror shall describe any controls available to enforce or limit power usage below wall plate power and the reaction time of this mechanism (e.g., what duration and magnitude can power usage exceed the imposed limits).

3.10.9 The Offeror shall describe the status of the system when in an Idle State (describe all Idle States if multiple are available) and the time to transition from the Idle State (or each Idle State if there are multiple) to the start of job execution.

3.11 Facilities and Site Integration

3.11.1 The system shall use 3-phase 480V AC. Other system infrastructure components (e.g., disks, switches, login nodes, and mechanical subsystems such as CDUs) must use either 3-phase 480V AC (strongly preferred), 3-phase 208V AC (second choice), or single-phase 120/240V AC (third choice). The total number of individual branch circuits and phase load imbalance shall be minimized.

3.11.2 All equipment and power control hardware of the system shall be Nationally Recognized Testing Laboratories (NRTL) certified and bear appropriate NRTL labels.

3.11.3 Every rack, network switch, interconnect switch, node, and disk enclosure shall be clearly labeled with a unique identifier visible from the front of the rack and/or the rear of the rack, as appropriate, when the rack door is open. These labels will be high quality so that they do not fall off, fade, disintegrate, or otherwise become unusable or unreadable during the lifetime of the system. Nodes will be labeled from the rear with a unique serial number for inventory tracking. It is desirable that motherboards also have a unique serial number for inventory tracking. Serial numbers shall be visible without having to disassemble the node, or they must be able to be queried from the system management console.

3.11.4 The Offeror shall describe the features of the system related to facilities and site integration, including:

- Description of the physical packaging of the system, including dimensioned drawings of individual cabinets types and the floor layout of the entire system.
- Remote environmental monitoring capabilities of the system and how it would integrate into facility monitoring.
- Emergency shutdown capabilities.
- Detailed descriptions of power and cooling distributions throughout the system, including power consumption for all subsystems.

- Description of parasitic power losses within Offeror's equipment, such as fans, power supply conversion losses, power-factor effects, etc. For the computational and platform storage subsystems separately, give an estimate of the total power and parasitic power losses (whose difference should be power used by computational or platform storage components) at the minimum and maximum ITUE, which is defined as the ratio of total equipment power over power used by computational or platform storage components. Describe the conditions (e.g. "idle") at which the extrema occur.
- OS distributions or other client requirements to support off-system access to the platform storage (e.g. LANL File Transfer Agents).

Table 3 Crossroads and NERSC-9 Facility Requirements

	Crossroads	**NERSC-9**
Location	Los Alamos National Laboratory, Los Alamos, NM. The system will be housed in the Strategic Computing Complex (SCC), Building 2327	National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory, Berkeley, CA. The system will be housed in Wang Hall, Building 59 (formerly known as the Computational Theory and Research Facility).
Altitude	7,500 feet	650 feet
Seismic	N/A	System to be placed on a seismic isolation floor. System cabinets shall have an attachment mechanism that will enable them to be firmly attached to each other and the isolation floor. When secured via these attachments, the cabinets shall withstand seismic design accelerations per the California Building Code and LBNL Lateral Force Design Criteria policy in

	Crossroads	NERSC-9
		effect at the time of subcontract award. (The CBC currently specifies 0.49g but is expected to be updated in 2016.)
Water Cooling	The system shall operate in conformance with ASHRAE Class W2 guidelines (dated 2011). The facility will provide operating water temperature that nominally varies between 60-75°F, at up to 35PSI differential pressure at the system cabinets However, Offeror should note if the system is capable of operating at higher temperatures. Note: LANL facility will provide inlet water at a nominal 75°F. It may go to as low as 60°F based on facility and/or environmental factors. Total flow requirements may not exceed 9600GPM.	Same Note: NERSC facility will provide inlet water at a nominal 65°F. It may go as high as 75°F based on facility and/or environmental factors. Total flow requirements may not exceed 9600GPM.
Water Chemistry	The system must operate with facility water meeting basic ASHRAE water chemistry. Special chemistry water is not available in the main building loop and would require a separate tertiary loop provided with the system. If	Same

	Crossroads	NERSC-9
	tertiary loops are included in the system, the Offeror shall describe their operation and maintenance, including coolant chemistry, pressures, and flow controls. All coolant loops within the system shall have reliable leak detection, temperature, and flow alarms, with automatic protection and notification mechanisms.	
Air Cooling	The system must operate with supply air at 76°F or below, with a relative humidity from 30%-70%. The rate of airflow is between 800-1500 CFM/floor tile. No more than 3MW of heat shall be removed by air cooling.	The system must operate with supply air at 76°F or below, with a relative humidity from 30%-80%. The current facility can support up to 60K CFM of airflow, and remove 500KW of heat. Expansion is possible to 300K CFM and 1.5MW, but at added expense.
Maximum Power Rate of Change	The hourly average in system power shall not exceed the 2MW wide power band negotiated at least 2 hours in advance.	N/A
Power Quality	The system shall be resilient to incoming power fluctuations at least to the level guaranteed by the ITIC power quality curve.	Same
Floor	42" raised floor	48" raised floor
Ceiling	16 foot ceiling and an 18'	17'10" ceiling however maximum cabinet height

	Crossroads	NERSC-9
	6" ceiling plenum	is 9'5"
Maximum Footprint	8000 square feet; 80 feet long and 100 feet deep.	64'x92', or 5888 square feet (inclusive of compute, platform storage and service aisles). This area is itself surrounded by a minimum 4' aisle that can be used in the system layout. It is preferred that cabinet rows run parallel to the short dimension.
Shipment Dimensions and Weight	No restrictions.	For delivery, system components shall weigh less than 7000 pounds and shall fit into an elevator whose door is 6ft 6in wide and 9ft 0 in high and whose depth is 8ft 3in. Clear internal width is 8ft 4 in.
Floor Loading	The average floor loading over the effective area shall be no more than 300 pounds per square foot. The effective area is the actual loading area plus at most a foot of surrounding fully unloaded area. A maximum limit of 300 pounds per square foot also applies to all loads during installation. The Offeror shall describe how the weight will be distributed over the footprint of the rack (point loads, line loads, or evenly distributed over the entire footprint). A	The floor loading shall not exceed a uniform load of 500 pounds per square foot. Raised floor tiles are ASM FS400 with an isolated point load of 2000 pounds and a rolling load of 1200 pounds.

	Crossroads	NERSC-9
	point load applied on a one square inch area shall not exceed 1500 pounds. A dynamic load using a CISCA Wheel 1 size shall not exceed 1250 pounds (CISCA Wheel 2 – 1000 pounds).	
Cabling	All power cabling and water connections shall be below the access floor. It is preferable that all other cabling (e.g., system interconnect) is above floor and integrated into the system cabinetry. Under floor cables (if unavoidable) shall be plenum rated and comply with NEC 300.22 and NEC 645.5. All communications cables, wherever installed, shall be source/destination labeled at both ends. All communications cables and fibers over 10 meters in length and installed under the floor shall also have a unique serial number and dB loss data document (or equivalent) delivered at time of installation for each cable, if a method of measurement exists for cable type.	Same
External network interfaces supported by the site for connectivity requirements specified	1Gb, 10Gb, 40Gb, 100Gb, IB	Same

	Crossroads	**NERSC-9**
below		
External bandwidth on/off the system for general TCP/IP connectivity	> 100 GB/s per direction	Same
External bandwidth on/off the system for accessing the system's PFS	> 100 GB/s	Same
External bandwidth on/off the system for accessing external, site supplied file systems. E.g. GPFS, NFS	> 100 GB/s	Same

4   Non-Recurring Engineering

The APEX team expects to award two (2) Non-Recurring Engineering (NRE) subcontracts, separate from the two (2) system subcontracts. It is expected that Crossroads and NERSC personnel will collaborate in both NRE subcontracts. It is anticipated that the NRE subcontracts will be approximately 10%-15% of the combined Crossroads and NERSC-9 system budgets. The Offeror is encouraged to provide proposals for areas of collaboration they feel provide substantial value to the Crossroads and NERSC-9 systems with the goals of:

▪ Increasing application performance.
▪ Increasing workflow performance.
▪ Increasing the resilience, and reliability of the system.

