September 3, 2024

RFP Technical Requirements Document

for

ATS-5 System

Version 1.0

LA-UR-23-30985

TABLE OF CONTENTS

REQUIREMENTS DEFINITIONS

Technical requirements have priority designations, which are defined as follows:

(a) Mandatory Requirements designated as (MR)

Mandatory Requirements (designated MR) are performance features that are essential to ATS-5 requirements, and an Offeror shall satisfactorily propose all Mandatory Requirements in order to have its proposal considered responsive.

(b) Mandatory Option Requirements designated as (MO)

Mandatory Option Requirements (designated MO) are features, components, performance characteristics, or upgrades whose availability as options are mandatory, and an Offeror shall satisfactorily propose all Mandatory Option Requirements in order to have its proposal considered responsive. The Laboratory may or may not elect to include such options in the resulting subcontract(s). Therefore, each MO shall appear as a separately identifiable item in Offeror's proposal.

(c) Target Requirements designated as (TR-1, TR-2, or TR-3)

Target Requirements (designated TR-1, TR-2, or TR-3) are features, components, performance characteristics, or other properties that are important to the Laboratory but will not result in a nonresponsive determination if omitted from a proposal. Target Requirements are prioritized by dash number. TR-1 is most desirable to the Laboratory and forms the baseline system, while TR-2 is more desirable and adds additional capabilities or increases productivity. TR-3s are stretch goals. Target Requirement responses will be considered as part of the proposal evaluation process.

(d) Technical Option Requirements designated as (TO-1, TO-2, or TO-3)

Technical Option Requirements (designated TO-1, TO-2, or TO-3) are features, components, performance characteristics, or upgrades that are important to the Laboratory but will not result in a nonresponsive determination if omitted from a proposal. Technical Options add value to a proposal. Technical Options are prioritized by dash number. TO-1 is most desirable to the Laboratory, while TO-2 is more desirable than TO-3. Technical Option responses will be considered as part of the proposal evaluation process; however, the Laboratory may or may not elect to include Technical Options in the resulting subcontract(s). Each proposed TO should appear as a separately identifiable item in an Offeror's proposal response.

Note: There are no mandatory requirements or mandatory options in the ATS-5 technical requirements document.

1.0 INTRODUCTION

The Department of Energy (DOE) National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) Program requires a computing system be deployed in 2027 to support the Stockpile Stewardship Program. In response to this requirement, Triad National Security, LLC (TNS) is releasing this Request for Proposal (RFP) for a next generation system, ATS-5.

The successful Offeror shall be responsible for delivering, installing, supporting, and maintaining the ATS-5 system.

The scope of work and technical specifications for any subcontracts resulting from this RFP will be negotiated based on this Technical Requirements Document and the Offeror's responses/proposed solutions.

ATS-5 has maximum funding limits over its system life to include all design and development, site preparation, maintenance, support, and analysts. Total Cost of Ownership (TCO) costs will be considered in system selection. The Offeror shall respond with a configuration and pricing for both the primary and alternate point designs.

Application performance and improved efficiency are essential to this procurement. Performance on the Linpack benchmark is not a metric of success. Success will be defined as meeting ATS-5 mission need (1.1) and, where necessary, achieving the architectural advancements (1.2) to meet these needs.

Supporting information can be found on the ATS-5 website: https://mission.lanl.gov/advanced-simulation-and-computing/platforms/ats-5/ .

Additional information on proposal preparation will be provided in the ***Proposal Submittal Requirements Section of the RFP***.

1.1 ATS-5 Mission Need

In the 2027 timeframe, Crossroads, the third ASC Advanced Technology System (ATS-3), will be nearing the end of its useful lifetime. The proposed ATS-5 system provides a replacement tri-lab computing resource for existing and new simulation codes and provides a resource for ever-increasing computing requirements to support the weapons program. The ATS-5 system, to be sited at Los Alamos, New Mexico, is projected to provide a large portion of the ATS resources for the NNSA ASC tri-lab simulation community of Los Alamos National Laboratory (LANL), Sandia National Laboratories (SNL), and Lawrence Livermore National Laboratory (LLNL) during the 2027–2031 timeframe. ATS-5 will support current and future simulation codes and will tackle some of the largest-scale 3D simulation workloads in support of the stockpile stewardship mission. These large-scale simulations are known as "hero"-class simulations, and ATS-5 will reduce these hero simulations' time-to-completion from months to days. Additionally, ATS-5 will provide the ability to run multiple of these hero-class simulations simultaneously. Vastly reduced time-to-solutions, combined with the ability to run multiple large-scale simulations, will dramatically improve NNSA's ability and agility to manage the stockpile.

The ASC Program is faced with significant challenges resulting from the ongoing technology revolution. The program must continue to meet mission needs while adapting to sometimes radical changes in technology. Codes running on NNSA Advanced Technology Systems (Crossroads and El Capitan) in the 2024–2027 timeframe are expected to run efficiently on ATS-5.

1.2 Architectural Advancements

Numerous architectural advancements in ATS-5 system design will support the NNSA Office of Defense Programs mission needs. ATS-5 needs to be designed with the following architectural advancements and major project goals:

- Overcoming the memory wall—continued memory bandwidth performance improvements for tri-lab applications
- Improved efficiency—programmer productivity, energy usage, and increased processor utilization
- Architectural diversity—ensuring that the high-performance computing ecosystem remains vibrant with multiple advanced technology solutions
- Time-to-solution—advancing strong scaling improvements to tackle the most pressing challenge of major improvements in time-to-solution for NNSA's largest and most complex stockpile simulations

Throughout this document, the term "architectural advancements" will refer to these four goals unless otherwise specified.

Crossroads[1] (ATS-3) will be used as the baseline for evaluating these goals. Continued memory bandwidth performance improvements are required for NNSA mission codes. Improving performance for sparse and semi-sparse memory accesses is of particular importance.

Improving efficiency includes programmer productivity, energy usage, and increased processor utilization. Programmer productivity is defined as the ease in which current NNSA mission codes can be ported to execute on the proposed architecture and the ease in which new codes can be written for the proposed architecture. Energy efficiency is defined as the application performance that is achieved as a function of the total dissipated power of the system. Increased processor utilization is defined as the percentage of available memory bandwidth and floating point operations that are utilized by the NNSA mission codes.

[1] Crossroads is a Cray HPE system with 6,144 compute nodes, each with dual socket Intel Xeon MAX CPUs (56 cores / socket) with 128 GBytes of HBM (per node). Nodes are interconnected with HPE Slingshot 11 network.

Architecture diversity is defined as the types of core technologies (Processor, Accelerators, Memory, Network) that are available to meet NNSA mission requirements. It is highly desirable to have a variety of technologies available across the tri-lab and a healthy vendor ecosystem to develop and deliver these technologies to the NNSA.

Time to insight is defined as the wall time it takes to complete NNSA's largest and most complex stockpile simulations. The ability to strong-scale these simulations to achieve significant reductions in total wall time while maintaining reasonable levels of processor utilization is of high value.

When evaluating proposals our goals of architecture advancements will be considered together.

1.3 Schedule

The following is the tentative schedule for the ATS-5 system.

Table 1. ATS-5 high-level schedule.

	ATS-5
RFP released	September 3, 2024
RFP responses due	45 days after release
Subcontracts (NRE/Build) awarded	Q3 CY 2025
Test or Early Access System	CY2026
Onsite system delivery	Q1 CY2027
Production	Q3 CY2027

1.4 Optional Advanced Technology Exploration System

In additional to selecting a successful offeror for the ATS-5 system, TNS, at its discretion, may also select a small-scale system based on the proposals for pursuing advanced technologies that are of interest to the future success of the overall NNSA computing strategy. The selection and award will be based on the technologies proposed and availability of funding beyond the ATS-5 funding profile.

2.0 HIGH-LEVEL SYSTEM REQUIREMENTS

This section describes the high-level technical requirements for the ATS-5 system proposals. In addition, the RFP provides system and node tables for the Offeror to complete and to submit separately as well as to include here.

In addition to the TRs identified in this document, the Offeror may choose to propose any additional features (i.e., Offeror-proposed features) consistent with the objectives of the ATS-5 procurement and the Offeror's roadmap, which the Offeror believes will be of value to TNS.

2.0.1 [TR-1] Each response/proposed solution within this document shall clearly describe the role of any lower-tier subcontractor(s) and the technology or technologies, both hardware and software, and value added that the lower-tier subcontractor(s) provide(s), where appropriate.

The potential improvement in performance of tri-lab workloads relative to the Crossroads supercomputer will be assessed in part using a scaled single node improvement (SSNI) metric defined in the Section 3, Benchmarks.

2.0.2 [TR-1] The Offeror should provide projections for SSNI of the ATS-5 system. Acceptance of the ATS-5 system will include a variety of criteria, including SSNI as well as a scalable system improvement (SSI) in which a subset of the benchmarks will be run multi-node. An SSI target of ten times (10x) over the Crossroads system will be negotiated prior to award and is not required in a response.

2.1 Architectural Description

2.1.1 [TR-1] The ATS-5 system will be sited at the LANL Strategic Computing Complex (SCC), Building 2327, on the LANL campus in Los Alamos, New Mexico. The Offeror should provide details of the physical footprint of the system and all of the supporting components to be sited at the LANL data center to meet the facility requirements in Section 9.

2.1.2 [TR-1] The Offeror should provide a detailed full system architectural description of the ATS-5 system that will **deliver at least a ten times (10x) SSNI over the Crossroads baseline benchmark results and at least 9 pebibytes (PiB) of compute partition memory**. The description should include diagrams and text describing the following details as they pertain to the Offeror's proposed system architecture(s) plus any unique features in the design. Include quantities and define any minimum scalable unit sizing to maintain optimal performance and productivity across the system.

- Component architecture—details of all processor(s), memory technologies, storage technologies, network interconnect(s), and any other applicable components. Details should include peak performance projections of each component and performance as a function of parallelism for each component.