Proposed collaboration areas should focus on topics that provide added value beyond planned roadmap activities. Proposals should not focus on one-off point solutions or gaps created by their proposed design that should be otherwise provided as part of a vertically integrated solution. It is expected that NRE collaborations will have impact on both the Crossroads and NERSC-9 systems and follow-on systems procured by the U.S. Department of Energy's NNSA and Office of Science.

NRE topics of interest include, but are not limited to, the following:

▪ Development and optimization of hardware and software capabilities to increase the performance of MPI+OpenMP and future task-based asychronous programming models.

- Development and optimization of hardware and software capabilities to increase the performance of application workflows, including consideration of consistency requirements, data-migration needs, and system-wide resource management.
- Development of scalable system management capabilities to enhance the reliability, resilience, power, and energy usage of Crossroads/NERSC-9.

5   Options

The APEX team expects to have future requirements for system upgrades and/or additional quantities of components based on the configurations proposed in response to this solicitation. The Offeror shall address any technical challenges foreseen with respect to scaling and any other production issues. Proposals should be as detailed as possible. The evaluation committee will make no presumption of technical capability when evaluating Offeror responses to Options. Offerors that do not address the Options in a materially responsive manner will be downgraded.

5.1   Upgrades, Expansions and Additions

5.1.1   The Offeror shall propose and separately price upgrades, expansions or procurement of additional system configurations by the following fractions of the system as measured by the Sustained System Improvement (SSI) metric.

- 25%
- 50%
- 100%
- 200%

5.1.2   The Offeror shall propose a configuration or configurations which double the baseline memory capacity.

5.1.3   The Offeror shall propose upgrades, expansions or procurement of additional platform storage capacity (per tier if multiple tiers are present) in increments of 25%.

5.2   Early Access Development System

To allow for early and/or accelerated development of applications or development of functionality required as a part of the statement of work, the Offeror shall propose options for early access development systems. These systems can be in support of the baseline requirements or any proposed options.

5.2.1   The Offeror shall propose an Early Access Development System. The primary purpose is to expose the application to the same programming environment as will be found on the final system. It is acceptable for the early access system to not use the final processor, node, or high-speed interconnect architectures. However, the programming and runtime environment must be sufficiently similar that a port to the final system is trivial. The early access system shall contain similar functionality of the final system, including file systems, but scaled down to the appropriate configuration. The Offeror shall propose an option for the following configurations based on the size of the final Crossroads/NERSC-9 systems.

- 2% of the compute partition.
- 5% of the compute partition.
- 10% of the compute partition.

5.2.2   The Offeror shall propose development test bed systems that will reduce risk and aid the development of any advanced functionality that is exercised as a part of the statement of work. For example, any topics proposed for NRE.

5.3   Test Systems

The Offeror shall propose the following test systems. The systems shall contain all the functionality of the main system, including file systems, but scaled down to the appropriate configuration. Multiple test systems may be awarded.

5.3.1   The Offeror shall propose an Application Regression test system, which shall contain at least 200 compute nodes.

5.3.2   The Offeror shall propose a System Development test system, which shall contain at least 50 compute nodes.

5.4   On Site System and Application Software Analysts

5.4.1   The Offeror shall propose and separately price two (2) System Software Analysts and two (2) Applications Software Analysts for each site. Offerors shall presume each analyst will be utilized for four (4) years. For Crossroads, these positions require a DOE Q-clearance for access.

5.5   Deinstallation

The Offeror shall deinstall, remove and/or recycle the system and supporting infrastructure at end of life. Storage media shall be wiped or destroyed to the satisfaction of ACES and NERSC, and/or returned to ACES and NERSC at their request.

5.6 Maintenance and Support

The Offeror shall propose and separately price maintenance and support with the following features:

5.6.1 Maintenance and Support Period

The Offeror shall propose all maintenance and support for a period of four (4) years from the date of acceptance of the system. Warranty shall be included in the 4 years. For example, if the system is accepted on April 1, 2021 and the Warranty is for one year, then the Warranty ends on March 30, 2022, and the maintenance period begins April 1, 2022 and ends on March 30, 2025. Offeror shall also propose additional maintenance and support extension for years 5-7.

5.6.2 Maintenance and Support Solutions

The Offeror shall propose the following maintenance and support solutions and propose pricing separately for each solution. ACES and NERSC may purchase either one of the solutions or neither of the solutions, at its discretion. Different maintenance solutions may be selected for the various test systems and final system.

5.6.2.1 Solution 1 – 7x24

The Offeror shall price Solution 1 as full hardware and software support for all Offeror provided hardware components and software. The principal period of maintenance (PPM) shall be for 24 hours by 7 days a week with a four hour response to any request for service.

5.6.2.2 Solution 2 – 5x9

The Offeror shall price Solution 2 as full hardware and software support for all Offeror provided hardware components and software. The principal period of maintenance (PPM) shall be on a 9 hours by 5 days a week (exclusive of holidays observed by ACES or NERSC). The Offeror shall provide hardware maintenance training for ACES/NERSC staff so that staff are able to provide hardware support for all other times the Offeror is unable to provide hardware repair in a timely manner outside of the PPM. The Offeror shall supply hardware maintenance procedural documentation, training, and manuals necessary to support this effort.

All proposed maintenance and support solutions shall include the following features and meet all requirements of this section.

5.6.3 General Service Provisions

The Offeror shall be responsible for repair or replacement of any failing hardware component that it supplies and correction of defects in software that it provides as part of the system.

At its sole discretion, ACES or NERSC may request advance replacement of components which show a pattern of failures which reasonably indicates that future failures may occur in excess of reliability targets, or for which there is a systemic problem that prevents effective use of the system.

Hardware failures due to environmental changes in facility power and cooling systems which can be reasonably anticipated (such as brown-outs, voltage-spikes or cooling system failures) are the responsibility of the Offeror.

5.6.4   Software and Firmware Update Service

The Offeror shall provide an update service for all software and firmware provided for the duration of the Warranty plus Maintenance period. This shall include new releases of software/firmware and software/firmware patches as required for for normal use. The Offeror shall integrate software fixes, revisions or upgraded versions in supplied software, including community software (e.g. Linux or Lustre), and make them available to ACES and NERSC within twelve (12) months of their general availability. The Offeror shall provide prompt availability of patches for cybersecurity defects.

5.6.5   Call Service

The Offeror shall provide contact information for technical personnel with knowledge of the proposed equipment and software. These personnel shall be available for consultation by telephone and electronic mail with ACES/NERSC personnel. In the case of degraded performance, the Offeror's services shall be made readily available to develop strategies for improving performance, i.e. patches, workarounds.

5.6.6   On-site Parts Cache

The Offeror shall maintain a parts cache on-site at both the ACES and NERSC facilities. The parts cache shall be sized and provisioned sufficiently to support all normal repair actions for two weeks without the need for parts refresh. The initial sizing and provisioning of the cache shall be based on Offeror's Mean Time Between Failure (MTBF) estimates for each FRU and each rack, and scaled based on the number of FRU's and racks delivered. The parts cache configuration will be periodically reviewed for quantities needed to satisfy this requirement, and adjusted if necessary, based on observed FRU or node failure rates. The parts cache will be resized, at the Offeror's expense, should the on-site parts cache prove to be insufficient to sustain the actually observed FRU or node failure rates.

5.6.7   On-Site Node Cache

The Offeror shall also maintain an on-site spare node inventory of at least 1% of the total nodes in all of the system.   These nodes shall be maintained and tested for hardware integrity and functionality utilizing the Hardware Support Cluster defined below if provided.

The following features and requirements are specific to responses for ACES requirements.

5.6.8   Hardware Support Cluster

The Offeror shall provide a Hardware Support Cluster (HSC). The HSC shall support the hot spare nodes and provide functions such as hardware burn-in, problem diagnosis, etc. The Offeror shall supply sufficient racks, interconnect, networking, storage equipment and any associated hardware/software necessary to make the HSC a stand-alone system capable of running diagnostics on individual or clusters of HSC nodes. ACES will store and inventory the HSC and other on-site parts cache components.

5.6.9   DOE Q-Cleared Technical Service Personnel

The Crossroads system will be installed in security areas that require a DOE Q-clearance for access. It will be possible to install the system with the assistance of uncleared US citizens or L-cleared personnel, but the Offeror shall arrange and pay for appropriate 3rd party security escorts. The Offeror shall obtain necessary clearances for on-site support staff to perform their duties.

6   Delivery and Acceptance

Testing of the system shall proceed in three steps: pre-delivery, post-delivery, and acceptance. Each step is intended to validate the system and feeds into subsequent activities. Sample Acceptance Test plans (Appendix A) shall be provided as part of the Request for Proposal.

6.1 Pre-delivery Testing

The APEX team and the Offeror shall perform pre-delivery testing at the factory on the hardware to be delivered. Any limitations for performing the pre-delivery testing shall be identified in the Offeror's proposal, including scale and licensing limitations (if any). During pre-delivery testing, the Offeror shall:

- Demonstrate RAS capabilities and robustness using simple fault injection techniques, such as disconnecting cables, powering down subsystems, or installing known bad parts.

- Demonstrate functional capabilities on each segment of the system built, including the capacity to build applications, schedule jobs, and run them using a customer-provided testing framework. The root cause of application failure must be identified prior to system shipping.
- Provide a file system sufficiently provisioned to support the suite of tests.
- Provide onsite and remote access to the APEX team to monitor testing and analyze results.
- Instill confidence in the ability to conform to the statement of work.

6.2 Site Integration and Post-delivery Testing

The APEX team and the Offeror staff shall perform site integration and post-delivery testing on the fully delivered system. Limitations and/or special requirements may exist for access to the onsite system by the Offeror.

- During post-delivery testing, the pre-delivery tests shall be run on the full system installation.
- Where applicable, tests shall be run at full scale.

6.3 Acceptance Testing

The APEX team and the Offeror staff shall perform onsite acceptance testing on the fully installed system. Limitations and/or special requirements may exist for access to the onsite system by the Offeror.

6.3.1   The Offeror shall demonstrate that the delivered system conforms to the subcontract's Statement of Work.

7   Risk and Project Management

The Offeror shall propose a risk management strategy and project management plan for the Crossroads and NERSC-9 systems that is closely coordinated between the subcontracts for LANS and UC.

7.1.1   The Offeror shall Propose a risk management strategy for the system in the event of technology problems or scheduling delays that affect delivery of the system or achievement of performance targets in the proposed timeframe. Offeror shall describe the impact of substitute technologies (if any) on the overall architecture and performance of the system in particular addressing the four technology areas listed below:

- Processor
- Memory
- High-speed interconnect
- Platform storage

7.1.2   The Offeror shall identify any other high-risk areas and accompanying mitigation strategies for the system.

7.1.3   The Offeror shall provide a clear plan for effectively responding to software and hardware defects and system outages at each severity level and document how problems or defects will be escalated.

7.1.4   The Offeror shall propose a roadmap showing how their response to this Request for Proposal aligns with their plans for exascale computing.

7.1.5   The Offeror shall identify additional capabilities, including:

- Its ability to produce and maintain the system for the life of the system
- Its ability to achieve specific quality assurance, reliability, availability and serviceability goals
- Its in-house testing and problem diagnosis capability, including hardware resources at appropriate scale

7.1.6   The Offeror shall provide project management specifics for the APEX team shall be detailed as part of the Request for Proposal document. Please see Appendix B for further information.

8   Documentation and Training

The Offeror shall provide documentation and training to effectively operate, configure, maintain, and use the systems to the APEX team and users of the Crossroads and NERSC-9 systems. The APEX team may, at their option, make audio and video recordings of presentations from the Offeror's speakers at public events targeted at the APEX user communities (e.g., user training events, collaborative application events, best practices discussions, etc.). The Offeror will grant the APEX team user and distribution rights of documentation provided by the Offeror, session materials, and recorded media to be shared with other DOE Labs' staff and all authorized users and support staff for Crossroads and NERSC-9.

8.1 Documentation

8.1.1   The Offeror shall provide documentation for each delivered system describing the configuration, interconnect topology, labeling schema, hardware layout, etc. of the system as deployed before the commencement of system acceptance testing.

8.1.2   The Offeror shall supply and support system and user-level documentation for all components before the delivery of the system. Upon request by the laboratories, the Offeror shall supply additional documentation necessary for operation and maintenance of the system. All user-level documentation shall be publically available.

8.1.3   The Offeror shall distribute and update all documentation electronically and in a timely manner. For example, changes to the system shall be accompanied by relevant documentation. Documentation of changes and fixes may be distributed electronically in the form of release notes. Reference manuals may be updated later, but effort should be made to keep all documentation current.

8.2   Training

8.2.1   The Offeror shall provide the following types of training at facilities specified by ACES or NERSC:

	Number of Classes	
Class Type	ACES	NERSC
System Operations and Advanced Administration	2	2
User Programming	3	3

8.2.2   The Offeror shall describe all proposed training and documentation relevant to the proposed solutions utilizing the following methods:

- Classroom training
- Onsite training
- Online documentation
- Online training

9   References

APEX schedule and high-level information can be found at the primary APEX website http://apex.lanl.gov.

Crossroads/NERSC-9 benchmarks and workflows whitepaper can be found at the APEX Benchmark and Workflows website https://www.nersc.gov/research-and-development/apex/apex-benchmarks-and-workflows.

High Performance Computing Power Application Programming Interface
Specification http://powerapi.sandia.gov.

Appendix A: Sample Acceptance Plans

Appendix A-1: LANS Sample Acceptance Plan

Testing of the system shall proceed in three steps: pre-delivery, post-delivery and acceptance. Each step is intended to validate the system and feeds into subsequent activities.

Pre-delivery (Factory) Test

The Subcontractor shall demonstrate all hardware is fully functional prior to shipping. If the system is to be delivered in separate shipments, each shipment shall undergo pre-delivery testing. If the Subcontractor proposes a development system subcomponent, LANS recognizes that the development system is not part of the pre-delivery acceptance criteria.

LANS and Subcontractor staff shall perform pre-delivery testing at the factory on the hardware to be delivered. Any limitations for performing the pre-delivery testing need to be identified including scale and licensing limitations.

- Demonstrate RAS capabilities and robustness, using simple fault injection techniques such as disconnecting cables, powering down subsystems, or installing known bad parts.

- Demonstrate functional capabilities on each segment of the system built, including the capability to build applications, schedule jobs, and run them using the customer-provided testing framework. The root cause of any application failure must be identified.

- The Offeror shall provide a file system sufficiently provisioned to support the suite of tests.

- Provide onsite and remote access for LANS staff to monitor testing and analyze results.

- Instill confidence in the ability to conform to the statement of work.

Pre-Delivery Assembly

- The Subcontractor shall perform the pre-delivery test of Crossroads or agreed-upon sub-configurations of Crossroads at the Subcontractor's location prior to shipment. At its option, LANS may send a representative(s) to observe testing at the Subcontractor's facility. Work to be performed by the Subcontractor includes:
  - o All hardware installation and assembly
  - o Burn in of all components
  - o Installation of software

- o Implementation of the ACES-specific production system-configuration and programming environment

- o Perform tests and benchmarks to validate functionality, performance, reliability, and quality

- Run benchmarks and demonstrate that benchmarks meet performance commitments.

Pre-Delivery Configuration

- TBD

Pre-Delivery Test

Subcontractor shall provide LANS on-site access to the system in order to verify that the system demonstrates the ability to pass acceptance criteria.

The pre-delivery test shall consist of (but is not limited to) the following tests:

Name of Test	Pass Criteria
System power up	All nodes boot successfully
System power down	All nodes shut down
Unix commands	All UNIX/Linux and vendor specific commands function correctly
Monitoring	Monitoring software shows status for all nodes
Reset	"Reset" functions on all nodes
Power On/Off	Power cycle all components of the entire system from the console
Fail Over/Resilience	Demonstrate proper operation of all fail-over or resilience mechanisms
Full Configuration Test	Pre-delivery system can efficiently run applications that use the entire compute resource of the pre-delivery system. The applications to be run will be drawn from the 72-hour test runs, scaled to the pre-delivery configuration
Benchmarks	Benchmarks shall achieve performance within the limits of pre-delivery configuration

Name of Test	Pass Criteria
72 Hour test	100% availability of the pre-delivery system for a 72 hour test period while running an agreed-upon workload that exercises at least 99% of the compute resources

Post-delivery Integration and Test

Post-delivery Integration

During Post-Delivery Integration, the Subcontractor's system(s) shall be delivered, installed, fully integrated, and shall undergo Subcontractor stabilization processes. Post-delivery testing shall include replication of all of the pre-delivery testing steps, along with appropriate tests at scale, on the fully integrated platform. Where applicable, tests shall be run at full scale.