- Compute 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. Details should be provided for each compute node type in the system. ATS-5 may contain multiple node types and include CPU-only or CPU+GPU (or other accelerator type) nodes, with the balance of node types to be determined in negotiations. SSNI shall be provided for homogenous systems and/or scalable units (SUs) of each node type proposed. A subset of the SSNI benchmarks may be used for different node types proposed as long as the full set of SSNI benchmarks is provided for at least the primary node type proposed for the ATS-5 solution.

- 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.

- Interconnect—details of the system's high speed network topology, network architecture, and connectivity across all system components (compute nodes, workflow environment nodes, platform storage, management system). Details shall include bandwidth and latency specifications between each computing element in the architecture.

- Storage systems—details of how the platform storage system is integrated with the ATS-5 system, including an architectural diagram and gateway nodes if applicable.

- 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 storage systems.

- Proposed floor plan—including details of the physical footprint of the system and all of the supporting components, including details of site and facility integration requirements (e.g., power, cooling, network).

- Management node(s)—details of hardware to support management and services to operate the ATS-5 system. Management node types can include, but are not limited to, master nodes for orchestration of system services, worker nodes for the deployment of services, Slurm resource manager, and storage nodes for system management and administration. Multiple node types may be needed to optimize for different uses described in Section 5.0.

- Workflow environment node(s) (WEN)—details of hardware to support user access and user-driven workflow activities. A pool of WENs will be needed to address the different requirements described in Section 4.0.

2.1.3 [TR-1] The Offeror should describe how the proposed architecture does or does not fit into the Offeror's long-term product roadmap and a potential follow-on system acquisition in the 2031 and beyond timeframe.

2.1.4 [TR-1] It is desirable to have the ability to intercept incorporate technologies that significantly improve application performance over the life of the ATS-5 system. The Offeror should provide descriptions, configurations, performance improvements for processor and accelerator technologies that may offer significant improvements in application performance but are not immediately available at the time of system delivery. It is highly desirable to have the flexibility to incorporate new processor, accelerator, and memory technologies during the lifetime of the ATS-5 system either as new compute partitions or integrated within the existing system. This flexibility should allow either high volume mass-market technologies or more targeted, tailored, and even specialized technologies that improve performance of one or more ATS-5 application targets. Technologies that further foster innovation of US-based technology providers are highly desirable. A subset of the SSNI benchmarks may be used for different node types proposed as long as the full set of SSNI benchmarks is provided for at least the primary node type proposed for the ATS-5 solution.

2.2 Software Description

2.2.1 [TR-1] The Offeror should provide a detailed description of the proposed system software and user programming environment, including a high-level software architecture diagram, the provenance of the software component (for example, open source or proprietary), support mechanism and licensing, if applicable (for the lifetime of the system including updates).

2.2.2 [TR-1] The Offeror should describe the high-level roadmap for the following:

- System software and tools provided for management and operation of the ATS-5 system.
- Provided user programming environment, including the ability to utilize new hardware features.
- How the software does or does not fit into the Offeror's long-term product roadmap and a potential follow-on system acquisition in the 2031 and beyond timeframe.

2.3 Non-Recurring Engineering (NRE)

TNS expects to award Non-Recurring Engineering (NRE) subcontracts, separate from the system build subcontract. It is expected that LANL personnel will collaborate in NRE subcontracts. It is anticipated that the NRE subcontracts could be approximately 10%–20% of the ATS-5 system budgets. The Offeror is encouraged to provide proposals for areas of collaboration they feel provide substantial value to the ATS-5 system. The goals of the NRE efforts should prioritize our required architectural advancements:

- Overcoming the memory wall—continued memory bandwidth performance improvements for tri-lab applications.
- Improved efficiency—programmer productivity, energy usage, and increased processor utilization.
- Architectural diversity—ensuring that the high-performance computing ecosystem remains vibrant with multiple advanced technology solutions.
- Time to insight—advancing strong scaling improvements to tackle the most pressing challenge of major improvements in time-to-solution for NNSA's largest and most complex stockpile simulations.

Advancements in overcoming the memory wall, improved performance of sparse memory access, and improved performance of branchy code is of particularly high value. Other advancements such as workflow enablement are of significantly lower priority.

Proposed NRE, which accelerates time-to-market of a solution, enhances features of a solution, or creates new solutions that then transition into planned roadmap activities are encouraged.

2.4 Upgrades, Expansions, and Additions

TNS 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 should propose separately priced options using whatever is the natural unit (such as SUs) for the proposed architecture design as determined by the Offeror. For example, for system size, the unit may be the number of racks or some other unit appropriate for incrementally increasing the system. The Offeror should identify any thresholds requiring increased component infrastructure (e.g., extra spine switches), any technical challenges foreseen with respect to scaling and any other production issues. Proposals should be as detailed as possible.

2.4.1 [TO-1] The Offeror should propose and separately price upgrades, expansions, or procurement of additional system configurations by the following fractions of the proposed system.

2.4.1.1 10%
2.4.1.2 25%
2.4.1.3 50%
2.4.1.4 100%
2.4.1.5 200%

2.4.2 [TO-1] The Offeror should propose upgrades, expansions, or procurement of additional Platform Storage System capacity in increments of 20% for the scalable units described in Section 7.

2.4.3 [TO-1] The Offeror should propose and price procurement of additional Application Regression Systems (ARS).

2.5 Early Access Systems

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 should propose options for Early Access Systems (EAS). The early access systems should contain similar functionality to the final system, including storage systems, management, and workflow environment nodes, but scaled down to the appropriate configuration.

2.5.1 [TO-1] The Offeror shall propose an EAS. 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 not to use the final processor, node, or high-speed interconnect architectures, but it is desirable for the EAS include any final architectures as close as possible. 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 ATS-5 system.

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

2.5.2 [TO-1] The Offeror shall propose development Test Bed EAS that will reduce risk and aid the development of any advanced functionality that is exercised as a part of the statement of work.

2.6 Test and Development Systems (TDS)

The test and development system (TDS) is a very small-scale version of the final system used for admins only. All patches/changes/etc. will be vetted on the TDS before pushing out to the ARS and ATS-5.

The TDS shall contain all the functionality of production ATS-5 systems, including storage systems, networks, all accelerator types, but scaled down to the appropriate configuration. It is desirable for ATS-5 production systems and TDS to be able to dynamically attach and detach from the same resources to allow scale testing on the test system by temporarily moving these resources from the production system to the test system. The Offeror should propose TDS for any production system delivered in support of the ATS-5 system. The TDS should be delivered before the production resource they are designed to support.

2.6.1 [TO-1] The Offeror shall propose a production TDS, which should contain at least 64 compute nodes.

2.7 Application Regression System (ARS)

The application regression system (ARS) is a smaller scale version the final ATS-5 that is used by developers and performance engineers to build, test, and optimize their codes. Users typically debug, optimize, and scale to modest node counts on the ARS. The ARS is also the first to get the patches/changes that have been vetted on the TDS prior to deployment on ATS-5.

The ARS shall contain all the functionality of production ATS-5 systems, including storage systems, networks, and all accelerator types, but scaled down to the appropriate configuration. It is desirable for ATS-5 production systems and ARS to be able to dynamically attach and detach from the same resources to allow scale testing on the test system by temporarily moving these resources from the production system to the test system. The Offeror should propose ARSs for any production system delivered in support of the ATS-5 system. The ARS should be delivered before the production resource they are designed to support.

2.7.1 [TO-1] The Offeror shall propose a production ARS, which should contain at least 200 compute nodes.

2.8 Risk Mitigation Design Point

[TO-1] The Offeror shall provide a summary of an alternate risk mitigation point design. The alternate point design shall be based on an architecture that reduces the risk of successful on-time deployment, for example, poses less schedule risk for delivery. It is of great importance that a viable platform (primary or alternate) is delivered in the ATS-5 timeframe capable of supporting mission needs regardless of unforeseen technology disruptions. The Offeror shall not submit a full alternative point design proposal. Instead, its summary of the alternate risk mitigation point design shall clearly describe any differences from the primary design point and how each of the noted differences satisfy the technical requirements contained in this document and reduces schedule risk for delivery of ATS-5. Multiple alternative risk mitigation point designs can be provided as long as one of the designs meets the requirements of 2.1.3.

3.0 BENCHMARKS

Assuring that real workflows perform well on the ATS-5 system is key to the success of the system. A suite of Benchmarks that are representative of the workloads of the NNSA laboratories has been developed. These Benchmarks, listed in Table 3-1, will be used to evaluate application performance as part of both the RFP response and system acceptance. The Benchmarks are supplemented by a collection of Micro-benchmarks listed in Table 3-2.

Final Benchmark acceptance performance targets for Scalable System Improvement (SSI) and SSNI will be negotiated prior to subcontract award. All performance tests must continue to meet acceptance criteria throughout the lifetime of the system.

The Benchmarks, Micro-benchmarks, and supplemental materials can be found on the ATS-5 benchmarks website: https://lanl.github.io/benchmarks/index.html. Benchmark results can be submitted for three categories of optimization (baseline, ported, and optimized), which are defined in the Run Rules.

SSNI is defined as follows:

Given two platforms using one as a reference, SSNI is defined as a weighted geometric mean using the following equation.

$$SSNI = N\left(\prod_{i=1}^{M}(S_i)^{w_i}\right)^{\frac{1}{\sum_{i=1}^{M} w_i}}$$

Where:

- N = Number of nodes on ATS-5 system / Number of nodes on reference system,
- M = total number of Benchmarks,
- S = application speedup (Figure of Merit on ATS-5 system / Figure of Merit on reference system Crossroads); S must be greater than 1, and
- w = weighting factor.

SSI is defined as follows:

Given two platforms using one as a reference, SSI is defined as a weighted geometric mean using the following equation.

$$SSI = \left(\prod_{i=1}^{M}(U_i S_i)^{w_i}\right)^{\frac{1}{\sum_{i=1}^{M} w_i}}$$

Where:

- M = total number of Benchmarks,
- S = application speedup (Figure of Merit on ATS-5 system / Figure of Merit on reference system); S must be greater than 1.
- U = utilization factor (n_ref / n) x (N / N_ref),

- n is the total number of nodes used for the application,
- N is the total number of nodes in the respective platform,
- ref refers to the reference system (Crossroads), and
- w = weighting factor.