Site Integration

When the Subcontractor has declared the system to be stable, the Subcontractor shall make the system available to LANS personnel for site-specific integration and customization. Once the Subcontractor's system has undergone site-specific integration and customization, the acceptance test shall commence.

Acceptance Test

The Acceptance Test Period shall commence when the system has been delivered, physically installed, and undergone stabilization and site-specific integration and customization completed. The duration of the Acceptance Test period is defined in the Statement of Work.

All tests shall be performed on the initial production configuration as defined by LANS.

The Subcontractor shall supply source code used, compile scripts, output, and verification files for all tests run by the Subcontractor. All such provided materials become the property of LANS.

All tests shall be performed on the initial production configuration of the Crossroads system as it will be deployed to the ACES user community. LANS may run all or any portion of these tests at any time on the system to ensure the Subcontractor's compliance with the requirements set forth in this document.

The acceptance test shall consist of a Functionality Demonstration, a System Boot Test, a System Resilience Test, a Performance Test, and an Availability Test, performed in that order.

Functionality Demonstration

Subcontractor and LANS will perform the Functionality Demonstration on a dedicated system. The Functionality Demonstration shall show that the system is configured and functions in accordance with the statement of work. Demonstrations shall include, but are not limited to, the following:

- Remote monitoring, power control and boot capability

- Network connectivity

- File system functionality

- Batch system

- System management software

- Program building and debugging (e.g. compilers, linkers, libraries, etc.)

- Unix functions

System Boot Test

Subcontractor and LANS will perform the System Boot Test on a dedicated system. The System Boot Test shall show that the system is configured and functions in accordance with the statement of work. Demonstrations shall include, but are not limited to, the following:

Two successful system cold boots to production state, with no intervention to bring the system up. Production state is defined as running all system services required for production use and being able to compile and run parallel jobs on the full system. In a cold boot, all elements of the system (compute, login, I/O) are completely powered off before the boot sequence is initiated. All components are then powered on.

- Single node power-fail/reset test: Failure or reset of a single compute node shall not cause system-wide failure.

System Resilience Test

Subcontractor and LANS will perform the System Resilience Test on a dedicated system. The System Resilience Test shall show that the system is configured and functions in accordance with the statement of work.

All system resilience features of Crossroads shall be demonstrated via fault-injection tests when running test applications at scale. Fault injection operations should include both graceful and hard shutdowns of components. The metrics for resilience operations include correct operation, any loss of access or data, and time to complete the initial

recovery plus any time required to restore (fail-back) a normal operating mode for the failed components.

Performance Test

Crossroads system performance and benchmark tests are fully documented in the Statement of Work along with guidance and test information found at this website: https://www.nersc.gov/research-and-development/apex/apex-benchmarks-and-workflows.

The Subcontractor shall run the Crossroads tests and application benchmarks, full configuration test, external network test and file system metadata test as described in the Application and Benchmark Run Rules document. Benchmark answers must be correct, and each benchmark result must meet or exceed performance commitments in the performance requirements section.

Benchmarks must be run using the supplied resource management and scheduling software. Except as required by the run rules, benchmarks need not be run concurrently. If requested by LANS, Subcontractor shall reconfigure the resource management software to utilize only a subset of compute nodes, specified by LANS.

JMTTI and System Availability Testing

The JMTTI and System Availability Test will commence after successful completion of the Functionality Demonstration, System Test and Performance Test. LANS will perform the JMTTI and Availability Test.

The Crossroads system must demonstrate the JMTTI and availability metrics defined in the Statement of Work, within an agreed-upon period of time. An automated job launch and outcome analysis tool, such as the Pavilion HPC Testing Framework, shall be used to manage an agreed-upon workload that will be used to measure the reliability of individual jobs. These jobs shall be a mixture of benchmarks from the Performance Test and other applications.

Every test in the JMTTI and System Availability Test workload shall obtain a correct result in both dedicated and non-dedicated modes:

- In dedicated mode, each benchmark in the Performance Test shall meet the performance commitment specified in the Statement of Work. In non-dedicated mode, the mean performance of each performance test shall meet or exceed the performance commitment specified in the Statement of Work

- During the JMTTI and System Availability Test, LANS shall have full access to the system and shall monitor the system. LANS and

users designated by LANS shall submit jobs through the Crossroads resource management system.

- During the JMTTI and System Availability Test, the Subcontractor shall adhere to the following requirements:

  o All hardware and software shall be fully functional at the end of the JMTTI and Availability Test. Any down time required to repair failed hardware or software shall be considered an outage unless it can be repaired without impacting system availability.

  o Hardware and software upgrades shall not be permitted during the last 7 days of the JMTTI and Availability Test. The system shall be considered down for the time required to perform any upgrades, including rolling upgrades.

  o No significant (i.e. levels 1, 2 or 3) problems shall be open during the last 7 days.

- During the JMTTI and Availability Testing period, if any system software upgrade or significant hardware repairs are applied, the Subcontractor shall be required to run the Performance Tests and demonstrate that the changes incur no loss of performance. At its option, LANS may also run any deemed necessary. Time taken to run the Performance and other tests shall not count as downtime, provided that all tests perform to specifications.

Definitions for Node and System Failures

The baseline of interrupts, as used in the JMTTI and SMTBI calculations, shall include, but may not be limited to, the following circumstances:

- A node shall be defined as down if a hardware problem causes Subcontractor supplied software to crash or the node is unavailable. Failures that are transparent to Subcontractor-supplied software because of redundant hardware shall not be classified as a node being down as long as the failure does not impact node or system performance. Low severity software bugs and suggestions (e.g. wrong error message) associated with Subcontractor supplied software will not be classified as a node being down.

- A node shall be classified as down if a defect in the Subcontractor supplied software causes a node to be unavailable. Communication network failures external to the system, and user application program bugs that do not impact other users shall not constitute a node being down.

- Repeat failures within eight hours of the previous failure shall be counted as one continuous failure.

- The Subcontractor's system shall be classified as down (and all nodes shall be considered down) if any of the following requirements cannot be met ("system-wide failures"):

  o Complete a POSIX `stat' operation on any file within all Subcontractor-provided file systems and access all data blocks associated with these files.

  o Complete a successful interactive login to the Subcontractor's system. Failures in the ACES network do not constitute a system-wide failure.

  o Successfully run any part of the performance test. The Performance Test consists of the Crossroads Benchmarks, the Full Configuration Test and the External Network Test.

  o Full switch bandwidth is available. Failure of a switch adapter in a node does not constitute a system-wide failure. However, failure of a switch would constitute failure, even if alternate switch paths were available, because full bandwidth would not be available for multiple nodes.

  o User applications can be launched and/or completed via the scheduler.

- Other failures in Subcontractor supplied products and services that disrupt work on a significant portion of the nodes shall constitute a system-wide outage.

- If there is a system-wide outage, LANS shall turn over the system to the Subcontractor for service when the Subcontractor indicates they are ready to begin work on the system. All nodes are considered down during a system-wide outage.

- Downtime for any outage shall begin when LANS notifies the Subcontractor of a problem (e.g. an official problem report is opened) and, for system outages, when the system is made available to the Subcontractor. Downtime shall end when:

  o For problems that can be addressed by bringing up a spare node or by rebooting the down node, the downtime shall end when a spare node or the down node is available for production use.

  o For problems requiring the Subcontractor to repair a failed hardware component, the downtime shall end when the failed component is returned to LANS and available for production use.

For software downtime, the downtime shall end when the Subcontractor supplies a fix that rectifies the problem or when LANS reverts to a prior copy of the failing software that does not exhibit the same problem. A failure due to ACES or to other causes out of the Subcontractor's control shall not be counted against the Subcontractor unless the failure demonstrates a defect in the system. If there are any disagreements as to whether a failure is the fault of the Subcontractor or ACES, they shall be resolved prior to the end of the acceptance period.