Table 3-1: Benchmarks

Benchmark	Description	Language	Parallelism
Branson	Implicit Monte Carlo transport.	C++	MPI + CUDA/HIP
AMG2023	AMG solver of sparse matrices using Hypre.	C	MPI + CUDA/HIP/SYCL OpenMP on CPU
MiniEM	Electro-Magnetic solver.	C++	MPI + Kokkos
MLMD	ML Training of interatomic potential model using HIPYNN on VASP Simulation data. ML inference using LAMMPS, Kokkos, and HIPYNN trained interatomic potential model.	Python, C++, C	MPI + CUDA/HIP
Parthenon-VIBE	Block structured AMR proxy using the Parthenon framework.	C++	MPI + Kokkos
Sparta	Direct Simulation Monte Carlo.	C++	MPI + Kokkos
UMT	Deterministic (Sn) transport.	Fortran	MPI + OpenMP and OpenMP Offload

3.0.1 [TR-1] The Offeror should provide baseline or ported performance results for the proposed system and platform storage system for all the Benchmarks listed in Table 3-1. If baseline results cannot be obtained, ported results may be provided in their place.

3.0.2 [TR-2] The Offeror should provide ported results in addition to baseline results for the proposed system if minor code changes enable substantial performance gain.

3.0.3 [TR-2] The Offeror should provide optimized results for the proposed system to showcase system capabilities.

3.0.4 [TR-2] The Offeror should provide code changes for ported and/or optimized results.

3.0.5 [TR-1] The Offeror should state a minimum SSNI for the ATS-5 system relative to the baseline of the Crossroads system, to be measured using baseline results of the Benchmarks provided by LANL. If baseline results cannot be obtained, ported results may be provided in their place.

3.0.6 [TR-2] The Offeror should state a minimum SSNI for the ATS-5 system relative to the Crossroads system, to be measured using any combination of baseline, ported, or optimized versions of the Benchmarks.

3.0.7 [TR-1] The Offeror should provide results or estimates of strong scaling for CPU architectures and throughput curves for GPU architectures for problem sizes that meet percentage of memory requirements as defined in the run rules for each Benchmark for each compute node type proposed.

3.0.8 [TR-1] The Offeror should provide performance results for the proposed system for the Micro-benchmarks listed in Table 3-2. Some Micro-benchmarks can be run on multiple subdomains of the system; results should be provided for each configuration listed in the table.

3.0.9 [TR-2] The Offeror may use a subset of the SSNI benchmarks for different node types proposed as long as the full set of SSNI benchmarks is provided for at least the primary node type proposed for the ATS-5 solution. If a subset of benchmarks are used for any node types the weights in the SSNI calculation for those benchmarks shall be normalized to 1.

3.0.10 [TR-1] The Offeror should provide licenses for the delivered system for all software required to achieve Benchmark and Micro-benchmark performance, including, but not limited to, compilers and libraries.

Table 3-2: Micro-benchmarks

Micro-benchmark	Description	Language	Parallelism	Multi-Node
Stream	Streaming memory bandwidth test	C/Fortran	OpenMP	No
Spatter	Sparse memory bandwidth test driven by application memory access patterns.	C++	MPI + OpenMP/CUDA/OpenCL	No
OSU MPI + Sandia SMB Message Rate	MPI Performance Benchmarks	C++	MPI	Yes
DGEMM	Single node floating-point performance on matrix multiply.	C/Fortran	Various	No
IOR	Performance testing of parallel file system using various interfaces and access patterns.	C	MPI	Yes
mdtest	Metadata benchmark that performs open/stat/close operations on files and directories.	C	MPI	Yes

4.0 WORKFLOW ENVIRONMENT

4.1 Scalable and Reliable Workflow Services

[TR-1] The system should support running jobs up to the full scale of the compute node resources. The Offeror should describe factors (for example, executable size, number of shared libraries used) that may affect application launch time.

4.1.1 [TO-1] LANL anticipates that SchedMD's Slurm resource job management scheduler will be the primary scheduler and policy engine of the system. LANL will continue to evaluate alternative solutions such as the Flux resource manager (see https://computing.llnl.gov/projects/flux-building-framework-resource-management) and will remain open to such solutions. An alternative to Slurm can be proposed as a technical option, the Offeror should describe how the alternative solution meets the requirements.

4.1.2 [TR-1] LANL will directly procure the necessary software licenses and ongoing maintenance support from SchedMD. The Offeror will work with LANL and/or SchedMD to resolve operational problems with Slurm that may be caused by the Offeror's products. The Offeror will provide the necessary integration interfaces to support scalable job launch, including node placement, topology-aware scheduling, rank reordering, and node configuration and re-provisioning of nodes if supported by the hardware. The system design should not limit Slurm's ability to support thousands of concurrent users and more than 20,000 concurrent batch jobs.

4.1.3 [TR-2] The system should support a container orchestration platform such as Kubernetes or similar to provide staff-managed and user-generated and user-supported services on the workflow environment nodes. It should be capable of operating with Slurm to provide a unified workflow environment in which users can securely and performantly launch job tasks on the compute resources from the workflow environment nodes or services running on them (e.g., the ability for processes within a Kubernetes domain to be able to launch jobs to Slurm). The Offeror should describe any specialized hardware [or software] that may be required to support or enhance this unified workflow environment capability.

4.1.4 [TR-1] The system workflow environment nodes shall support interactive user access modes, including the following:

- command-line interface (CLI) through ssh and web-based user access modes for login, code compilation (cross-compilation is not desirable), application development, container builds, job lifecycle management, small-scale data analysis, and data transfer
- long-lived user services and frameworks (e.g., JupyterHub, databases, API services, and message brokers), that are staff-managed or user self-supported

The Offeror should describe mechanisms to enable these access models and how they are managed.

4.1.5 [TR-1] The system should provide correct numerical results. The Offeror should describe mechanisms and strategies for ensuring correct numerical results.

4.1.6 [TR-2] The Offerer should provide tests for correctness that can be used during the lifetime of ATS-5.

4.1.7 [TR-1] The Offeror should describe strategies for minimizing runtime variability in production and state a percent runtime variability that will not be exceeded. The Offerer should provide configuration recommendations (Slurm, MPI, etc.) to minimize runtime variability.

4.1.8 [TR-2] The Offeror should describe strategies for achieving no more than 10% runtime variability in production.

4.1.9 [TR-3] The Offeror should describe strategies for achieving no more than 4% runtime variability in production.

4.1.10 [TR-2] The scheduler should support job workflows with data stage-in and stage-out from local file systems and storage systems accessible only from a data transfer system.

4.2 Software Tools and Programming Environment

[TR-1] The system should support building and executing C17 code, C++20, C++23 code and Fortran 2008 + Further interoperability of Fortran with C - TS 29113 (henceforth referred to as Fortran2008+) code including code utilizing OpenMP directives 5.2 or latest. The Offeror should describe all supported compilers, including any enhancements or limitations that can be expected in meeting full support of the standards and other native language features for expressing parallelism including, but not limited to, support for C++ parallel STL, and Fortran do concurrent.

4.2.1 [TR-2] The system should support vendor-provided LLVM backends for each processing element (e.g., CPU, GPU, specialized accelerator) that can be utilized with both vendor-provided frontends and the open Clang/Flang projects.

4.2.2 [TR-2] The Offeror should describe the capability of the system to compile and run applications using Kokkos and RAJA. If present, describe any performance enhancements.

4.2.3 [TR-2] The Offeror should describe the capability of the system to compile and run applications using Legion and GASNet. If present, describe any performance enhancements.

4.2.4 [TR-3] The Offeror should describe the capability of the system to compile and run applications using SYCL and/or OpenACC 3.x. If present, describe any performance enhancements.

4.2.5 [TR-3] The Offeror should describe the capability of the system to compile and run CUDA based applications for GPU based solutions. Solutions that involve code transformations or source to source compilation should describe what functionality in CUDA will be supported.

4.2.6 [TR-1] The system should provide MPI libraries that support MPI 4.0 or higher and where applicable make GPU-aware MPI available wherever this is supported by a GPU vendor, and this shall be capable of running a job at full-system scale. The Offeror should describe any extensions or limitations to the MPI standard in the available MPI libraries.

4.2.7 [TR-1] The system shall support the Process Management Interface-X (PMI-X). The Offeror should describe the version and supported integrations.

4.2.8 [TR-2] The Offeror should describe any optimized BLAS, LAPACK, ScaLAPACK, and FFT libraries for CPUs, GPUs, and specialized accelerators.

4.2.9 [TR-3] The Offeror should describe any provided communication libraries (e.g., PGAS libraries and task-based programming libraries).

4.2.10 [TR-3] The Offeror should describe support for optimized scientific I/O libraries for CPUs, GPUs, and specialized accelerators (e.g., HDF5, NetCDF).

4.2.11 [TR-3] The Offeror should describe any limitations of using provided libraries (used to expose high-performance CPU, GPU, specialized accelerators, I/O, or communication capabilities in 4.2.5-4.2.9) through standard Python packages (e.g., NumPy, SciPy, h5py, mpi4py) or using standard Python packaging tools.

4.2.12 [TR-3] The Offeror should describe performance optimizations for distributed or parallel execution of Python programs through Python multiprocessing, Dask, Ray, or other similar non-MPI-based runtimes.

4.2.13 [TR-2] The Offeror should describe any provided optimized libraries for execution of machine learning and AI workloads for CPUs, GPUs, and specialized accelerators such as those required for optimal execution of deep learning frameworks like PyTorch and Tensorflow.

4.2.14 [TR-3] The Offeror should describe any support for distributed deep learning libraries (e.g., Horovod, PyTorch DDP) that enable scaling of training workloads across the full system and any provided tools that accelerate ML/AI/DL based workflows. (e.g., hyperparameter optimization, tracking experiments and integration with simulation and data pipelines).

4.2.15 [TR-2] The Offeror should describe, where applicable, support for direct data movement and access from the GPU to improve network bandwidth and latency (GPU-to-GPU) and I/O performance (GPU-to-Storage).

4.2.16 [TR-1] The Offeror should describe all provided profiling tools that include MPI and OpenMP profiling, and support for all user-accessible hardware (CPUs, GPUs, and specialized accelerators) and any provided compilers.

4.2.17 [TR-2] The Offeror should describe support for APIs, including Linux perf, that enable profilers and other performance optimization tools to access CPU, GPU, and specialized accelerator performance counters on the system. Include any restrictions on perf_event_paranoid, required kernel modules, or other security considerations.

4.2.18 [TR-1] The Offeror should describe all provided debugging tools for applications running on all user-accessible hardware such as gdb for CPUs and equivalents for GPUs and specialized accelerators.

4.2.19 [TR-2] The Offeror should provide a mechanism for users to build and run Open Container Initiative (OCI) compliant containers on the system without requiring privileged access to the system or allowing a user to escalate privilege.

4.2.20 [TR-1] The Offeror should describe how software and hardware dependencies, such as device driver libraries, MPI libraries, and libfabric/ucx/portals, can be accessed by containers, including dynamic mechanisms to maintain accessibility of these dependencies when software updates are made.

4.2.21 [TR-2] The Offeror should describe any provided container images for users (e.g., libraries, applications), the licensing model and how they can be distributed (e.g., can we distribute a container build on top of an Offeror-provided container), and the image registry where these container images may be published (e.g., an internal or public registry).

4.2.22 [TR-2] The Offeror shall describe any slowdowns and scaling limitations that would be observed due to running an application in a container up to the full scale of the system.

4.2.23 [TR-2] The Offeror should describe shared library support that prevents long-tail job launch times including but not limited to any restrictions on data access location, dedicated hardware requirements and/or any limitations of this solution.

4.2.24 [TR-1] The Offeror should describe an approach for a scalable mount solution of site-managed NFS including any limitations on supported versions or features.

4.3 Workflow Readiness Support

4.3.1 [TR-1] The Offeror should include in their proposal a separately priced plan to assist in transitioning select tri-lab applications to the system and shall propose a vehicle (e.g., a Center of Excellence [COE]) for supporting the successful execution of this plan. Support could be provided by the Offeror and processor(s) technology providers. The Offeror should provide access to experts in the areas of compilers and application performance in the form of staff training and deep-dive interactions with a set of teams. The deep-dive interactions should include, but not be limited to, support in building the application and its dependencies, profiling select problems, and advice on compiler options and optimizations.

4.3.2 [TR-2] The Offeror should include in their COE plan support for transitioning select workflows to the system. Support could be provided by the Offeror and/or key technology providers (e.g., the processor(s), storage, networking, third-party software) addressing overall workflow performance. The Offeror should include how they will collaborate with third-party developers from open source communities.

4.3.3 [TR-1] The Offeror should propose user training available during and outside of the COE for the lifetime of the system. Activities should target effective use of the user environment, performance, and optimization. The description should include topics, frequency, and format (e.g., classroom training or online training, hackathons).

4.4 Programming the Data Center

4.4.1 [TR-1] The system should support complex workflows through REST API interfaces or other mechanisms that expose functionality to users and automated services. The Offeror should describe the capabilities that their REST APIs expose, as well as how they are documented and tested. This description may include, but is not limited to, the following capabilities:

- System and subsystem status and health
- Data transfer, management, and archiving
- Orchestration of workflows, persistent services, CI/CD workflows (including container deployment), and complex science workflows
- Dynamic reconfiguration of storage, compute, and networking hardware
- Any authentication and authorization requirements/expectations for their REST APIs.

4.4.2 [TR-2] The Offeror should describe any capabilities that enable or improve multi-tenancy support (on compute and/or WEN nodes) that goes beyond the status quo (shared node jobs). This description may include, but is not limited to, the following capabilities:

- Protecting users from one another through minimized privileges and other mitigation techniques to prevent escalations in privileges.
- Virtualization and container networking (e.g., SR-IOV, VXLAN), including details of hardware offload capabilities, the number of tenants supported and guarantees of isolation between tenants.

4.4.3 [TR-3] The Offeror should describe capabilities to support "server-less" or Function-as-a-Service, including how it could integrate with the system and scheduler, security model, scaling, and performance. Describe any capabilities to support an event-based message model that can be used to publish and subscribe to system events, job events, data-related events, and other event types that can be integrated into and used to support complex workflows.

5.0 SYSTEM SOFTWARE & MANAGEMENT

5.0.1 [TR-1] The ATS-5 system should include management capabilities that facilitate integration with the evolving HPC environment. The management system should employ the following:

- Configuration management to ensure reproducibility and automation of critical tasks (e.g., continuous deployment of operating system images, container images, microservices on compute nodes, server nodes, any support devices, and reinstallation when necessary for operational reasons).
- Software components that should not restrict the evolution of the ATS-5 ecosystem, comply with open standards when available, and provide documented programming interfaces. For some components, LANL may choose to use open source or third-party software over the course of the production lifecycle.

5.0.2 [TR-1] The Offeror should provide an overview of applicable open standards and commodity open-source software that will fundamentally contribute to the solution.

5.0.3 [TR-1] For each open-source component of the management system, the Offeror should indicate the community governance model used.

5.0.4 [TR-1] The Offeror should provide a high-level overview of their proposed system management solution and any limitations toward achieving a modular environment.

5.1 Infrastructure Services

5.1.1. [TR-1] The Offeror should describe remote manageability capabilities of the compute nodes, network switches, fabric management, platform storage, power distribution units and servers comprising the system, including power control and console access, firmware updates, zero-touch provisioning, diagnostics, event logs, and alert capabilities. These capabilities should be accessible via documented APIs, preferably based on open standards, and a user interface.

5.1.2. [TR-1] The Offeror should describe any features provided for scalable full-platform management software that automates the management of all hardware, provides a comprehensive overview of system operations, and automates whole-system maintenance actions. Relevant features include, but are not limited to, sequenced power up and power down of the system; summarization of temperature, power, and other sensors; automating firmware and configuration updates; maintaining an inventory of field-replaceable units over the system lifetime; and collecting alert and error information from hardware.

5.2 Operating System

5.1.3. [TR-1] The system should 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.

5.1.4. [TR-1] The system should 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.

5.1.5. [TR-1] The Offeror shall describe the security capabilities of all operating systems proposed (e.g., compute, service).

5.1.6. [TR-1] The Offeror should enable all provided device drivers or kernel modules to be rebuildable and manageable.

5.1.7. [TR-1] The Offeror should provide access to source code and necessary build environment for all software except for firmware, compilers, and third-party products. The Offeror should provide updates of source code and any necessary build environment for all software over the life of the subcontract.

5.3 Platform Management

5.3.1. [TR-1] The Offeror should describe the system configuration management and diagnostic capabilities of the system that address the following details of system management:

- Any effect or overhead of software management tool components on the CPU or memory available on compute nodes.
- 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.

5.3.2. [TR-1] The Offeror should provide access to source code and necessary build environment for all software except for firmware, compilers, and third-party products. The Offeror should provide updates of source code and any necessary build environment for all software over the life of the subcontract.

5.3.3. [TR-1] The system should have no single points of failure that would cause a system outage (testing is described in Appendix A). The system and network should remain in an operational or degraded state after the unexpected failure of, or planned maintenance on, any single FRU, server, or switch and during any repair or other maintenance action. The Offeror should describe RAS capabilities to mitigate single points of failure (hardware or software) and the potential effect on running applications and system availability.

5.3.4. [TR-2] The Offeror should describe the resilience, reliability, and availability mechanisms and capabilities of the system to mitigate any condition or event that can potentially cause a job interrupt and how a job maintains its resource allocation and is able to relaunch an application after an interrupt.

5.4 System Software Deployment

5.4.1 [TR-1] The system should include the ability to perform rolling upgrades and rollbacks on a subset of the system while at least half of the system remains in production operation. The Offeror should describe the mechanisms and limitations of the continuous deployment framework.

5.4.2 [TR-2] The Offeror should describe the process for scalable boot, reconfiguring and rebooting of compute, server, and any other node types in the system. The description should include an overview of the node boot process (warmboot and coldboot), including secure boot, stateless/stateful node provisioning, and infrastructure automation for customization and configuration of a node, the coordination, ordering and parallelism of the boot process, and techniques to provide rapid configuration and rebooting. Include how the time required to reboot scales with the number of nodes being rebooted.

5.4.3 [TR-3] The Offeror should describe any suggested system development tools to make deployments easier (e.g., container registry, container image management, automated testing and version control) and describe how it integrates with the system management.

5.5 Data Collection and Monitoring

The ATS-5 system should include a mechanism to collect and provide metrics and logs that monitor the status, health, utilization, and performance of the system, subsystems, and all major components, including, but not limited to the following:

5.5.1 [TR-1] Environmental measurement capabilities for all systems and peripherals and their sub-systems and supporting infrastructure, including power and energy consumption and control.

5.5.2 [TR-2] Internal high speed network and management network performance counters, including measures of network congestion and network resource consumption.

5.5.3 [TR-3] Information enabling traffic and congestion attribution, with explanation of the attribution logic.

5.5.4 [TR-2] Hardware performance counters enabling application performance assessment with the ability to integrate these with system metric (e.g., network performance counters) data.

5.5.5 [TR-2] All levels of integrated and attached platform storage.

5.5.6 [TR-2] The system as a whole, including hardware performance counters for metrics for all levels of integrated and attached platform storage.

6.0 SYSTEM NETWORKS

6.0.1 [TR-1] The Offeror should propose a high-speed interconnect that will support a high messaging bandwidth, high injection rate, low latency, and high throughput.

6.0.2 [TR-1] The Offeror should provide a description of the system interconnect in detail, including topology, latencies, tail latencies, bandwidths, bi-section bandwidth, global bandwidth, message rates, routing algorithm, management and monitoring capabilities, and congestion mitigation.