Appendix A-2: NERSC Sample Acceptance Plan

Integration, Installation, and Acceptance Testing

Pre-delivery

> The Subcontractor shall demonstrate all hardware is fully functional prior to shipping. If the system is to be delivered in separate shipments, each shipment should undergo pre-delivery testing. If the Subcontractor proposes a development system subcomponent, the University recognizes that the development system is not part of the pre-delivery acceptance criteria. Deliverables of any NRE effort that are integrated into the build system should be considered part of the pre-delivery acceptance criteria.

Pre-Delivery Assembly

> The Subcontractor shall perform the pre-delivery test of the NERSC-9 system or agreed-upon sub-configurations of NERSC-9 at the Subcontractor's location prior to shipment. At its option, the University may send a representative(s) to observe testing at the Subcontractor's facility. Work to be performed by the Subcontractor includes:

- All hardware installation and assembly

- Burn in of all components

- Installation of software

- Successful integration of any NRE components into the build system.

- Implementation of the University-specific production system-configuration and programming environment

- Perform tests and benchmarks to validate functionality, performance, reliability, and quality

- Run benchmarks and demonstrate that benchmarks meet performance commitments

Pre-Delivery Test

> Subcontractor shall provide the University on-site access to the system in order to verify that the system demonstrates the ability to pass acceptance criteria.

> The pre-delivery test shall consist of (but is not limited to) the following tests:

Name of Test	Pass Criteria
System power up	All nodes boot successfully

Name of Test	Pass Criteria
System power down	All nodes shut down
Unix commands	All UNIX/Linux and vendor specific commands function correctly
Monitoring	Monitoring software shows status for all nodes
Reset	"Reset" functions on all nodes
Power On/Off	Power cycle all components of the entire system from the console
Fail Over/Resilience	Demonstrate proper operation of all fail-over or resilience mechanisms
Full Configuration Test	Full Configuration Test runs successfully on the system
Benchmarks	The system shall demonstrate the ability to achieve the required performance level on all benchmark requirements
72 Hour test	High availability of the production system for a 72 hour test period under constant throughput load

Post-delivery Integration and Test

▪ The Subcontractor's system(s) shall be delivered, installed, fully integrated, and shall undergo Subcontractor stabilization processes. Post-delivery testing shall include replication of all of the pre-delivery testing steps, along with appropriate tests at scale, on the fully integrated system.

Site Integration

When the Subcontractor has declared the system to be stable, the Subcontractor shall make the system available to University personnel for site-specific integration and customization. Once the Subcontractor's system has undergone site-specific integration and customization, the acceptance test shall commence.

Acceptance Test

The Acceptance Test Period shall commence when the system has been delivered, physically installed, and undergone stabilization and site-specific integration and customization. The duration of the Acceptance Test Period shall not exceed 60 days.

All tests shall be performed on the production configuration as defined by the University.

The Subcontractor shall not be responsible for failures to meet the performance metrics set or the availability metrics set forth in this Section, if such failure is the direct result of modifications made by the University to Subcontractor source code. Such suspension will be only for those requirements that fail due to the modification(s) and only for the length of time the modification(s) result(s) in the failure.

The Subcontractor shall supply source code used, compile scripts, output, and verification files for all tests. All such provided materials become the property of The University.

All tests shall be performed on a production configuration of the NERSC-9 system, as it will be deployed to the University user community. The University may run all or any portion of these tests at any time on the system to ensure the Subcontractor's compliance with the requirements set forth in this document.

The acceptance test shall consist of Functionality Demonstrations, System Tests, System Resiliency Tests, Performance Tests, and an Availability Test, performed in that order.

Functionality Demonstration

Subcontractor and the University will perform the Functionality Demonstration on a dedicated system. The Functionality Demonstration shall show that the system is configured and functions in accordance with the statement of work. Demonstrations shall include, but are not limited to, the following:

- Remote monitoring, power control and boot capability

- Network connectivity

- File system functionality

- Batch system

- System management software

- Program building and debugging (e.g. compilers, linkers, libraries, etc.)

- Unix functions

System Test

Subcontractor and the University will perform the System Test on a dedicated system. The System Test shall show that the system is configured and functions in accordance with the statement of work. Demonstrations shall include, but are not limited to, the following:

Two successful system cold boots to production state in accordance with required timings, with no intervention to bring the system up. Production state is defined as running all system services required for production use and being able to compile and run parallel jobs on the full system.

In a cold boot, all elements of the system (compute, login, I/O, network) are completely powered off before the boot sequence is initiated. All components are then powered on.

Single node power-fail/reset test: Failure or reset of a single compute node shall not cause a system-wide failure. A node shall reboot to production state after reset in accordance with required timings.

System Resilience Test

Subcontractor and the University will perform the System Resilience Test on a dedicated system. The System Resilience Test shall show that the system is configured and functions in accordance with the statement of work.

All system resilience features of the NERSC-9 system shall be demonstrated via fault-injection tests when running test applications at scale. Fault injection operations should include both graceful and hard shutdowns of components. The metrics for resilience operations include correct operation, any loss of access or data, and time to complete the initial recovery plus any time required to restore (fail-back) a normal operating mode for the failed components.

Performance Test

The Subcontractor shall run the NERSC-9 tests and application benchmarks, full configuration test, external network test and file system metadata test, a minimum of five times each as described in the Benchmark Run Rules section. Benchmark answers must be correct, and each benchmark result must meet or exceed performance commitments.

Benchmarks must be run using the supplied resource management and scheduling software. Except as required by the run rules, benchmarks need not be run concurrently. If requested by the

University, Subcontractor shall reconfigure the resource management software to utilize only a subset of compute nodes, specified by the University. Performance must be consistent from run to run.
Availability Test

The Availability Test will commence after successful completion of the Functionality Demonstration, System Test and Performance Test. The Subcontractor shall perform the Availability Test; at this time or before, the University will add user accounts to the system. The Availability Test shall be 30 contiguous days in a sliding window within the Acceptance Test Period. The NERSC-9 system must demonstrate the required availability of the system.

During the Availability Test, the University shall have full access to the system and shall monitor the system. The University and users designated by the University shall submit jobs through the NERSC-9 resource management system. These jobs shall be a mixture of benchmarks from the Performance Test and other applications.

The Subcontractor shall adhere to the System Availability and Reliability requirements as defined below:

- All hardware and software shall be fully functional at the end of the Availability Test. Any down time required to repair failed hardware or software shall be considered an outage unless it can be repaired without impacting system availability.

- Hardware and software upgrades shall not be permitted during the last 7 days of the Availability Test. The system shall be considered down for the time required to perform any upgrades, including rolling upgrades.

- No significant (i.e. levels 1, 2 or 3) problems shall be open during the last 7 days.

- During the Availability Testing period, if any system software upgrade or significant hardware repairs are applied, the Subcontractor shall be required to run the Benchmark Tests and demonstrate that the changes incur no loss of performance. At its option, the University may also run any deemed necessary. Time taken to run the Benchmark and other tests shall not count as downtime, provided that all tests perform to specifications.

- Every test in the Functionality Test, Performance Test and NERSC-defined workload shall obtain a correct result in both dedicated and non-dedicated modes.

- In dedicated mode, each benchmark in the Performance Test shall meet or exceed the performance commitment and variation requirement.

- In non-dedicated mode, the mean performance of each performance test shall meet or exceed the performance commitment. The measured Coefficient of Variation (standard deviation divided by the mean) of results from each performance test shall not be greater than 5%.

- Node and system availability will be measured on a node hour basis as follows.

$$SystemAvailability = \frac{\sum_i^N (S_i - D_i)}{\sum_i^N (S_i)}$$

where:

S_i is the number of scheduled hours for node i (wallclock time minus downtime scheduled by the University)

D_i is the number of hours of downtime for node i

Node and system outages are defined in the following section.

Definition of Node and System Failures

- A node shall be defined as down if a hardware problem causes Subcontractor supplied software to crash or the node is unavailable. Failures that are transparent to Subcontractor-supplied software because of redundant hardware shall not be classified as a node being down as long as the failure does not impact node or system performance. Low severity software bugs and suggestions (e.g. wrong error message) associated with Subcontractor supplied software will not be classified as a node being down.

- A node shall be classified as down if a defect in the Subcontractor supplied software causes a node to be unavailable. Communication network failures external to the system, and user application program bugs that do not impact other users shall not constitute a node being down.

- Repeat failures within eight hours of the previous failure shall be counted as one continuous failure.