6.0.3 [TR-1] The Offeror should describe the balance of per-node computational and memory bandwidth performance relative to that of the high-speed interconnect for each proposed node configuration.

6.0.4 [TR-1] The Offeror should describe link failure resilience throughout the network, including the number of links, network-interfaces, and switch failures that can occur while maintaining connectivity and how performance degrades as links fail.

6.0.5 [TR-1] The Offeror should describe support IPv4 and IPv6 for all proposed TCP/IP networks including high-performance network, management plane, external connectivity, and platform storage system.

6.1 High Performance Interconnect Hardware Requirements

6.1.1 [TR-1] The Offeror should 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.

6.1.2 [TR-1] The Offeror should describe how the interface will allow all processors in the system to simultaneously communicate synchronously or asynchronously with the high-speed interconnect.

6.1.3 [TR-1] The Offeror should describe how the interconnect will enable low-latency communication for one- and two-sided communication paradigms.

6.1.4 [TR-2] The Offeror should provide a detailed description of the interconnect's capabilities to accelerate MPI operations in hardware. This may include such features as collective offloads, hardware tag matching, hardware endpoint features, and "instant on" support (i.e., overhead for preparing network for use by a large-scale job).

6.1.5 [TR-1] The Offeror should report or project the proposed systems' node injection/ejection bandwidth.

6.1.6 [TR-1] The Offeror should report or project the proposed 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.

6.2 Communication/Computation Overlap

6.2.1 [TR-3] The Offeror should describe how both hardware and software components of the interconnect support effective computations and communication overlap for both point-to-point operations and collective operations through independent progress (i.e., the ability of the interconnect subsystem to progress outstanding communication requests in the background of the main computation thread).

6.3 Programming Model Requirements

6.3.1 [TR-1] The Offeror should provide the necessary system software to enable programming models including, but not limited to, MPI.

6.3.2 [TR-2] The Offeror should provide a lower-level communication API (LLCA) that supports a rich set of functionalities, including Remote Memory Access (RMA) and a Scalable Messaging Service (SMS), such as libfabric, UCX, or Portals, to meet this requirement.

6.3.3 [TR-3] The Offeror should provide a high performance and scalable TCP implementation.

6.3.4 [TR-2] The Offeror should provide a user-accessible mechanism to permit lower-level communication between processes that are launched independently, regardless of whether those processes are part of the same resource manager allocation or distinct resource manager allocations (i.e., to facilitate multi-executable and multi-job workflows).

6.3.5 [TR-2] If the system has multiple processor types (i.e., accelerators or coprocessors), each processor type should be able to initiate LLCA operations without explicit host activity and the LLCA will support RMA communication to each processor type's memory without explicit activity by the remote node host.

6.4 Quality of Service/Message Classes

6.4.1 [TR-3] The Offeror's interconnect should provide QoS capabilities (e.g., in the form of traffic classes/virtual channels or other sub-system QoS capabilities such as congestion control) including, but not limited to (1) the ability to prevent core communication traffic from interfering with other classes of communication such as debugging and performance tools or with I/O traffic, (2) additional virtual channels for efficient adaptive routing, and (3) prevention of different application traffic from interfering with each other (either through QoS capabilities or appropriate job partitioning).

6.4.2 [TR-3] The Offeror should provide an explanation of any sub-system QoS capabilities (e.g., platform storage QoS features).

6.5 Network Telemetry

6.5.1 [TR-3] The Offeror's interconnect should provide real-time telemetry (including network performance, congestion, and network resource consumption), FRU, and RAS metrics. The Offeror will describe which metrics are accessible and describe any provided software tools for collating and analyzing the data. Telemetry data will be user-level accessible.

6.5.2 [TR-3] The Offeror should provide access to the telemetry data for a job from the application processes to enable profiling support for the network. This support will include switch and/or router data, congestion state, throttling, bandwidth and throughput, and latency information for select packets traversing the network.

6.5.3 [TR-3] If the Offeror supports offloading of collectives to the network, the profiling capabilities should include the offloaded operations.

7.0 STORAGE SYSTEMS

7.1 Platform Storage System

The platform storage system (PSS) should accommodate the I/O and storage needs of multiple workload types. The PSS should be designed for scale, capable of serving I/O to whole-system jobs (as well as smaller jobs), maximize bandwidth, and have a data protection scheme in place (e.g., a form of RAID or erasure coding). The PSS will also be accessed via other systems, for example, external data transfer nodes, and as such should not require the compute system for it to be accessible and vice versa.

7.1.1 [TO-1] The Offeror should provide a separately priced option for the PSS.

7.1.2 [TR-2] The system should include platform storage capable of retaining all application input, output, and working data for 12 weeks (84 days), estimated at a minimum of 15% of baseline system memory per day.

7.1.3 [TR-1] The system should include platform storage with a warranted durability or a maintenance plan such that the platform storage is capable of absorbing approximately two times the systems baseline memory per day for a nominal five years.

7.1.4 [TR-1] The Offeror should describe how the system provides sufficient bandwidth to support a JMTTI/Delta-Ckpt ratio of greater than 200.

7.1.5 [TR-1] The Offeror should describe how the system satisfies a minimum storage bandwidth requirement capable of writing 25% of baseline system memory in less than 300 seconds.

7.1.6 [TR-1] The Offeror should describe the scalable unit for capacity, bandwidth, IOPS, and provide details for increasing each, and any associated limitations, including number of concurrent clients. Describe projected characteristics of primary storage devices such as media type, usable capacity, storage interfaces (e.g., NVMe, PCIe), and media durability.

7.1.7 [TR-1] The Offeror should describe all available interfaces to platform storage for the system, including but not limited to POSIX, Kubernetes CSI, and other APIs. Describe any exceptions to POSIX compliance, time to consistency, any potential delays for reliable data consumption.

7.1.8 [TR-1] The Offeror should propose a method(s) for PSS to be accessible externally and from login or interactive nodes even when the compute portion of the system is unavailable.

a. [TR-1] Describe any scenarios where a rebalance of resources is required after a reconfiguration and any scenarios where downtime for the entire PSS is required.

7.1.9 [TR-1] The Offeror should describe data protection schemes, mechanisms provided to ensure data integrity, reliability, and availability characteristics including any single points of failure, mechanisms for recovery, and related performance impact. Describe any anticipated loss of performance over time as the file system ages or reaches capacity.

7.1.10 [TR-1] The Offeror shall provide features to enforce and report upon soft (accounting) and hard (enforcement) quotas based on uid, gid, or other constructs.

7.1.11 [TR-1] The Offeror should provide system features for metadata scanning, metadata processing, and file/object purging across the entire PSS to allow for purging of older data and provide data to feed user data management tools. Describe the expected rate and elapsed time of a full-system scan and any performance impact while metadata scan or purge is happening.

7.1.12 [TR-2] The Offeror should describe any special capabilities that would mitigate user performance issues and/or allow the enumeration to complete in fewer than four hours; expect at least one billion objects. Describe the minimum resources required to achieve this target. Describe options for improving this scan performance.

7.1.13 [TR-2] The Offeror should describe the ability of the PSS to purge based upon individual files' last access time, modification time, owner, group, location, and size including capability to explicitly include or exclude files and directories, and a dry-run reporting mode.

7.1.14 [TR-1] The Offeror should 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.

7.1.15 [TR-3] The Offeror should describe their support of computational storage devices and how they are used in the storage tier (e.g., compression, erasure coding, data filtering, query pushdown support).

7.1.16 [TR-2] The Offeror should describe any capabilities to comprehensively collect platform storage usage data and describe specifically what data can be collected in-band versus out-of-band.

7.1.17 [TR-3] The Offeror should describe any features, alternative interfaces, and/or specialized hardware required for optimized performance of the following:

a. Traditional checkpoint/restart workloads (shared file and file-per-process)

b. Scientific data analysis workloads

c. Machine learning and inference workloads (both data parallel and model parallel strategies)

7.1.18 [TR-2] The Offeror shall describe available capabilities for storage server and media disaggregation including hardware, network, and software features to support such a design.

7.1.19 [TR-3] The Offeror shall describe all software translations required in the context of nvme-rdma, nfs-rdma, lnet, and any other supported efficient storage protocols. Please describe areas where TCP or other software is used to implement rdma read/write and explain expected impacts on performance.

8.0 SYSTEM OPERATION

8.1 Resilience, Reliability, and Availability Metrics

The ability to achieve the ATS-5 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 should provide a detailed description of how they arrived at their estimate(s).**

8.1.1 [TR-1] The system is available if it meets the system outage criteria in the glossary. The Offeror should propose a system availability.

8.1.2 [TR-1] The Offeror should propose a System Mean Time Between Interrupt (SMTBI).

8.1.3 [TR-2] The minimum System Mean Time Between Interrupt (SMTBI) should be greater than 720 hours.

8.1.4 [TR-1] The Offeror should propose a Job Mean Time To Interrupt (JMTTI) for a single job running on the entire system.

8.1.5 [TR-2] The minimum Job Mean Time To Interrupt (JMTTI) should be greater than 24 hours. Automatic restarts do not mitigate a job interrupt for this metric.

8.1.6 [TR-1] The system should complete power on and power off in a timely manner. The Offeror should describe the sequence of steps and timings for full system initialization and full system shutdown. Include any dependencies and how timings may scale with the size of the system.

8.1.7 [TR-2] A complete system initialization should take no more than 30 minutes.

8.2 System Security

8.2.1 [TR-1] The Offeror should describe the security capabilities of the proposed compute node and service partition operating systems.

8.2.2 [TR-1] The Offeror should describe how the system may be configured to support Zero-Trust requirements as described in the CISA Zero Trust Maturity Model (https://www.cisa.gov/zero-trust-maturity-model).

8.2.3 [TR-1] The Offeror should describe how their implementation of Identity Management (including Federated Identity) for the system works, what protocols and standards are being utilized (e.g., SpiFFE, SPIRE) and what system services and/or service accounts will use that framework.