- The Subcontractor's system shall be classified as down (and all nodes shall be considered down) if any of the following requirements can not be met ("system-wide failures"):

- Complete a POSIX 'stat' operation on any file within all Subcontractor-provided file systems and access all data blocks associated with these files.

- Complete a successful interactive login to the Subcontractor's system. Failures in the University network do not constitute a system-wide failure.

- Successfully run any part of the performance test. The Performance Test consists of the NERSC-9 Benchmarks, the Full Configuration Test and the External Network Test.

- Full switch bandwidth is available. Failure of a switch adapter in a node does not constitute a system-wide failure. However, failure of a switch would constitute failure, even if alternate switch paths were available, because full bandwidth would not be available for multiple nodes.

- User applications can be launched and/or completed via the scheduler.

- Other failures in Subcontractor supplied products and services that disrupt work on a significant portion of the nodes shall constitute a system-wide outage.

- If there is a system-wide outage, the University shall turn over the system to the Subcontractor for service when the Subcontractor indicates they are ready to begin work on the system. All nodes are considered down during a system-wide outage.

- Downtime for any outage shall begin when the University notifies the Subcontractor of a problem (e.g. an official problem report is opened) and, for system outages, when the system is made available to the Subcontractor. Downtime shall end when:

  - For problems that can be addressed by bringing up a spare node or by rebooting the down node, the downtime shall end when a spare node or the down node is available for production use.

  - For problems requiring the Subcontractor to repair a failed hardware component, the downtime shall end when the failed component is returned to the University and available for production use.

  - For software downtime, the downtime shall end when the Subcontractor supplies a fix that rectifies the problem or when the University reverts to a prior copy of the failing software that does not exhibit the same problem.

  - A failure due to the University or to other causes out of the Subcontractor's control shall not be counted against the Subcontractor unless the failure demonstrates a defect in the system. If there are disputes as to whether a failure is the fault of the Subcontractor or the University, they shall be resolved prior to the end of the acceptance period.

Appendix B: LANS/UC Specific Project Management Requirements

Appendix B-1: LANS Project Management Requirements

NOTE:

The following requirements apply to the project management of the delivery of the system proposed by the subcontractor. However, since there is a Non-Recurring Engineering (NRE) component to this Request for Proposal, NRE areas will also have similar project management requirements, should the proposals be negotiated into contracts. Key aspects will include the subcontractor NRE point of contact, NRE delivery milestone schedules, regular updates and reviews, and milestone approvals. The specific requirements for NRE project management will be negotiated by LANS and the selected subcontractor and will be based on the technical NRE areas proposed for evaluation.

Project Management

The development, pre-shipment testing, installation and acceptance testing of the Crossroads system is a complex endeavor and will require close cooperation between the Subcontractor, Los Alamos National Security, LLC (LANS), and ACES. There shall be quarterly executive reviews by corporate officers of the Subcontractor, ACES, and representatives of DOE/DP, to assess the progress of the project.

Project Planning Workshop

- LANS and Subcontractor shall schedule and complete a workshop to mutually understand and agree upon project management goals, techniques, and processes.

- The workshop shall take place no later than award + 45 days

Project Plan

- Delivery Milestone: no later than award + 60 days

Subcontractor shall provide the LANS with a detailed Project Plan – which includes a detailed Work Breakdown Structure (WBS). The Project Plan shall contain all aspects of the proposed Subcontractor's solution and associated engineering (hardware and software) and support activities.

The Project Plan shall address or include:

- Program Management

- High Assurance Delivery Process

    WBS:

- o Facilities Planning (e.g., floor, power & cooling, cabling);
- o Computer Hardware Planning;
- o Installation & Test Planning;
- o Deployment and Integration Milestones
- o System Stability Planning;
- o System Scalability Planning;
- o Software Plan
- o NRE deliverables
- o Testing (Build and NRE)
- o Development
- o Interdependencies between Build and NRE
- o Testing
- o Deployment
- o Risk Assessment & Risk Mitigation (Build and NRE)
- o Staffing;
- o On-site Warranty and Maintenance and Support Planning;
- o Training & Education;

Project Plan – Program Management

At a minimum, the Project Plan – Program Management Section shall:

- o Identify, by name, the Program Management Team members;
- o Identify, by name, the lead Crossroads System Architect
- o Identify, by name, the Crossroads System RAS Point of Contact
- o Describe the roles and responsibilities of the Team members;
- o List Subcontractor's Management Contacts;
- o Define and institutionalize the Periodic Progress Review process with regard to frequency (daily, weekly, monthly, quarterly, and annually), level (support, technical, and executive), and escalation procedures.

- • Additionally, the Project Plan – Program Management Section shall detail the joint activities of the Subcontractor and LANS to monitor and assess the overall Program Performance.

- LANS will furnish the Subcontractor with a top-10 list of problems and issues. The Subcontractor is responsible for appointing a point of contact for each of the items on the list. This list shall be reviewed weekly.

- All Subcontractor Program Management shall interface with the designated LANS Crossroads project manager.

- The WBS will be updated by the Subcontractor monthly and reviewed for approval by LANS

- The Subcontractor Project Plan shall be updated by the Subcontractor quarterly and reviewed for approval by LANS

Project Plan - High Assurance Hardware Delivery Process

Subcontractor shall provide the LANS with a high assurance delivery process and certification program for hardware deliverables of all stages of the deployment and operational use by the ASC Applications Community of the systems.

All assets delivered shall be, at a minimum, factory-tested and field–certified;

A "pre-delivery test" shall take place at the factory prior to each shipment. Functional diagnostics and agreed upon LANS applications shall be executed to verify the proper functioning of each system prior to shipment. Problems identified as a result of these tests shall be corrected prior to shipment. Assets that have successfully completed this pre-delivery test are "pre-verified."

Project Plan - High Assurance Software Delivery Process

Subcontractor shall provide LANS with a high assurance delivery process and certification program for software deliverables of all stages of the deployment and operational use by the ASCI Applications Community of the systems. In addition, Subcontractor shall provide LANS with documentation of Subcontractor's anticipated software release schedules during lifetime of the subcontract. This includes major and minor releases, updates, and fixes as well as expected beta-level availability.

- While Beta software and/or pre-GA software is anticipated to be installed and run on these systems, however all such installations are subject to LANS approval;

- Subcontractor shall provide LANS with a list of interdependencies between hardware and software as they pertain to the delivered systems;

Project Plan – WBS, Milestones

Subcontractor shall define appropriate high-level Milestones for the execution of the delivery and acceptance of the Crossroads system.

Project Plan – WBS, Facilities Planning

Compliant with the requirements of the Facilities described in the Technical Requirements.

Project Plan – WBS, System Stability Planning

Scalable systems of the size being delivered can at times prove difficult to predict in terms of stability. The number of components can have a significant effect on the stability and may provide some scalability problems in terms of stability of the system. The LANS requires a plan to progressively qualify a series of configurations of increasing complexity, in terms of both processor counts and interconnect topology.

Subcontractor shall be responsible for delivering a Stabilization Plan that includes the following:

- Plan objectives
- Target Goals for Stability, as agreed to jointly with the LANS
- Technical Strategy
- Roles and responsibilities
- Testing Plan
- Progress Evaluation Checkpoints
- Contingencies

Project Plan – Staffing:

- Staff Support shall be for the life of the subcontract.
- Subcontractor shall identify its members of the Project Team.

Project Plan – On-site Warranty and Maintenance and Support Planning

- On-site Warranty and Maintenance and Support shall be for the life of the subcontract
- On-site Warranty and Maintenance and Support shall include Subcontractor's preventive maintenance schedule.
- On-site Warranty and Maintenance and Support shall include logging and weekly reporting of all interruptions to service. At a minimum, the Subcontractor shall enter all interrupt logging into the LANS tracking system.

Project Plan – Training and Education

- In addition to Subcontractor's usual and customary customer Training and Education program, Subcontractor shall allow the LANS's staff access to Subcontractor's internal Training & Education program;

- Training and Education Support shall be for the life of the subcontract.

Project Plan – Risk Assessment and Risk Mitigation

- Subcontractor shall provide the LANS with a Risk Management Plan that identifies and addresses all identified risks.