8.2.4 [TR-1] The Offeror should describe how they will implement a Root of Trust to assure a boot environment for the system that is secure.

8.2.5 [TR-1] Security vulnerabilities in the software supplied by the vendor should be addressed with patches and/or updates in a timely manner depending on the classification and severity of the vulnerability. The Offeror should describe the process for handling security vulnerabilities and the time to provide patches from confirmed availability for

- Non-critical vulnerabilities,
- Known exploited vulnerabilities as defined by CISA, and
- Common Vulnerabilities and Exposures (CVEs) in the National Vulnerability Database (NVD) as defined by NIST with a score of Critical, High, or Medium as defined by the latest version of CVSS.

8.2.6 [TR-1] The Offeror should describe how the system validates its software and data components in compliance with the Presidential Executive Order on Software Bill of Materials (Cybersecurity Executive Order 14028.)

8.2.7 [TR-2] The Offeror should describe how the baseline configuration is tested, validated, and documented. Describe the process of re-validation, at any time, of a running system against the baseline, including how the system can be audited, so its current state can be captured and documented.

8.2.8 [TR-1] The Offeror should describe the security model, tools, and functionality for observability for any system applications and services that are containerized.

8.2.9 [TR-1] The Offeror should describe the available security controls and how they would be deployed and used for access methods to the system and/or services (e.g., interactive access using ssh, non-interactive methods using APIs or Web interfaces)

8.3 Power and Energy

8.3.1 [TR-1] The maximum power consumed by the system and its peripheral systems, including the proposed storage systems, will not exceed 20 megawatts. The maximum power consumption includes all equipment provided by the proposal. The Offeror should describe how its proposal fits within the power budget.

8.3.2 [TR-1] The Offeror should describe the power management capabilities of the system available to users or administrators.

8.3.3 [TR-3] The Offeror shall describe the overall system measurement capabilities for power, current, and voltage.

8.4 Maintenance and Support (Hardware/Software)

8.4.1 [TR-1] The Offeror should propose and separately price maintenance and support for all systems for a period of four years from the date of acceptance of the system in coordination with TNS and subject to review. The maintenance and support will include all features outlined in the Key Elements of the Maintenance and Support Plan described in Appendix B or describe how the proposed plan differs.

8.4.2 Maintenance and Support Solutions

[TR-1] The Offeror shall propose the following maintenance and support solutions and propose pricing separately for each solution. TNS 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.

Solution 1: 7x24
[TR-1] 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. Hardware service requests require the Offeror to be on site within four hours of the request.

Solution 2: 5x9
[TR-1] 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 TNS). The Offeror shall provide hardware maintenance training for TNS 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.

8.4.3 [TR-2] The system should include a means for tracking and analyzing all software updates, software and hardware failures, and hardware replacements over the lifetime of the system.

8.4.4 [TO-1] Offeror should also propose additional maintenance and support extension for years 5–8.

8.4.5 [TR-1] The Offeror should provide at least one Systems Operations and Advanced Administration training for each system delivered at facilities specified by TNS. The Offeror should describe additional training available for systems operations and advanced administration available for the lifetime of the system, including topics, duration, and proposed timing.

8.4.6 [TO-1] Offeror should provide pricing for additional training.

8.4.7 General Service Provisions

[TR-1] The Offeror shall be responsible, at its own expense, for the 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.

[TR-1] At its sole discretion, TNS may request advance replacement of components that show a pattern of failures that 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.

[TR-1] Hardware failures (damage to equipment) 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.

8.4.8 Software and Firmware Update Service

[TR-1] The Offeror shall provide an update service for all software and firmware provided for the duration of the Warranty plus Maintenance period. This service shall include new releases of software/firmware and software/firmware patches as required for normal use. The Offeror shall integrate software fixes, revisions, or upgraded versions in supplied software, including community software (e.g., Linux, Lustre), and make them available to TNS within twelve months of their general availability. The Offeror shall provide prompt availability of patches for cybersecurity defects.

8.4.9 Call Service

[TR-1] 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 TNS personnel. In the case of degraded performance, the Offeror's services shall be made readily available to develop strategies for improving performance (e.g., patches, workarounds).

8.4.10 Onsite Parts Cache

[TR-1] The Offeror shall maintain a parts cache on site at the TNS 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, 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 if the onsite parts cache proves to be insufficient to sustain the observed FRU or node failure rates.

8.4.11 Onsite Node Cache

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

8.4.12 Hardware Support Cluster

[TR-1] 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 or problem diagnosis. 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. TNS will store and inventory the HSC and other onsite parts cache components.

8.4.13 DOE Q-Cleared Technical Service Personnel

[TR-1] The ATS-5 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 third-party security escorts. The Offeror shall obtain necessary clearances for onsite support staff to perform their duties.

8.5 Onsite System and Application Software Analysts

[TO-1] The Offeror shall propose and separately price two System Software Analysts and two Applications Software Analysts. Offerors shall presume each analyst will be utilized for four years. For ATS-5, these positions require a DOE Q-clearance for access.

8.6 Deinstallation

[TO-1] The Offeror shall propose an option to deinstall, remove, and/or recycle the system and supporting infrastructure at end of life. Storage media shall be destroyed to the satisfaction of TNS and/or returned to TNS at its request.

9.0 FACILITIES AND SITE INTEGRATION

The following section addresses the facility-based requirements for the proposed system. It includes pertinent information and vendor requirements for the physical, electrical, cooling, seismic, safety, and transportation aspects of designing, delivering, installing, and integrating the system at the facility.

9.0.1 [TR-1] The system should use 3-phase Delta 480V AC (four-wire system, three phases and one ground) or 3-phase Wye 277/480V AC (five-wire system, three phases, one neutral, and one ground). 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 120/208V AC (third choice). The total number of individual branch circuits and phase load imbalance should be minimized.

9.0.2 [TR-1] All equipment and power control hardware of the system should be Nationally Recognized Testing Laboratories (NRTL) certified and bear appropriate NRTL labels.

9.0.3 [TR-1] Every rack, network switch, interconnect switch, node, and disk enclosure should be clearly labeled with a unique identifier visible from the front of the rack and 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.

9.0.4 [TR-1] All components in a rack intended to be serviced while the rack has power shall be fully serviceable without danger of touching an exposed conducting surface. Consider power switches for individual components that may need to be powered-off/-on individually. Consider minimizing the number of connecting cables that need to be removed to power-off/-on a component. Consider the placement of connectors, handles, etc., with respect to conducting services, that must be used to remove and replace a component.

9.0.5 [TR-1] Table 4 below shows target facility requirements identified by TNS for the ATS-5 system. The Offeror shall describe the features of its proposed systems relative to site integration at the respective facilities, including the following:

- Description of the physical packaging of the system, including dimensioned drawings of individual cabinet 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, and power-factor effects. 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 4: Target facility requirements

Location	Los Alamos National Laboratory, Los Alamos, New Mexico. The system will be housed in the Strategic Computing Complex (SCC), Building 2327
Elevation	7,500 feet

Seismic	N/A
Water Cooling	The system should operate in conformance with ASHRAE Class W2 guidelines (dated 2011). The facility will provide operating water temperature of 80.6°F, at up to 35PSI differential pressure at the system cabinets However, the Offeror should note if the system is capable of operating at higher temperatures. Note: LANL facility will provide inlet water at a nominal 80.6°F, per system design. 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 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 should have reliable leak detection, temperature, and flow alarms, with automatic protection and notification mechanisms.
Air Cooling	The system must operate with supply air at 75°F–65°F, with a relative humidity from 35%–60%. The rate of airflow is between 800–1400 CFM/floor tile. No more than 3MW of heat should be removed by air cooling.
Maximum Power Rate of Change	The hourly average in system power should not exceed the 2MW wide power band negotiated at least 2 hours in advance.
Power Quality	The system should be resilient to incoming power fluctuations at least to the level guaranteed by the ITIC power quality curve.
Floor	42-inch raised floor
Ceiling	16-foot ceiling and 16-foot ceiling plenum
Maximum Footprint	8,000 square feet; 80 feet long and 100 feet deep.
Shipment Dimensions and Weight	No restrictions.
Floor Loading	The average floor loading over the effective area should 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 point load applied on a one square inch area should not exceed 1,500 pounds. A dynamic load using a CISCA Wheel 1 size should not exceed 1,250 pounds (CISCA Wheel 2 – 1,000 pounds).
Cabling	All power cabling and water connections should 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) should be plenum rated and comply with NEC 300.22 and NEC 645.5. All communications cables, wherever installed, should be source/destination labeled at both ends. All communications cables and fibers over 10 meters in length and installed under the floor should 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.
External network interfaces supported by the site for connectivity requirements specified below	1 Gb, 10 Gb, 40 Gb, 100 Gb, 200 Gb, 400 Gb, IB EDR, IB HDR, IB NDR. The network infrastructure is continuously upgraded moving to the latest Ethernet and IB capabilities.
External bandwidth on/off the system for general TCP/IP connectivity	Minimum of 100 GB/s per direction with a preference for 300 GB/s per direction. Describe how 100 GB/s per direction could be expanded to 300 GB/s per direction.
External bandwidth on/off the system for accessing the system's PSS	Minimum of 100 GB/s with a preference for 300 GB/s. Describe how 100 GB/s could be expanded to 300 GB/s.
External bandwidth on/off the system for accessing external, site supplied file systems. E.g. Lustre, NFS	Minimum of 100 GB/s with a preference for 300 GB/s. Describe how 100 GB/s could be expanded to 300 GB/s.

10.0 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) are provided as part of the Request for Proposal.

10.1 Pre-delivery Testing

[TR-1] The TNS 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 should be identified in the Offeror's proposal, including scale and licensing limitations (if any). During pre-delivery testing, the Offeror must achieve the following:

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

10.1.2 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.

10.1.3 Provide a file system sufficiently provisioned to support the suite of tests.

10.1.4 Provide onsite and remote access to the TNS team to monitor testing and analyze results.

10.1.5 Instill confidence in the ability to conform to the Statement of Work.

10.2 Site Integration and Post-delivery Testing

10.2.1 [TR-1] TNS 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.

10.2.2 [TR-1] During post-delivery testing, the pre-delivery tests shall be run on the full system installation.

10.2.3 [TR-1] Where applicable, tests shall be run at full scale.

10.3 Acceptance Testing

10.3.1 [TR-1] TNS 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.

10.3.2 [TR-1] The Offeror shall demonstrate that the delivered system conforms to the subcontract's Statement of Work.

11.0 PROJECT AND RISK MANAGEMENT

The development, pre-shipment testing, installation, and acceptance testing of the ATS-5 system and the management of the Non-Recurring Engineering (NRE) subcontract are complex endeavors and will require close cooperation between the Offeror and TNS. The documents described in this section are not required in the RFP response; however, a commitment from the Offeror to deliver the documents within the timeframe described is required.

11.0.1 [TR-1] The Offeror shall Propose a risk management strategy documented in the Risk Management Plan 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

11.0.2 [TR-1] The Offeror should commit to develop, maintain, and submit to TNS the Planning Deliverables described in the **Key Project Planning Deliverables** section and form the relevant **Project Working Groups**.

11.0.3 [TR-1] The Offeror should provide in its RFP response a set of milestones in this section as described in the **Key Milestone Dates.**

11.0.4 [TR-1] LANL and Offeror shall schedule and complete a **Project Planning Kickoff Meeting** to mutually understand and agree upon project management goals, techniques, and processes for the ATS-5 system and the NRE subcontract. The kickoff meeting shall take place no later than 45 days after subcontract award.

Key Project Planning Deliverables

[TR-1] The Offeror will develop, deliver, submit for approval, and maintain the following Planning Deliverables.

These plans are described in more detail in Appendix B: Project and Risk Management—Key Planning Elements Descriptions. Initial versions and updates of these plans shall be provided in specific agreed upon time frames. Each of the plans and any revisions will be submitted for comment and approval to TNS project leadership.

- Project Plan
- Communication Plan
- Risk Management Plan
- Site Preparation Facilities Plan
- Chemical Management Plan
- Network Plan
- Early Access System Plan

- Acceptance Test Plan
- System Delivery and Installation Plan
- Training and Education Plan
- Center of Excellence for Workflow Readiness Plan
- Maintenance and Support Plan

Project Working Groups

As described above, the ATS-5 subcontracts must represent a partnership that is committed to delivering the most useful system possible. Upon subcontract award, the selected Offeror and TNS will assess the project for areas in which deep collaboration is necessary to ensure meeting that goal. The partnership will form working groups (WGs) for these topics. Each WG will interact in regard to all details of the technical topic; the selected Offeror will not attempt to limit the scope of these interactions. Specific details include NRE deliverables related to the technical topic and deployed software and hardware in the systems being built. The WGs will serve as a key conduit to identify, to refine and to understand LANL requirements in detail and to ensure that the delivered system meets those requirements to the greatest extent possible. WGs will establish a regular schedule for electronic meetings (e.g., telecons). Each Quarterly face-to-face meeting may include WG breakout sessions. WG breakout sessions will be determined by project management, including TNS and selected Offeror representatives. Project management will regularly assess WG progress and identify topics for which WGs are no longer required or additional topics for which new WGs are needed.

Key Milestone Dates

[TR-1] Offeror will provide TNS, in its proposed response, a proposed set of milestones for this section and, for each milestone, a proposed associated payment that is applicable to Offeror's proposed development and deployment timeline and methodology. Offeror is encouraged to identify milestones for each year of the project that merit revenue that the Offeror can legally recognize in that year.

The successful Offeror and TNS will hold a Technical Decision Point evaluation and joint planning meeting 9–12 months before System delivery. At the Technical Decision Point meeting, the final configuration of the System will be determined based on technology status and evaluations and component pricing. Performance targets will be re-evaluated and converted into requirements.

Prior to award, the Offeror and TNS will finalize the list of Key Milestone Dates, including dates for necessary Technical Decision Point evaluations. Following is a list of the kinds of key dates of importance to TNS. Other key dates may be needed for phased installations or deployments featuring major upgrades during the subcontract. *Early completion is highly desired*.

- Project Liaisons assigned to include the Offeror Project Manager, Executive Point of Contact, Service Manager, Contract Manager, and Account Manager
- Project Plan complete
- Technical Decision Point to exercise any proposed system options
- PSS late binding decision(s) to exercise any proposed tier options
- System Delivery and Installation Plan to include date of onsite support personnel to arrive on site (e.g., hardware, storage, software specialists)
- Begin delivery and installation of system and exercised storage and network options
- System Installation and Integration complete including PSS
- System Accepted including PSS

12.0 DOCUMENTATION AND TRAINING

12.1 Documentation

[TR-1] The Offeror should describe documentation provided to effectively administer and use the system, including types of documentation, format (e.g., user manuals, man pages, release notes, stable URLs, plain text vs pdf, vendor websites, any interactive elements), initial delivery, and frequency of updates, for the following:

- 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,
- Documentation of the proposed solution to the operators and system administrators to effectively operate and configure the platform, and
- Documentation for users describing the programming environment and software tools, including compatibility across system software updates.
- [TR-2] The Offeror should grant TNS use and distribution rights for provided documentation, training session materials and recorded media to be shared with DOE Lab staff and all authorized users and LANL support staff. TNS may, at their option, make audio and video recordings of presentations from Offeror at public events targeted at the LANL user communities (e.g., user training events, collaborative application events, Hackathons, Best Practices discussions) and make them available to all LANL users.
- [TR-1] All documentation shall be distributed and updated electronically and in a timely manner that maintains the productivity and performance of the system. For

example, changes to the system should be accompanied by relevant documentation, such as binary compatibility after major OS upgrades.

- [TR-1] 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.

12.2 Training

12.2.1 [TR-1] The Offeror shall provide the following types of training at facilities specified by TNS:

Class Type	Number of Classes
System Operations and Advanced Administration	2
User Programming	3
Networking and Storage	2

12.2.2 [TR-1] The Offeror shall describe all proposed training and documentation in the Training and Education plan relevant to the proposed solutions utilizing the following methods:

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

13.0 REFERENCES

Appendix A: ATS-5 Sample Acceptance Test 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, TNS recognizes that the development system is not part of the pre-delivery acceptance criteria.

TNS 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 TNS 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 ATS-5 or agreed-upon sub-configurations of ATS-5 at the Subcontractor's location prior to shipment. At its option, TNS may send a representative(s) to observe testing at the Subcontractor's facility. Work to be performed by the Subcontractor includes the following:

- All hardware installation and assembly
- Burn in of all components
- Installation of software
- Implementation of the TNS-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 Configuration

- TBD

Pre-Delivery Test

Subcontractor shall provide TNS onsite 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 shall be drawn from the 72-hour test runs, scaled to the pre-delivery configuration

Name of Test	Pass Criteria
Benchmarks	Benchmarks shall achieve performance within the limits of pre-delivery configuration
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 TNS 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 TNS.

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 TNS.

All tests shall be performed on the initial production configuration of the ATS-5 system as it will be deployed to the TNS user community. TNS 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 TNS 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)
- Unix functions

System Boot Test

Subcontractor and TNS 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 TNS 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 ATS-5 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

ATS-5 system performance and benchmark tests are fully documented in the Statement of Work along with guidance and test information found at the ATS-5 website: https://mission.lanl.gov/advanced-simulation-and-computing/platforms/ats-5/ .

The Subcontractor shall run the ATS-5 tests and 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 TNS, Subcontractor shall reconfigure the resource management software to utilize only a subset of compute nodes, specified by TNS.

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. TNS will perform the JMTTI and Availability Test.

The ATS-5 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, TNS shall have full access to the system and shall monitor the system. TNS and users designated by TNS shall submit jobs through the ATS-5 resource management system.
- During the JMTTI and System Availability Test, the Subcontractor shall adhere to the following requirements:
    - 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.
    - Hardware and software upgrades shall not be permitted during the last seven days of the JMTTI and 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 seven 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, TNS may also run any test 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 shall 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"):
  - 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 TNS network do not constitute a system-wide failure.
  - Successfully run any part of the performance test. The Performance Test consists of the ATS-5 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, TNS 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 TNS 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 TNS and available for production use.

For software downtime, the downtime shall end when the Subcontractor supplies a fix that rectifies the problem or when TNS reverts to a prior copy of the failing software that does not exhibit the same problem. A failure due to TNS 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 TNS, they shall be resolved prior to the end of the acceptance period.

Appendix B: Project and Risk Management—Key Planning Elements Descriptions

Each key planning element is described in detail below. The specific details are designed to help the successful Offeror successfully meet its commitment, to help TNS track the ATS-5 project, and to help TNS and the selected Subcontractor to understand and to mitigate risks successfully.

Key Elements of the Project Plan

The Offeror and TNS shall have a joint project planning call no more than 30 days after subcontract award. The Offeror should provide TNS with a detailed Project Plan no later than 90 days after subcontract award that addresses, at a minimum, the following:

- The Offeror should appoint a Program Manager (PM) for the purposes of executing the Project Plan for the system and NRE subcontract on behalf of the Offeror. The Program Manager shall serve as the primary interface for TNS, managing all aspects of the Subcontract in response to the program requirements.
- Project Management Organization Breakdown Structure (OBS) with management team's roles and responsibilities clearly defined
- Points of Contacts to include the Offeror's Technical Contact(s), Service Manager, Contract Manager, and Account Manager
- Work Breakdown Structure (WBS) to include all major subsystems, each software product, and each major equipment deliverable to TNS
- Full Project Schedule Gantt chart for the duration of the subcontract

The Project Plan should be updated as timelines for delivery and installation become firm.

Key Elements of the Communication Plan

The project planning kick-off meeting shall take place no later than 45 days after subcontract award. A Communication Plan shall be developed and shall describe the types of communications, meetings, and progress reviews as described below:

- Daily Communication (System subcontract)

  The Offeror's PM (or designate) is the owner of this meeting with a target duration one-half hour. Both Offeror and TNS may submit agenda items for this meeting. These daily communications shall commence 30 days before expected system delivery and continue until both parties agree they are no longer needed.