- Subcontractor shall provide a risk management strategy for the proposed system in case of technology problems or scheduling delays that affect availability or achievement of performance targets in the proposed timeframe. Subcontractor shall describe the impact of substitute technologies on the overall architecture and performance of the system. In particular, the subcontractor shall address the technology areas listed below:

    o Processor

    o Memory

    o High-Speed Interconnect

    o Platform Storage and all other I/O subsystems

- Subcontractor shall continuously monitor and assess the risks involved for those major technology components that Subcontractor identifies to be on the Critical Path (i.e., Risk Assessment);

- Subcontractor shall provide the LANS with timely and regular updates regarding Subcontractor's Risk Assessment;

- Subcontractor shall provide the LANS with a Risk Mitigation Plan. Each risk mitigation strategy shall be subject to LANS approval. Such Risk Mitigation Plan shall include:

    o Risks Categorization – Risks shall be categorized according to

    o Probability of occurrence (Low, medium, or high)

    o Impact to the program if they occur (low, medium, or high)

    o Dates for Risk Mitigation Decision Points Identified

    o Execution of mitigation plans are subject to LANS approval and may include:

        ▪ *Technology Substitution* – subject to the condition that substituted technologies shall not have aggregate performance, capability, or capacity less than originally proposed;

        ▪ *3rd Party Assistance* – especially in areas of critical software development;

- *Source Code Availability* – especially in the areas of Operating Systems, Communication Libraries;

- *Performance Compensation* – possibility of compensating for performance shortfalls via additional deliveries.

o Subcontractor's Risk Mitigation Plan will be reviewed quarterly by the LANS.

Appendix B-2: UC Project Management Requirements

Project Management

The development, pre-shipment testing, installation and acceptance testing of the NERSC-9 system and the management of the Non-Recurring Engineering (NRE) subcontract(s) are complex endeavors and will require close cooperation between the Subcontractor and the Laboratory. There shall be quarterly executive reviews by corporate officers of the Subcontractor and UC to assess the progress of the project.

Project Planning Workshop

- LBNL and Subcontractor shall schedule and complete a workshop to mutually understand and agree upon project management goals, techniques, and processes.

- The workshop shall take place no later than 45 days after contract award

- The workshop shall address management goals, techniques and processes for the "Build" (NERSC-9) subcontract and the "NRE" subcontract.

Project Plan

- Subcontractor shall provide the University with detailed Project Plans – which include a detailed Work Breakdown Structure (WBS) for the "Build" and the "NRE" contracts. The Project Plans shall contain all aspects of the proposed Subcontractor's solution and associated engineering (hardware and software) and support activities.

- The Project Plans shall be submitted no later than 60 days after contract award

- The Project Plans shall address or include:
  - o Project Management
  - o Work Breakdown Structure for each of the projects
  - o Facilities Planning information (e.g., floor, power & cooling, cabling requirements) as applicable to the Build contract
  - o Computer Hardware Planning
  - o Installation & Test Planning (including pre-delivery factory tests and acceptance tests)
  - o Deployment and Integration
  - o System Stability Planning

- o System Scalability Planning

- o Software Plan

- o Development

- o NRE deliverables

- o Interdependencies between Build and NRE

- o Testing (Build and NRE)

- o Risk Assessment & Risk Mitigation (Build and NRE)

- o Staffing (for the life of the subcontracts)

- o On-site Support and Services Planning (for the life of the subcontracts)

- o Training & Education

Project Management Team

- The Subcontractor shall appoint a Project Manager (PM) for the purposes of executing the Project Management Plan for the "Build" system on behalf of the Subcontractor.   The PM for the ACES/Crossroads system and the NERSC-9 system shall be the same individual.

- The NRE contract(s) shall also have a Project Manager assigned to oversee the execution of the NRE contract on behalf of the Subcontractor. The PM for the ACES/Crossroads NRE and the NERSC-9 NRE shall be the same individual.

- The PMs for the system build and NRE subcontracts shall closely coordinate the projects. It is desireable that the same individual be the lead PM for all 4 subcontracts.

- The PMs shall be assigned for the duration of the subcontract. The PM for the "Build" system shall be based in the Bay area through the installation and acceptance of the delivered System. When the PMs are unavailable due to vacation, sick leave, or other absence, the Subcontractor shall provide backups who are knowledgeable of the NERSC-9 "Build" and "NRE" projects and have the authority to make decisions in the absence of the PM. The PMs or backups shall be available for emergency situations via phone on a 24x7 basis.

Subcontractor Management Contacts

The following positions in the Subcontractor management chain are responsible for performance under this subcontract:

- Technical Contact(s)

- Service Manager(s)

- Contract Manager(s)

- Account Manager(s)

Roles and Responsibilities for each of the PMs and management chain (Build and NRE)

The PM has responsibility for overall customer satisfaction and subcontract performance. It is anticipated that he/she shall be an experienced Subcontractor employee with working knowledge of the products and services proposed. The Subcontractor's PM can and shall:

- Delegate program authority and responsibility to Subcontractor personnel

- Establish internal schedules consistent with the subcontract schedule and respond appropriately to schedule redirection from the designated University authority

- Establish team communication procedures

- Conduct regularly scheduled review meetings

- Approve subcontract deliverables for submittal to the University

- Obtain required resources from the extensive capabilities available from within the Subcontractor and from outside sources

- Act as conduit of information and issues between the University and the Subcontractor

- Provide for timely resolution of problems

- Apprise the University of new hardware and software releases and patches within one week of release to the general market place and provide the University with said software within two weeks of request

The PM shall serve as the primary interface for the University into the Subcontractor, managing all aspects of the Subcontractor in response to the program requirements.

- The Technical Contacts shall be responsible for:

  o Developing (Build) System configurations to technical design requirements

  o Translation of NRE requirements to deliverables and tracking said deliverables

  o Updating the University on the Subcontractor's products and directions

- o Working with the respective PMs to review the Subcontractor's adherence to the Subcontracts

- The Contract Managers are:
  - o The Subcontractor's primary interface for subcontract matters
  - o Is authorized to sign subcontract documents committing the Subcontractor
  - o Supports the Project Manager by submitting formal proposals and accepting subcontract modifications.

- The Service Managers have the responsibility for:
  - o Compliance with the Subcontractor's hardware service requirements.
  - o Determining workload requirements and assigning services personnel to support the University
  - o Managing the Subcontractor's overall service delivery to the University
  - o Meeting with University personnel regularly to review whether the Subcontractor's service is filling the University's requirements
  - o Helping Subcontractor's service personnel understand University business needs and future directions

Periodic Progress Reviews

Daily Communication (Build Contract)

- For the Build contract, the Subcontractor's PM or designate shall communicate daily with the University's Technical Representatives or designate and appropriate University staff. These daily communications shall commence shortly after subcontract award and continue until both parties agree they are no longer needed. The topics covered in this meeting include:
  - o System problems – status including escalation
  - o Non-system problems
  - o Impending deliveries
  - o Other topics as appropriate

- The Subcontractor's PM (or designate) is the owner of this meeting. Target duration for this meeting is one-half hour. Both Subcontractor and the University may submit agenda items for this meeting.

Weekly Status Meeting (Build and NRE contracts)

- The Subcontractor's PM shall schedule this meeting. Target duration is one hour. Attendees normally include the Subcontractor's PM, Service Manager, University's Procurement Representative, Technical Representative and System Administrator(s) as well as other invitees. Topics covered in this meeting include:

  o Review of the past seven days and the next seven days with a focus on problems, resolutions, and impending milestones

  o Review of the University's top-10 list of problems and issues.

  o Specifically for the Build system

     ▪ System reliability

     ▪ System utilization

     ▪ System configuration changes

     ▪ Open issues (hardware/software) shall be presented by the Subcontractor's PM. Open issues that are not closed at this meeting shall have an action plan defined and agreed upon by both parties by close of this meeting

  o Specifically for the NRE contract(s)

     ▪ Progress towards deliverables

     ▪ Progress towards meeting technical milestones in the Build contract

     ▪ Implications of NRE deliverables for the Build system configuration

     ▪ Other topics as appropriate

Extended Status Review Meeting (Build and NRE contracts)

- Periodically, but no more than once per month and no less than once per quarter, an Extended Status Review Meeting will be conducted in lieu of the Weekly Status Meeting.

- A separate meeting for the NRE and Build contracts shall be conducted.