- Weekly Status Meeting (System and NRE subcontract)

The Offeror's PM shall schedule this meeting with a target duration of one hour. Attendees normally include the Offeror's PM, Service Manager, TNS's Procurement Representative, Technical Representative and System Administrator(s) as well as other invitees.

- Quarterly Business Reviews (System and NRE subcontract)

  The Offeror's PM shall schedule this meeting with a target duration of no less than six hours. Attendees normally include the following: Offeror's PM, Offeror's Senior Management, TNS's Procurement Representative, Technical Representative, selected Management, selected Technical Staff and other invitees. Topics covered will cover both NRE and System subcontract issues that will include the following:

  - Program status (Offeror to present)
  - TNS satisfaction (TNS to present)
  - Partnership issues and opportunities (joint discussion)
  - Future hardware and software product plans and potential impacts for TNS
  - Participation by Offeror's suppliers as appropriate
  - Other topics as appropriate
  - Both Offeror and TNS may submit agenda items for this meeting.

Key Elements of Risk Management (System and NRE subcontract)

The Offeror should provide TNS with a Risk Management Plan (RMP) for the technology, schedule and business risks of the ATS-5 project 30 days after award of the Subcontract. The RMP describes the Subcontractor's approach to managing ATS-5 project risks by identifying, analyzing, mitigating, contingency planning, tracking, and ultimately retiring project risks. The RMP must address both the System and the NRE portions of the project. The purpose of this RMP, as detailed below, is to document, assess, and manage Subcontract's risks affecting the ATS-5 project:

- Document procedures and methodology for identifying and analyzing known risks to the ATS-5 project along with tactics and strategies to mitigate those risks.

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

- 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). Once the plan is approved by TNS, TNS shall review the Offeror's RMP annually.

The Offeror should also maintain a formal Risk Register (RR) documenting all individual risk elements that may affect the successful completion of the ATS-5 project (both System and NRE subcontract). The RR is a database managed using an application and format approved by TNS. 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. Items in the RR include, at a minimum, mitigation strategies, impact to the ATS-5 project, severity rating, and probability of the risk occurring.

Key Elements of Site Preparation Facilities Plan

Site planning shall be conducted by the Offeror's Site Engineering department. Site planning consists of the exchange of system specification documents, site floor plans, and is followed by a physical inspection of facilities. The plan shall include the floor plan and the Machine Unit Specification (MUS). The Offeror should provide Preliminary Site Preparation Facilities Plan at least one year prior to the delivery of the first equipment. At least nine months prior to the first equipment delivery, the Final Site Preparation Plan will be delivered to the Laboratory for approval. Details of the Facility requirements are in Section 9.0.

Key Elements of the Chemical Management Plan

The Offeror should develop and document a chemical management plan covering all aspects of transport, storage, filling, draining, and disposal of the chemicals, MSDS safety compliance, and required training and personal protective equipment needed for compliance with safety regulations, as required. The plan shall be documented and submitted for TNS review and concurrence prior to the delivery and deployment of the system.

Key Elements of the Network Plan

No less than six months prior to installation of the System, the Offeror should provide an initial draft of the system network configurations for TNS to review, including all network types provided, and showing compute-to-compute, compute-to-storage, and system-to-external components connectivity.

No less than four months prior to installation, TNS and the Offeror should finalize the network design and the Offeror should provide an up-to-date copy of the network configurations reflecting the expected-at-installation design.

Key Elements of the Acceptance Test Plan

TNS and the Offeror will create a detailed Acceptance Test Plan one year prior to the first equipment delivery and will be updated as necessary.

Items in the plan should, at a minimum, include the following:

- Pre-Delivery Assembly, Quality Assurance, and Preliminary Factory Testing. The plans will include how the Offeror qualifies their vendors, factory burn in and validation test plans and a pre-ship test plan for the ATS-5 system.
- Acceptance Testing. The plans shall consist of Functionality Demonstrations, System Tests, System Resilience Tests, Performance Tests, and an Availability Test, performed in that order.
- EAS, TDS, and ARS Testing. Details for testing systems in support of the ATS-5 production system.

An example Acceptance Test Plan is included in Appendix A.

Key Elements of System Delivery and Installation Plan

The Offeror should provide a Preliminary Delivery and Installation Plan to TNS one year prior to the first equipment delivery and will be updated as necessary. The Offeror will provide a Final Installation Plan no less than 90 days before delivery. The plan shall include the following:

- Core installation team and staffing plan
- Detailed delivery and installation schedule
- Equipment layout and installation sequence (multi-stage deliveries)
- Detailed integration and test plan addressing all equipment and software included in the delivery
- Safety documents required by LANL safety processes
- Diagrams showing the internal layout of cabinets

TNS shall review the plan and work with the Subcontractor to promptly resolve any issues or clarifications.

Key Elements of Maintenance and Support Plan

The Offeror should provide a maintenance, onsite support, and services plan for the life of the subcontract and shall include the following features:

- **Maintenance and Support Period**
  The Offeror should propose all maintenance and support for a period of four years from the date of acceptance of the system. Warranty shall be included in the four years. For example, if the system is accepted on April 1, 2027, and the Warranty is for one year, then the Warranty ends on March 30, 2028, and the maintenance period begins April 1, 2028, and ends on March 30, 2031. Offeror should also propose additional maintenance and support extension for years 5–7.
- **Maintenance and Support Solution**
  The Offeror should propose a maintenance and support solution with 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. The Offeror should provide/enable access to direct communication between LANL staff and any of their technology (cpu, gpu, network, software, etc.) providers technical staff.
- **Concurrent Maintenance Techniques**
  The Offeror should use continuous operations maintenance techniques (e.g., warm

swap) that avoid service disruptions. Continuous operations comprise both hardware (including servicing node hardware, cabinet hardware), and software upgrades to systems management nodes, workflow nodes, storage, and compute nodes. These actions shall not be deemed to cause a system outage if performed with the concurrence of TNS and completed in a timely manner. Six hours are permitted for cabinet-level repairs and two hours for all other repairs performed concurrently, node downtime due to concurrent maintenance is counted in calculating System availability.

- **General Service Provisions**
  The Offeror should 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, LANL 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.

  When a component has failed in service, the Offeror should replace the component with a newly manufactured or remanufactured/fully tested component. The Offeror should not place a component back into the main system in order to determine if a failure is transient. With TNS's concurrence, the Offeror may use the test system to test components.

- **Software and Firmware Update Service**
  The Offeror should 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 normal use. The Successful Offeror should integrate software fixes, revisions or upgraded versions in supplied software, including community software (e.g., Linux, Lustre), and make them available to LANL within 12 months of their general availability. The Offeror should provide prompt availability of patches for cybersecurity defects.

- **Call Service**
  The Offeror should 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 LANL 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).

- **Problem Escalation**
  The Offeror should document severity classifications and response for hardware and software problems. The description should include the technical problem escalation mechanism based either on time or the need for more technical support in the event issues are not being addressed to TNS's satisfaction. Problem escalation procedures

are the same for hardware and software problems. Problems should be searchable in a database and made accessible via a web interface or for download in a standard format (e.g., csv). This capability shall be made available to all individual TNS staff members designated by TNS.

- **Onsite Parts Cache**
  The Offeror should maintain a parts cache on site at LANL. 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 onsite parts cache prove to be insufficient to sustain the observed FRU or node failure rates.

- **Onsite Node Cache**
  The Offeror should also maintain an onsite spare node inventory of at least 1% of the total nodes for all of the onsite systems. These nodes shall be maintained and tested for hardware integrity and functionality utilizing the Hardware Support Cluster defined below if provided.

- **Hardware Support Cluster**
  A test and development system (TDS) described in Section 2.6 will be used as 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 should 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.

Key Elements of the Training and Education Plan

The Offeror should provide a Training and Education Plan within 120 days of the subcontract award. These plans will be updated throughout the life of the project to reflect the latest content, as needed. The plans will include the following:

- Description of available training activities that target effective use of the user environment, performance, and optimization. The description should include topics, frequency, format (such as classroom training or online training, in-person hackathons, and remote mini-hackathons on request by developers), and pricing.

- Description of collaboration with CPU and GPU vendors, other key technology providers, and LANL staff where appropriate for the proposed training activities.

Key Elements of the Center of Excellence for Workflow Readiness Plan

The Offeror should provide a Center of Excellence for Workflow Readiness Plan to assist in transitioning select LANL mission workflows to the system within 120 days of the subcontract award.

The plan will be updated quarterly throughout the life of the project as needed. The plan will include the following:

- Named support staff provided by the Offeror and the CPU and GPU vendor.
- Staff training and deep-dive interactions with a set of teams. The deep-dive interactions should include concrete deliverables such as building the code and its dependencies and profiling a selected problem.
- How application and workflow developers can begin porting and optimization activities using proposed early access systems described in Section 2 and how this content will be documented and shared with developers.
- Mutually agreed upon duration and level of effort. For example, at least 1.0 FTE equivalent support should be provided from the date of subcontract execution through two years after final acceptance of the system.

DEFINITIONS AND GLOSSARY

Baseline Memory: High performance memory technologies such as DDR-DRAM and HBM, 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.

Coldboot: Full power-on of a system from a non-energized state, such as a post–power outage situation. It can be assumed that facilities power, water, network, and site infrastructure services have been returned to service and timings for all offered file systems and cluster can begin from this point.

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 9 PiB, Delta-Ckpt is the time to checkpoint 7.2 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.

Nameplate Power: The maximum theoretical power the system could consume. This is a design limit, likely not achievable in operation, commonly specified on electrical equipment labels and used for power provisioning design per National Electrical Code (NEC, NFPA 70).

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

Operational Capability: Real, usable capabilities in production operation, not theoretical capabilities.

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.

Warmboot: The cluster/file system management servers being booted and configured.