- The Subcontractor's PM shall schedule this meeting with the agreement of the University's Technical Representative. Target duration is one to three hours. Attendees normally include: Subcontractor's PM, Technical Contact, Field Service Manager and Line Management, University's Procurement Representative,

Technical Representative and Line Management as well as other invitees. Topics covered in this meeting include:

- Review of the past 30 days and the next 30 days with a focus on problems, resolutions and impending milestones (Subcontractor PM to present)
    - Deliverables schedule status (Subcontractor PM to present)
    - High priority issues (issue owners to present)
    - For the "Build" system: Facilities issues (changes in product power, cooling, and space estimates for the to be installed products)
    - All topics that are normally covered in the Weekly Status Meeting
    - Other topics as appropriate

Both Subcontractor and the University may submit agenda items for this meeting.

Quarterly Executive Meeting (Build and NRE contracts)

- Subcontractor's PM shall schedule this meeting. Target duration is six hours. Attendees normally include: Subcontractor's PM, Subcontractor's Senior Management, University's Procurement Representative, Technical Representative, selected Management, selected Technical Staff and other invitees. Topics covered in this meeting include:
    - Program status (Subcontractor to present)
    - University satisfaction (University to present)
    - Partnership issues and opportunities (joint discussion)
    - Future hardware and software product plans and potential impacts for the University
    - Participation by Subcontractor's suppliers as appropriate
    - Other topics as appropriate
    - Both Subcontractor and the University may submit agenda items for this meeting.
- The meeting will cover both NRE and Build contract issues.

Hardware and Software Support (Build Contract)

- Severity Classifications

o The Subcontractor shall have documented problem severity classifications. These severity classifications shall be provided to the University along with descriptions defining each classification.

- Severity Response

  o The Subcontractor shall have a documented response for each severity classification. The guidelines for how the Subcontractor will respond to each severity classification shall be provided to the University.

Problem Search Capabilities (Build and NRE contracts)

- The Subcontractor shall provide the capability of searching a problem database via a web page interface. This capability shall be made available to all individual University staff members designated by the University.

Problem Escalation (Build and NRE contracts)

- The Subcontractor shall utilize a problem escalation system that initiates escalation based either on time or the need for more technical support. Problem escalation procedures are the same for hardware and software problems. A problem is closed when all commitments have been met, the problem is resolved and the University is in agreement.

- As applicable to either contract, the University initiates problem notification to onsite Subcontractor personnel, or designated Subcontractor on-call staff.

Risk Management (Build and NRE contracts)

- The Subcontractor shall continuously monitor and assess risks affecting the successful completion of the NERSC-9 project (Build and NRE contracts), and provide the University with documentation to facilitate project management, and to assist the University in its risk management obligations to DOE.

- The Subcontractor shall provide the University with a Risk Management Plan (RMP) for the technological, schedule and business risks of the NERSC-9 project. The RMP describes the Subcontractor's approach to managing NERSC-9 project risks by identifying, analyzing, mitigating, contingency planning, tracking, and ultimately retiring project risks.

- The Plan shall address both the Build and the NRE portions of the project.

- The initial plan is due 30 days after award of the Subcontract. Once approved by the University, the University shall review the Subcontractor's RMP annually.

- The Subcontractor shall also maintain a formal Risk Register (RR) documenting all individual risk elements that may affect the successful completion of the NERSC-9 project (both Build and NRE contracts). The RR is a database managed using an application and format approved by the University.

- The initial RR is due 30 days after award of the Subcontract. The RR shall be updated at least monthly, and before any Critical Decision (CD) reviews with DOE. After acceptance, the RR shall be updated quarterly.

- Along with each required update to the RR, the Subcontractor shall provide a Risk Assessment Report (RAR) summarizing the status of the risks and any material changes. The initial report and subsequent updates will be reviewed and approved by the University's Technical Representative or his/her designee.

Risk Management Plan

- The purpose of this RMP, as detailed below, is to document, assess and manage Subcontract's risks affecting the NERSC-9 project:

  o Document procedures and methodology for identifying and analyzing known risks to the NERSC-9 project along with tactics and strategies to mitigate those risks.

  o Serve as a basis for identifying alternatives to achieving cost, schedule, and performance goals.

  o Assist in making informed decisions by providing risk-related information.

The RMP shall include, but is not limited to, the following components: management, hardware, software; risk assessment, mitigation and contingency plan(s) (fallback strategies).

Risk Register

- The RR shall include an assessment of each likely risk element that may impact the NERSC-9 project.   For each identified risk, the report shall include:

  o Root cause of identified risk

  o Probability of occurrence (low, medium, or high)

  o Impact to the project if the risk occurs (low, medium, or high)

- o Impact identifies the consequence of a risk event affecting cost, schedule, performance, and/or scope.
- o Risk mitigation steps to be taken to reduce likelihood of risk occurrence and/or steps to reduce impact of risk.
- Execution of mitigation plans are subject to University approval and may include:
  - o Technology substitution - subject to the condition that substituted technologies shall not have aggregate performance, capability, or capacity less than originally proposed;
  - o 3rd party assistance - especially in areas of critical software development;
  - o Performance compensation - possibility of compensating for performance shortfalls via additional deliveries.
  - o Dates for risk mitigation decision points.
  - o Contingency plans to be executed should risk occur; subject to University approval
  - o Owner of the risk.

Risk Assessment Report

- The RAR shall include the following:
  - o Total number of risks grouped by severity and project area (NRE and Build).
  - o Summary of newly identified risks from last reporting period.
  - o Summary of any risks retired since the last report.
  - o Identification and discussion of the status of the Top 10 (watch list) risks.

Definitions and Glossary

Baseline Memory: High performance memory technologies such as DDR-DRAM, HBM, and HMC, for example, that may be included in the systems memory capacity requirement. It does not include memory associated with caches.

Coefficient of Variation: The ratio of the standard deviation to the mean.

Delta-Ckpt: The time to checkpoint 80% of aggregate memory of the system to persistent storage. For example, if the aggregate memory of the compute partition is 3 PiB, Delta-Ckpt is the time to checkpoint 2.4 PiB. Rationale: This will provide a checkpoint efficiency of about 90% for full system jobs.

Ejection Bandwidth: Bandwidth leaving the node (i.e., NIC to router).

Full Scale: All of the compute nodes in the system. This may or may not include all available compute resources on a node, depending on the use case.

Idle Power: The projected power consumed on the system when the system is in an **Idle State**.

Idle State: A state when the system is prepared to but not currently executing jobs. There may be multiple idle states.

Injection Bandwidth: Bandwidth entering the node (i.e., router to NIC).

Job Interrupt: Any system event that causes a job to unintentionally terminate.

Job Mean Time to Interrupt (JMTTI): Average time between job interrupts over a given time interval on the full scale of the system. Automatic restarts do not mitigate a job interrupt for this metric.

JMTTI/Delta-Ckpt: Ratio of the JMTTI to Delta-Ckpt, which provides a measure of how much useful work can be achieved on the system.

Nominal Power: The projected power consumed on the system by the APEX workflows (e.g., a combination of the APEX benchmark codes running large problems on the entire system).

Peak Power: The projected power consumed by an application that utilizes the maximum achievable power consumption such as DGEMM.

Platform Storage: Any nonvolatile storage that is directly usable by the system, its system software, and applications. Examples would include disk drives, RAID devices, and solid state drives, no matter the method of attachment.

Rolling Upgrades/Rolling Rollbacks: A rolling upgrade or a rollback is defined as changing the operating software or firmware of a system component in such a way that the change does not require synchronization across the entire system. Rolling upgrades and rollbacks are designed to be performed with those parts of the system that are not being worked on remaining in full operational capacity.

System Interrupt: Any system event, or accumulation of system events over time, resulting in more than 1% of the compute resource being unavailable at any given time. Loss of access to any dependent subsystem (e.g., platform storage or service partition resource) will also incur a system interrupt.

System Mean Time Between Interrupt (SMTBI): Average time between system interrupts over a given time interval.

System Availability: ((time in period – time unavailable due to outages in period)/(time in period – time unavailable due to scheduled outages in period)) * 100

System Initialization: The time to bring 99% of the compute resource and 100% of any service resource to the point where a job can be successfully launched.

Wall Plate (Name Plate) Power: The maximum theoretical power the system could consume. This is a design limit, likely not achievable in operation.